

---

# DynGFN: Towards Bayesian Inference of Gene Regulatory Networks with GFlowNets

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1       One of the grand challenges of cell biology is inferring the gene regulatory network  
2       (GRN) which describes interactions between genes and their products that control  
3       gene expression and cellular function. We can treat this as a causal discovery  
4       problem but with two non-standard challenges: (1) regulatory networks are inher-  
5       ently cyclic so we should not model a GRN as a directed acyclic graph (DAG),  
6       and (2) observations have significant measurement noise, so for typical sample  
7       sizes there will always be a large equivalence class of graphs that are likely given  
8       the data, and we want methods that capture this uncertainty. Existing methods  
9       either focus on challenge (1), identifying *cyclic* structure from dynamics, or on  
10      challenge (2) learning complex Bayesian *posteriors* over DAGs, but not both. In  
11      this paper we leverage the fact that it is possible to estimate the “velocity” of gene  
12      expression with *RNA velocity* techniques to develop an approach that addresses  
13      both challenges. Because we have access to velocity information, we can treat the  
14      Bayesian structure learning problem as a problem of sparse identification of a dy-  
15      namical system, capturing cyclic feedback loops through time. Since our objective  
16      is to model uncertainty over a discrete structures, we leverage Generative Flow  
17      Networks (GFlowNets) to estimate the posterior distribution over the combinatorial  
18      space of possible sparse dependencies. Our results indicate that our method learns  
19      posteriors that better encapsulate the distributions of cyclic structures compared to  
20      counterpart state-of-the-art Bayesian structure learning approaches.

## 21   1 Introduction

22    Inferring gene regulatory networks (GRNs) is a long standing problem in cell biology [25, 44]. If  
23    we knew the GRN, it would dramatically simplify the design of biological experiments and the  
24    development of drugs because it would serve as a map of which genes to perturb to manipulate  
25    protein and gene expression. GRNs concisely represent the complex system of directed interactions  
26    between genes and their regulatory products that govern cellular function through control of RNA  
27    (gene) expression and protein concentration. We can treat GRN inference as a causal discovery  
28    problem by treating the regulatory structure between genes (variables) as causal dependencies (edges)  
29    that we infer / rule out by using gene expression data. Structure learning methods aim to automate  
30    this task by inferring a set of directed acyclic graphs (DAGs) that are consistent with the conditional  
31    independencies that we can measure among the variables [13, 41, 42]. While there may be multiple  
32    DAGs in this set—the “Markov equivalence class”—when we are able to perturb the variables with  
33    enough experimental interventions, it is possible to uniquely identify a causal graph [17].

34    However, structure learning for inferring GRNs comes with two non-standard challenges: (1) gene  
35    regulation contains inherent cyclic feedback mechanisms, hence we should not model a GRN as a  
36    DAG, and (2) observations are limited and have significant measurement noise, hence there exists a

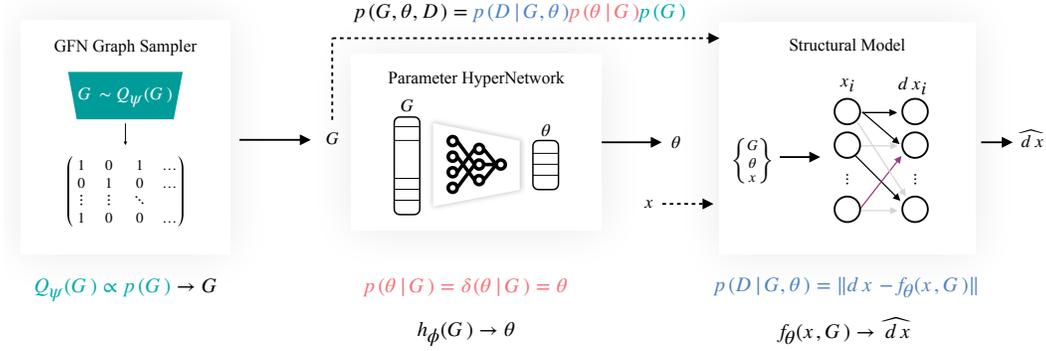


Figure 1: Architecture for Bayesian structure learning of dynamical systems. DynGFN consists of three main components: A GFlowNet modeling a posterior distribution over graphs  $Q(G|\mathcal{D})$ , a HyperNetwork modeling a posterior over parameters given a graph  $Q(\theta|G, D)$ , and the structural equation model scoring  $G$  and  $\theta$  according to how well they fit the data. Although the figure shows the case where  $Q(G|\mathcal{D})$  is modelled with a GFlowNet, this can be any arbitrary graph sampler that can sample discrete structures  $G \sim Q(G|\mathcal{D})$ .

37 large equivalence class of graphs that are likely given datasets with typical sample sizes. Existing  
 38 methods either focus on (1) – identifying graphs with *cyclic* structure by leveraging dynamics [14, 12]  
 39 or assuming the system is in equilibrium [36], or (2) – learning complex Bayesian *posteriors* over  
 40 explanatory DAGs [11], but not both. In this work, we address both challenges concurrently in a fully  
 41 differentiable end-to-end pipeline (see Figure 1).

42 To accomplish this, we treat structure learning as a problem of sparse identification of a dynamical  
 43 system. From a dynamical systems perspective, one can model both causal structure between variables  
 44 as well as their time-dependent system response with the drift function [37, 43]. We leverage the  
 45 fact that we can estimate the rate of change of a gene’s expression (velocity) with *RNA velocity*  
 46 methods [8]. This data takes the form of dynamic tuple pairs  $(x, dx)$ , which we can use to pose the  
 47 dynamical system learning problem as a regression task (see Figure 1). This significantly simplifies  
 48 the learning objective as we can model system dynamics while also learning structure without the  
 49 need for numerically intensive differential equation solvers. We view this as a step towards Bayesian  
 50 structure learning from continuous dynamics – we term this *Bayesian dynamic structure learning*.

51 Our approach estimates the posterior over the sparse dependencies and parameters of the dynamical  
 52 system. This is important in scientific applications because it is usually prohibitively expensive  
 53 to acquire a enough data to uniquely identify the true graph underlying a data generating process.  
 54 Capturing the complex distribution over candidate structure is critical for downstream scientific  
 55 applications and is an essential step in active causal discovery [39, 52, 18]. This is especially important  
 56 in settings where experiments are expensive, e.g. conducting genetic perturbations for inference of  
 57 GRNs. *Bayesian structure learning* is a class of methods that try to model this distribution over  
 58 structure from observed data. These methods model posteriors over admissible structures  $P(G|D)$   
 59 that explain the observations [32, 10, 4, 11, 31], but focus on modelling distributions over DAGs.

60 Our approach leverages Generative Flow Networks (GFlowNets) to model complex distributions over  
 61 *cyclic* structures. GFlowNets [6, 7] parameterize the distribution over any discrete object (e.g. graphs)  
 62 through a sequential policy, and as a result avoid needing to make restrictive parametric assumptions  
 63 on the distribution. This makes them a useful tool in structure learning, particularly in cases where  
 64  $P(G|D)$  is discrete and complex [11]. In this work, we use GFlowNets to learn posteriors over the  
 65 sparse structure in a dynamical system, and separately learn the posteriors over the parameters of the  
 66 drift function via a HyperNetwork [15] that conditions on inferred structures. Our main contributions  
 67 are summarized as follows:

- 68 • We develop a novel framework for Bayesian structure learning under the lens of dynamical  
 69 system identification for modelling complex posteriors over cyclic graphs. We consider flexible  
 70 parameterizations for the structural model such that we can capture both linear and non-linear  
 71 dynamic relationships.

- 72 • We design a novel GFlowNet architecture, Dynamic GFlowNet (DynGFN), tailored for modelling  
73 posteriors over cyclic structures. We propose a *per-node* factorization within DynGFN that enables  
74 efficient search over the discrete space of cyclic graphs.
- 75 • We empirically evaluated DynGFN on synthetic dynamic data designed to induce highly multi-  
76 modal posteriors over graphs.
- 77 • We showcase the use of DynGFN on a real biological system using single-cell RNA-velocity data  
78 for learning posteriors of GRNs.

## 79 2 Related Work

80 There is a lot of prior work on the problem of identifying causal structure  $G$  from either observa-  
81 tional [e.g. 49, 54, 35] or interventional [e.g. 26, 29, 38] data, but the majority of existing methods  
82 return only the most likely DAG under the observed data. Bayesian approaches attempt to explicitly  
83 model this distribution over admissible DAGs.

84 **Bayesian Structure Learning:** Recently, there has been significant interest in fully differentiable  
85 Bayesian methods for structure learning in the static case. DiBS [32], BCD-Nets [10], VCN [4], and  
86 DAG-GFlowNet [11] all attempt to learn a distribution over structural models from a fully observed  
87 system. The key difference is in how these methods parameterize the graph. DiBS is a particle  
88 variational inference method that uses two matrices  $U$  and  $V$  where  $G = \text{sigmoid}(U^T V)$  where the  
89 sigmoid is applied elementwise which is similar to graph autoencoders. BCD-Nets and DP-DAG  
90 use the Gumbel-Sinkhorn distribution to parameterize a permutation and direct parameterization of a  
91 lower triangular matrix. VCN uses an autoregressive LSTM to generate the graph as this gets rid of  
92 the standard uni-modal constraint of Gaussian distributed parameters. DAG-GFN has shown success  
93 for modelling  $P(G|\mathcal{D})$  [11]. However, it remains restrictive to assume the underlying structure of  
94 the observed system is a DAG as natural dynamical systems typically contain regulating feedback  
95 mechanisms. This can be particularly challenging for GFlowNets since including cycles in the  
96 underlying structure exponentially increases the discrete search space. We show that under certain  
97 assumptions we can in part alleviate this shortcoming for learning Bayesian posteriors over cyclic  
98 structures for dynamical systems. In small graphs, these methods can model the uncertainty over  
99 possible models (including over Markov equivalence classes).

100 **Dynamic and Cyclic Structure Learning:** There has been comparatively little work towards  
101 Bayesian structure learning from dynamics. Recent works in this direction based on NeuralODEs [9]  
102 propose a single explanatory structure [50, 5, 1, 2]. CD-NOD leverages heterogeneous non-stationary  
103 data for causal discovery when the underlying generative process changes over time [53, 20]. A  
104 similar approach uses non-stationary time-series data for causal discovery and forecasting [19].  
105 DYNOTEARS is a score-based approach that uses time-series to learn structure [40]. However,  
106 these methods do not attempt to explicitly model a distribution over the explanatory structure. Other  
107 methods aim to learn cyclic dependencies in the underlying graph [23, 36, 28, 3]. For instance, [23]  
108 propose an iterative method that leverages interventional data to learn directed cyclic graphs. It is  
109 suggested that CD-NOD is also extendable to learn cyclic structure [20]. But these methods do not  
110 model a posterior over structure. In general, there remains a gap for the problem of Bayesian structure  
111 learning over cyclic graphs.

112 We include further discussion on related work for GRN inference from single-cell transcriptomic  
113 data and cell dynamics in Appendix C.1.

## 114 3 Preliminaries

### 115 3.1 Bayesian Dynamic Structure Learning

116 **Problem Setup:** We consider a finite dataset,  $\mathcal{D}$ , of dynamic pairs  $(x, dx) \in \mathbb{R}^d \times \mathbb{R}^d$  where  $x$   
117 represents the state of the system sampled from an underlying time-invariant stochastic dynamical  
118 system governed by a latent drift  $\frac{dx}{dt} = f(x, \epsilon)$  where  $\epsilon$  is a noise term that parameterizes the SDE;  $x$   
119 and  $\epsilon$  are mutually independent. The latent drift has some fixed sparsity pattern i.e.  $\frac{\partial f_i}{\partial x_j} \neq 0$  for a  
120 small set of variables, which can be parameterized by a graph  $G$  such that  $g_{ij} = \mathbf{1}[\frac{\partial f_i}{\partial x_j} \neq 0]$ , where

121  $g_{ij} \in G, i = 1, \dots, d, j = 1, \dots, d$ . The variables  $x_j$  for which  $\frac{\partial f_i}{\partial x_j} \neq 0$  can be interpreted as the  
 122 causal parents of  $x_i$ , denoted  $\text{Pa}(x_i)$ . This lets us define an equivalent dynamic structural model  
 123 [37, 43] of the form,

$$\frac{dx_i(t)}{dt} = f_i(\text{Pa}(x_i), \epsilon_i), \quad (1)$$

124 for  $i = 1, \dots, d$ . For the graph  $G$  to be identifiable, we assume that all relevant variables are observed,  
 125 such that *causal sufficiency* is satisfied.

126 Our goal is to model our posterior over explanatory graphs  $Q(G|\mathcal{D})$  given the data. We aim to  
 127 jointly learn distribution over parameters  $\theta$  that parameterize the latent drift  $f(x)$ ; these parameters  
 128 will typically depend on the sparsity pattern such that, i.e.  $p(\theta|G) \neq p(\theta)$ . We can factorize this  
 129 generative model as follows,

$$p(G, \theta, \mathcal{D}) = p(\mathcal{D}|G, \theta)p(\theta|G)p(G) \quad (2)$$

130 This factorization forms the basis of our inference procedure. We learn a parameterized function  
 131  $f_\theta(x) : \mathbb{R}^d \rightarrow \mathbb{R}^d$  that approximates the structural model defined in (1). To model this joint  
 132 distribution, we need a way of representing,  $P(G)$ , a distribution over the combinatorial space of  
 133 possible sparsity patterns, and  $P(\theta|G)$ , the posterior over the parameters of  $f_\theta$ . We use GFlowNets  
 134 [6] represent  $P(G)$ , and a HyperNetwork to parameterize  $P(\theta|G)$ .

### 135 3.2 Generative Flow Networks

136 GFlowNets are an approach for learning generative models over spaces of discrete objects [6, 7].  
 137 GFlowNets learn a stochastic policy  $P_F(\tau)$  to sequentially sample an object  $\mathbf{x}$  from a discrete space  
 138  $\mathcal{X}$ . Here  $\tau = (s_0, s_1, \dots, s_n)$  represents a full Markovian trajectory over plausible discrete states,  
 139 where  $s_n$  is the terminating state (i.e. end of a trajectory) [34]. The GFlowNet is trained such that at  
 140 convergence, sequential samples from the stochastic policy over a trajectory,  $\mathbf{x} \sim P_F(\tau)$ , i.e.  $\mathbf{x} = s_n$ ,  
 141 are equal in distribution to samples from the normalized reward distribution  $P(x) = \frac{R(\mathbf{x})}{\sum_{\mathbf{x}' \in \mathcal{X}} R(\mathbf{x}')}.$   
 142 The GFlowNet policies are typically trained by optimizing either the *Trajectory Balance* (TB)  
 143 loss [34], *Subtrajectory Balance* (Sub-TB) loss [33], or the *Detailed Balance* (DB) loss [11]. In this  
 144 work, we exploit the DB loss to learn a stochastic policy for directed graph structure.

145 **Detailed Balance Loss:** The DB loss [11] leverages the fact that the reward function can be  
 146 evaluated for any partially constructed graph (i.e. any prefix of  $\tau$ ), and hence we get intermediate  
 147 reward signals for training the GFlowNet policy. The DB loss is defined as:

$$\mathcal{L}_{\text{DB}}(s_i, s_{i-1}) = \left( \log \frac{R(s_i)P_B(s_{i-1}|s_i; \psi)P_F(s_n|s_{i-1}; \psi)}{R(s_{i-1})P_F(s_i|s_{i-1}; \psi)P_F(s_n|s_i; \psi)} \right)^2, \quad (3)$$

148 where  $P_F(s_i|s_{i-1}; \psi)$  and  $P_B(s_{i-1}|s_i; \psi)$  represent the forward transition probability and backward  
 149 transition probability, and a trainable normalizing constant, respectively. Under this formulation,  
 150 during GFlowNet training the reward is evaluated at every state. For this reason, the DB formulation  
 151 is in general advantageous for the structure learning problem where any sampled graph can be viewed  
 152 as a complete state, hence more robustly inform gradients when training the stochastic policy than  
 153 counterpart losses. Previous work has shown GFlowNets are useful in settings with multi-modal  
 154 posteriors. This is of particular interest to us where many admissible structures can explain the  
 155 observed data equally well. We model  $Q_\psi(G)$  using  $P_F(s_i|s_{i-1}; \psi)$  and learn the parameters  $\psi$ .

## 156 4 DynGFN for Bayesian Dynamic Structure Learning

157 We present a general framework for Bayesian dynamic structure learning and propose a GFlowNet  
 158 architecture, DynGFN, tailored for modelling a posterior over discrete cyclic graphical structures.  
 159 We summarize our framework in Figure 1 and Algorithm 1. DynGFN consists of 3 key modules:

- 160 1. A graph sampler that samples graphical structures that encode the structural dependencies  
 161 among the observed variables. This is parameterized with a GFlowNet that iteratively adds  
 162 edges to a graph.

---

**Algorithm 1** Batch update training of DynGFN

---

```

1: Input: Data batch  $(x_b, dx_b)$ , initial NN weights  $\psi, \phi$ ,  $L^0$  sparsity prior  $\lambda_0$ , and learning rate  $\epsilon$ .
2:  $s_0 \leftarrow \mathbf{0}_{B \times d \times d}$  ▷ Training is paralleled over  $B$  graph trajectories
3:  $a \sim P_F(s_1|s_0; \psi)$ , ▷ Sample initial actions vector
4: while  $a$  not  $\emptyset$  do
5:   Compute  $P_F(s_i|s_{i-1}; \psi), P_B(s_{i-1}|s_i; \psi)$ 
6:    $\theta \leftarrow h_\phi(s_i)$ 
7:    $\widehat{dx}_b \leftarrow f_\theta(x, s_i)$ 
8:    $R_i(s_i) \leftarrow e^{-\|dx_b - \widehat{dx}_b\|_2^2 + \lambda_0 \|s_i\|_0}$ 
9:    $\psi \leftarrow \psi - \epsilon \nabla_\psi \mathcal{L}_{DB}(s_i, s_{i-1})$  ▷  $\mathcal{L}_{DB}(s_i, s_{i-1})$  computed as in Equation 3
10:   $a \sim P_F(s_i|s_{i-1}; \psi), s_i \rightarrow s_{i+1}$  ▷ Take action step to go to next state
11:  $\phi \leftarrow \phi - \epsilon \nabla_\phi \log R$ 
return Updated GFN weights  $\psi$  and updated HyperNetwork weights  $\phi$ .

```

---

- 163 2. A model that approximates the structural equations defined in (1) to model the functional  
164 relationships between the observed variables, indexed by parameters  $\theta$ . This is a class of  
165 functions that respect the conditional independencies implied by the graph sampled in step  
166 1. We enforce this by masking inputs according to the graph.
- 167 3. Because the functional relationships between variables may be different depending on which  
168 graph is sampled, we use a HyperNetwork architecture that outputs the parameters  $\theta$  of the  
169 structural equations as a function of the graph. We also show that under linear assumptions  
170 of the structural modules, we can solve for optimal  $\theta$  analytically (i.e. we do not need the  
171 HyperNetwork).

172 For training, we assume  $L^0$  sparsity of graphs  $G$  to constrain the large discrete search space over  
173 possible structures. We use a reward  $R$  for a graph  $G$  and  $L^0$  penalty of the form:  $R(G) =$   
174  $e^{-\|dx - \widehat{dx}\|_2^2 + \lambda_0 \|G\|_0}$ . We motivate this set-up so we can estimate  $\widehat{dx}$  close to  $dx$  in an end-to-end  
175 learning pipeline. Since estimates for  $\widehat{dx}$  are dependent on  $G$  and  $\theta$ , this reward informs gradients to  
176 learn a policy that can approximate  $Q(G)$  given dynamic data.

177 The main advantage of DynGFN comes when modelling complex posteriors with many modes. Prior  
178 work has shown GFlowNets are able to efficiently model distributions where we can share information  
179 between different modes [34]. The challenge we tackle is how to do this with a changing objective  
180 function, as the GFlowNet objective is a function of the current parameter HyperNetwork and the  
181 structural equations. We use multilayer perceptrons (MLPs) to parameterize the stochastic GFlowNet  
182 policy, HyperNetwork architecture, and the dynamic structural model<sup>1</sup>.

### 183 4.1 Graph Sampler

184 DynGFN models a posterior distribution over graphs  $Q(G|\mathcal{D})$  given a finite set of observations. To  
185 learn  $Q(G|\mathcal{D})$ , DynGFN needs to explore over a large discrete state space. Since we aim to learn  
186 a bipartite graph between  $x$  and  $dx$ , DynGFN needs to search over  $2^{d^2}$  possible structures, where  
187  $d$  denotes the dimensionality of the system and  $2^{d^2}$  the number of possible edges in  $G$ . For even  
188 moderate  $d$ , this discrete space is very large (e.g. for  $d = 20$  we have  $2^{400}$  possible graphs).

189 However, under the assumption of causal sufficiency, we can significantly reduce this search space,  
190 by taking advantage of the fact that  $Q(G|\mathcal{D})$  factorizes as follows,

$$Q(G|\mathcal{D}) = \prod_{i \in [1, \dots, d]} Q_i(G[\cdot, i]|\mathcal{D}) \quad (4)$$

191 By using this model, we reduce the search space from  $2^{d^2} \rightarrow d2^d$ . For  $d = 20$  this is  $\approx 2^{104}$ . While  
192 still intractable to search over, it is still a vast improvement over the unfactorized case. We call this  
193 model a *per-node* posterior, and we use a per-node GFlowNet going forward. We discuss details  
194 regarding encouraging forward policy exploration during training in Appendix B.6.

<sup>1</sup>When we assume linear structural relationships, we can solve for the parameters analytically, thus do not need MLPs for the HyperNetwork and dynamic structural model. This is further discussed in section 4.2

195 **4.2 HyperNetwork and Structural Model**

196 We aim to jointly learn the structural encoding  $G$  and parameters  $\theta$  that together model the structural  
 197 relationships  $dx = f_\theta(x, G)$  of the dynamical system variables. To accomplish this, we propose  
 198 learning an individual set of parameters  $\theta$  for each graph  $G$ , independent of the input data  $x$ . This  
 199 approach encapsulates  $P(\theta|G)$  in (2). We use a HyperNetwork architecture that takes  $G$  as input  
 200 and outputs the structural equation model parameters  $\theta$ , i.e.  $\theta = h_\phi(G)$  hence  $P(\theta|G) = \delta(\theta|G)$  –  
 201 allowing us to learn a separate  $\theta$  for each  $G$ . This HyperNetwork model does not capture uncertainty  
 202 in the parameters, however the formulation may be extended to the Bayesian setting by placing a  
 203 prior on the HyperNetwork parameters  $\phi$ . Although  $h_\phi$  allows for expressive parameterizations for  $\theta$ ,  
 204 it may not be easy to learn<sup>2</sup>. HyperNetworks have shown success in learning parameters for more  
 205 complex models (e.g. LSTMs and CNNs) [15], hence motivates their fit for our application.

206 **Linear Assumption on Dynamic Structural Model:** In addition to analytically modelling linear  
 207 systems, in some cases it may suffice to assume a linear differential form  $\frac{dx}{dt} = Ax$  to approximate  
 208 dynamics. In this setting, given a sampled graph  $G \sim Q(G)$  and  $n$  i.i.d. observations of  $(x, dx)$  we  
 209 can solve for  $\theta = A$  analytically. To induce dependence on the graph structure, we use the sampled  
 210  $G$  as a mask on  $x$  and construct  $\tilde{x}_i = G_i^T \odot x$ . Then we can solve for  $\theta$  on a per-node basis as

$$\theta_i = (\tilde{x}_i^T \tilde{x}_i + \lambda I)^{-1} \tilde{x}_i^T dx_i, \quad (5)$$

211 where  $i = 1, \dots, d$ ,  $\lambda > 0$  is the precision of an independent Gaussian prior over the parameters, and  
 212  $I$  is the identity matrix. We use  $\lambda = 0.01$  throughout this work.

213 **5 A Useful Model of Indeterminacy**

214 In order to evaluate DynGFN  
 215 ability to model complex pos-  
 216 teriors over graphs, we need a  
 217 structure learning problem with  
 218 a large equivalence class of ad-  
 219 missible graphs. We present a  
 220 simple way to augment a set  
 221 of identifiable dynamics under  
 222 some model to create a combi-  
 223 natorial number of equally  
 224 likely dynamics under the same  
 225 model. More specifically, this  
 226 creates a ground truth posterior  
 227  $Q^*(G|D) \propto \sum T(G^*)$  where  
 228  $T(\cdot) : \mathcal{G} \rightarrow \mathcal{G}$  is an analytically  
 229 computable transformation over graphs and  $G^*$  is the identified graph under the original dynamics.  
 230 We use this system to test how well we can learn a posterior over structures that matches what we see  
 231 in single-cell data.

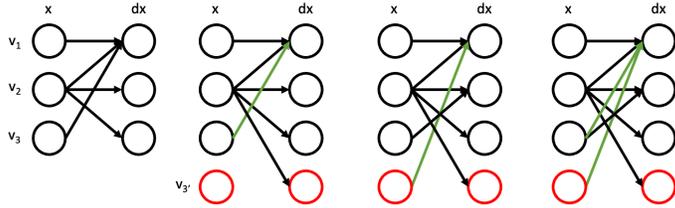


Figure 2: For an identifiable graph, we add a new variable which has the same values as  $v_3$  and creates three possible explanations for the data (green). If we consider a sparsity penalty, then we can eliminate the last possibility (which has two additional edges) for only two possible graphs.

232 Specifically, given a dataset of  $(x, dx) \in \mathbb{R}^d \times \mathbb{R}^d$  pairs, we create a new dataset with  $d + 1$  variables  
 233 where the ‘new’ variable  $v'$  is perfectly correlated with an existing variable  $v$ . In causal terms, this  
 234 new variable inherits the same parents as  $v$ , that is  $\text{Pa}(v') := \text{Pa}(v)$  and the same structural equations  
 235 as  $v$ , that is  $dv' = dv$ . This is depicted in Figure 2. This creates a number of new possible explanatory  
 236 graphs, which we generalize with the following proposition.

237 **Proposition 1.** *Given any  $d$  dimensional ODE system with  $\mathcal{G}^*$  identifiable under  $f \in \mathcal{F}$ , the*  
 238  *$D = d + a$  dimensional system  $\frac{dx}{dt} = Ax$ , denote the vector of multiplicities  $m \in \mathbb{N}^d$  with  $m_i$  as*  
 239 *the number of repetitions of each variable. Then this construction creates an admissible family of*  
 240 *graphs  $\mathcal{G}'$  where  $|\mathcal{G}'| = \prod_{i \in d} (2^{m_i} - 1)^{\mathcal{G}_i^*}$ . Furthermore, under an  $L^0$  penalty on  $G$ , this reduces to*  
 241  *$\prod_i (m_i)^{\mathcal{G}_i^*}$ .*

242 See Appendix A for full proof. The intuition behind this proposition can be seen from the case of  
 243 adding a single copied variable. This corresponds to  $A = [\delta_v I_d]$  where  $\delta_v$  is a vector with a 1 on

<sup>2</sup>We discuss training dynamics when using  $h_\phi$  in Appendix B.7.

Table 1: Bayesian dynamic structure learning of linear and non-linear systems with  $d = 20$  variables. The graphs representing the structural dynamic relationships of the linear and non-linear systems have 50 edges out of possible 400. The ground truth discrete distribution  $P(G^*)$  contains 1024 admissible graphs for each respective system. The  $\ell$  and  $h$  pre-fix denote usage of the analytic linear solver and HyperNetwork solver for structural model parameters, respectively. Results are reported on held out test data over 5 model seeds.

Model	Bayes-SHD ↓	Linear System		NLL ↓
		AUC ↑	KL ↓	
$\ell$ -DynBCD	$32.0 \pm 0.27$	$0.71 \pm 0.0$	$1707.45 \pm 9.66$	—
$\ell$ -DynDiBS	$29.2 \pm 0.78$	$0.71 \pm 0.0$	$6622.43 \pm 171.67$	—
$\ell$ -DynGFN	<b><math>22.8 \pm 1.4</math></b>	<b><math>0.75 \pm 0.01</math></b>	<b><math>1091.60 \pm 35.72</math></b>	—
$h$ -DynBCD	<b><math>5.5 \pm 1.1</math></b>	$0.89 \pm 0.04$	$701.19 \pm 46.99$	$(9.83 \pm 0.59)E - 5$
$h$ -DynDiBS	$28.5 \pm 4.2$	$0.51 \pm 0.07$	$7934.90 \pm 381.80$	<b><math>(8.17 \pm 1.30)E - 6</math></b>
$h$ -DynGFN	<b><math>6.7 \pm 0.0</math></b>	<b><math>0.94 \pm 0.0</math></b>	<b><math>350.92 \pm 30.15</math></b>	$(8.35 \pm 0.02)E - 3$

Model	Bayes-SHD ↓	Non-linear System		NLL ↓
		AUC ↑	KL ↓	
$\ell$ -DynBCD	$77.5 \pm 8.3$	$0.42 \pm 0.03$	$3814.86 \pm 354.56$	—
$\ell$ -DynDiBS	$75.7 \pm 7.7$	<b><math>0.59 \pm 0.01</math></b>	$5893.65 \pm 59.66$	—
$\ell$ -DynGFN	<b><math>45.7 \pm 0.6</math></b>	$0.55 \pm 0.0$	<b><math>226.25 \pm 6.58</math></b>	—
$h$ -DynBCD	$192.9 \pm 0.7$	$0.50 \pm 0.0$	$9108.69 \pm 51.34$	$(3.83 \pm 0.32)E - 4$
$h$ -DynDiBS	$48.1 \pm 9.0$	$0.53 \pm 0.10$	$8716.64 \pm 265.29$	<b><math>(4.06 \pm 0.10)E - 6</math></b>
$h$ -DynGFN	<b><math>32.6 \pm 0.9</math></b>	<b><math>0.67 \pm 0.01</math></b>	<b><math>193.28 \pm 8.53</math></b>	$(1.47 \pm 0.11)E - 3$

node  $v$  and zeros elsewhere, and  $I_d$  is the  $d$ -dimensional identity matrix. Let  $v$  have  $c$  children, such that  $v \in Pa(c)$  in the identifiable system, then any of those  $c$  child nodes could depend either on  $v$  or on the new node  $v'$  or both. This creates  $3^c$  possible explanatory graphs. If we restrict ourselves to the set of graphs with minimal  $L^0$  norm, then we eliminate the possibility of a child node depending on both  $v$  and  $v'$ , this gives  $2^c$  possible graphs, choosing either  $v$  or  $v'$  as a parent.

## 6 Experimental Results

In this section we evaluate the performance of DynGFN against counterpart Bayesian structure learning methods (see Appendix B.2 for details). Since our primary objective is to learn Bayesian posteriors over discrete structure  $G$ , we compare to Bayesian methods that can also accomplish this task, i.e. versions of BCD-Nets [10] and DiBS [32]. We show in certain cases, DynGFN is able to better capture the true posterior when there are a large number of modes. We evaluate methods according to four metrics: Bayes-SHD, area under the receiver operator characteristic curve (AUC), Kullback–Leibler (KL) divergence between learned posteriors  $Q(G)$  and the distribution over true graphs  $P(G^*)$ , and the negative log-likelihood (NLL)  $P(D|G, \theta)$  (in our setting this reduces to the mean squared error between  $\widehat{dx}$  and  $dx$ , given  $\theta$  and sampled  $G'$ s). Since the analytic linear solver requires data at run-time to compute optimal parameters for the structural model, we include the NLL metric only for models using the HyperNetwork solver. Bayes-SHD measures the average distance to the closest structure in the admissible set of graphs according to the structural hamming distance, which in this case is simply the hamming distance of the adjacency matrix representation to the closest admissible graph. We assume  $P(G^*)$  is uniform over  $G^*$  and include further details about evaluating the quality of learned posteriors in Appendix B.8.

### 6.1 Experiments with Synthetic Data

We generated synthetic data from two systems using our indeterminacy model presented in section 5: (1) a linear dynamical system  $dx = Ax$ , and (2) a non-linear dynamical system  $dx = \text{sigmoid}(Ax)$ . We consider  $\ell$ -DynGFN and  $h$ -DynGFN, i.e. DynGFN with the linear analytic parameter solver as shown in (5), and DynGFN with the HyperNetwork parameter solver  $h_\phi$ . Likewise, we compare  $\ell$ -DynGFN and  $h$ -DynGFN to counterpart Bayesian baselines which we call  $\ell$ -DynBCD,  $\ell$ -DynDiBS,  $h$ -DynBCD, and  $h$ -DynDiBS. To constrain the discrete search procedure, we assume a sparse prior on the structure the graphs  $G$ , specifically the  $L^0$  prior. Due to challenging iterative optimization dynamics present when using  $\theta = h_\phi(G)$  for DynGFN, to train initialize the forward policy  $P_F(s_i|s_{i-1}; \psi)$

274 using the  $\psi$  learned in  $\ell$ -DynGFN to provide a more admissible starting point for learning  $h_\phi$  (we  
 275 discuss further details in Appendix B.7). We do not need to do this for  $h$ -DynBCD and  $h$ -DynDiBS  
 276 as we are able to train both models end-to-end without iterative optimization. In Table 1 we show  
 277 results of our synthetic experiments for learning posteriors over multi-modal distributions of cyclic  
 278 graphs. We observe the DynGFN is most competitive on both synthetic systems for modelling the  
 279 true posterior over structure. Details about DynGFN, baselines, and accompanying hyper-parameters  
 280 can be found in Appendix B.

## 281 6.2 Ablations Over Sparsity and Linearity of Dynamic Systems

282 We conduct two ablations: (1) ablation over sparsity of the dynamic system structure, and (2) ablation  
 283 over  $\Delta t$ , the time difference between data points of dynamic simulation. For a sparsity level of 0.9,  
 284 the ground truth graphs have 50 edges out of  $d^2$  possible edges. In these experiments,  $P(G^*)$  has  
 285 1024 modes. We conduct the ablations over 5 random seeds for each set of experiments.

286 **Sparsity:** DynGFN uses the  $L^0$  prior on  
 287  $G$  throughout training. Under this setting,  
 288 system sparsity carries significant weight  
 289 on the ability to learn posteriors over the  
 290 structured dynamics of a system. We show  
 291 this trend in Table 2. We note that comput-  
 292 ing the KL-divergence for DynGFN, specif-  
 293 ically computing the probability generat-  
 294 ing a true  $G$ , becomes computationally in-  
 295 tractable as  $G$  is less sparse<sup>3</sup>. For systems  
 296 of 0.9 and 0.95 sparsity, we observe a de-  
 297 creasing trend in KL and Bayes-SHD, and  
 298 an increasing trend in AUC. This result is  
 299 expected as DynGFN can better traverse  
 300 sparse graphs as the combinatorial space  
 301 over possible trajectories is smaller relative  
 302 to denser systems.

303 **Linearity:** Training DynGFN via the lin-  
 304 ear solver for the structural model param-  
 305 eters is an easier objective due to simplified  
 306 training dynamics. Because of this, we explore the performance of  $\ell$ -DynGFN assuming  $f_\theta$  for  
 307 modelling equation (1) to be linear in the non-linear system. We do this by conducting an ablation  
 308 over  $\Delta t$  and find that the performance of  $\ell$ -DynGFN on the non-linear system improves as  $\Delta t \rightarrow 0$ .  
 309 We show a portion of this trend in Table 3.

## 310 6.3 Experiments on Single-Cell RNA-velocity Data

311 To show how DynGFN can be applied to single cell data we use a cell cycle dataset of human  
 312 Fibroblasts [46]. As a motivating example we show the correlation structure of single-cell RNA-seq  
 313 data from human Fibroblast cells [46] Figure 3. We show both the raw correlation and the correlation  
 314 over cell cycle time, which is significantly higher. With such a pure cell population whose primary  
 315 axis of variation is state in the cell cycle by aggregating over cell cycle time we expect observation  
 316 noise to be averaged out, leading to a “truer” view of the correlation between latent variables. Further  
 317 details for this experimental set-up are provided in Appendix C.1.3. Since there are many genes which  
 318 are affected by the cell cycle phase, there are many correlated variables that are downstream of the  
 319 true cell cycle regulators. This provides a natural way of using cell cycle data to evaluate a model’s  
 320 ability to capture the Bayesian posterior. In Table 4 we show results for learning posteriors over an  
 321 undetermined GRN using RNA velocity data. We find that  $\ell$ -DynGFN and  $h$ -DynGFN yield low KL

Table 2: Ablation for  $\ell$ -DynGFN on  $d = 20$  systems with varying levels of sparsity and fixed  $\Delta t = 0.05$ .

Sparsity	Bayes-SHD ↓	AUC ↑	KL ↓
0.95	16.4 ± 1.71	0.79 ± 0.0	889.57 ± 31.24
0.90	22.8 ± 1.41	0.75 ± 0.01	1091.60 ± 35.72
0.85	32.8 ± 0.72	0.71 ± 0.0	—
0.80	39.2 ± 0.69	0.71 ± 0.0	—
0.75	60.2 ± 1.17	0.66 ± 0.01	—

Table 3: Ablation for  $\ell$ -DynGFN on  $d = 20$  non-linear systems with varying  $\Delta t$  and fixed sparsity at 0.9.

$\Delta t$	Bayes-SHD ↓	AUC ↑	KL ↓
0.001	38.7 ± 0.80	0.61 ± 0.0	202.41 ± 9.95
0.005	39.0 ± 0.81	0.60 ± 0.0	206.83 ± 11.55
0.01	40.6 ± 1.13	0.59 ± 0.0	202.71 ± 7.74
0.05	45.7 ± 0.62	0.55 ± 0.0	226.25 ± 6.58
0.1	51.8 ± 0.18	0.50 ± 0.0	264.86 ± 2.17

<sup>3</sup>For example, since DynGFN constructs one object  $G$  sequentially over a state space distribution, we must compute probabilities of all combinatorial state trajectories for constructing  $G = (s_i, \dots, s_n)$ . The space of combinatorial state trajectories is  $n!$  in nature, hence this computation is only possible for small graphs and/or sparse graphs.

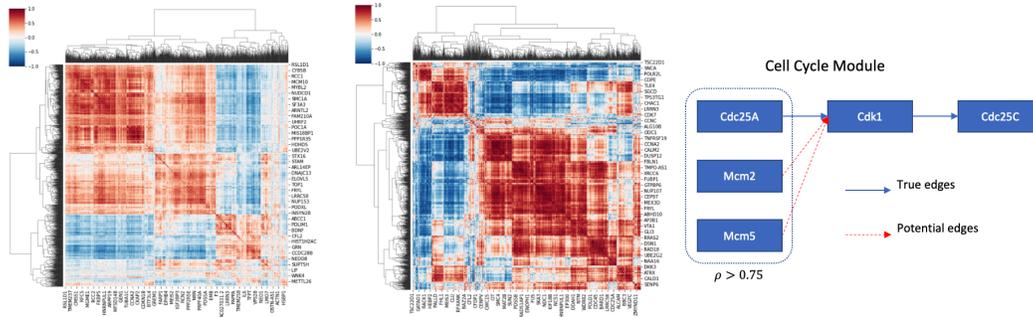


Figure 3: (Left) Correlation structure in the raw single cell data over 5000 cells and 2000 genes selected by scVelo [8] pre-processing. (Middle) Correlation structure among genes over (inferred) cell cycle times. This stronger correlation structure is more reflective of the correlation in the underlying system. (Right) Cdc25A is known to inhibit Cdk1 which is known to inhibit Cdc25C, while the Mcm complex is highly correlated with Cdc25A, they do not directly interact with Cdk1 [24].

Table 4: Bayesian dynamic structure learning 5-D cellular system using scRNA velocity data. The dynamics of this system are unknowns, however we identify 81 admissible graphs between variables (genes) that describe the data. We train models over 5 seeds. The graphs of this system contain of 7 true edges.

Model	Cellular System - RNA Velocity			NLL ↓
	Bayes-SHD ↓	AUC ↑	KL ↓	
$\ell$ -DynBCD	<b>2.6 ± 0.1</b>	0.56 ± 0.01	321.95 ± 3.34	—
$\ell$ -DynDiBS	6.5 ± 0.4	0.47 ± 0.01	550.17 ± 16.63	—
$\ell$ -DynGFN	<b>3.3 ± 0.4</b>	<b>0.59 ± 0.03</b>	<b>44.98 ± 18.60</b>	—
$h$ -DynBCD	10.1 ± 0.8	0.53 ± 0.03	587.41 ± 24.00	0.094 ± 0.003
$h$ -DynDiBS	9.6 ± 4.2	0.51 ± 0.13	560.85 ± 83.83	<b>0.084 ± 0.0</b>
$h$ -DynGFN	<b>5.1 ± 1.2</b>	<b>0.58 ± 0.05</b>	<b>39.82 ± 28.05</b>	0.109 ± 0.001

322 and moderate Bayes-SHD. While  $\ell$ -DynBCD performs well in terms of identify a small distribution  
 323 of true  $G$ 's, it falls short in modelling the true posterior (this can be seen from low Bayes-SHD, high  
 324 KL).

## 325 7 Conclusion

326 We presented DynGFN, a method for Bayesian dynamic structure learning. In low dimensions we  
 327 found that DynGFN is able to better model the distribution over possible explanatory structures than  
 328 counterpart Bayesian structure learning baseline methods. As a proof of concept, we presented an  
 329 example of learning the distribution over likely explanatory graphs for linear and non-linear synthetic  
 330 systems where complex uncertainty over explanatory structure is prevalent. We demonstrate the  
 331 use of DynGFN for learning gene regulatory structure from single-cell transcriptomic data where  
 332 there are many possible graphs, showing DynGFN can better model the uncertainty over possible  
 333 explanations of this data rather than capturing a single explanation.

334 **Limitations and Future Work:** We have demonstrated a degree of efficacy when using DynGFN  
 335 for Bayesian structure learning with dynamic observational data. A key limitation of DynGFN  
 336 is scaling to larger systems. To effectively model  $P(G, \theta, D)$ , DynGFN needs to search over an  
 337 environment state space of possible graphs. This state space grows exponentially with the number of  
 338 possible edges, i.e.  $2^{d^2}$  or  $d2^d$  for per-node-GFN where  $d$  is the number of variables in the system.  
 339 Therefore, DynGFN is currently limited to smaller systems. Nevertheless, there are many applications  
 340 where Bayesian structure learning, even over 5-20 dimensional examples that we explore here, could  
 341 be extraordinarily useful. We include further discussion of scaling DynGFN in Appendix C.3 with  
 342 some ideas on how to approach this challenge. We found that training DynGFN requires some  
 343 selection of hyper-parameters and in particular parameters that shape the reward function. Selecting  
 344 hyper-parameters for the baseline methods prove more difficult for this task.

## References

- 345
- 346 [1] Aliee, H., Theis, F. J., and Kilbertus, N. Beyond predictions in neural odes: Identification and  
347 interventions. *arXiv preprint 2106.12430*, 2021.
- 348 [2] Aliee, H., Richter, T., Solonin, M., Ibarra, I., Theis, F., and Kilbertus, N. Sparsity in continuous-  
349 depth neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- 350 [3] Améndola, C., Dettling, P., Drton, M., Onori, F., and Wu, J. Structure learning for cyclic linear  
351 causal models. *Uncertainty in Artificial Intelligence (UAI)*, 2020.
- 352 [4] Annadani, Y., Rothfuss, J., Lacoste, A., Scherrer, N., Goyal, A., Bengio, Y., and Bauer, S.  
353 Variational causal networks: Approximate bayesian inference over causal structures. *arXiv  
354 preprint*, 2021.
- 355 [5] Bellot, A. and Branson, K. Neural Graphical Modelling in Continuous Time: Consistency  
356 Guarantees and Algorithms. *International Conference on Learning Representations (ICLR)*,  
357 2022.
- 358 [6] Bengio, E., Jain, M., Korablyov, M., Precup, D., and Bengio, Y. Flow Network based Generative  
359 Models for Non-Iterative Diverse Candidate Generation. *Advances in Neural Information  
360 Processing Systems (NeurIPS)*, 2021.
- 361 [7] Bengio, Y., Deleu, T., Hu, E. J., Lahlou, S., Tiwari, M., and Bengio, E. GFlowNet Foundations.  
362 *arXiv preprint 2111.09266*, 2022.
- 363 [8] Bergen, V., Lange, M., Peidli, S., Wolf, F. A., and Theis, F. J. Generalizing RNA velocity to  
364 transient cell states through dynamical modeling. *BioRxiv preprint 820936*, 2019.
- 365 [9] Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. Neural Ordinary Differential  
366 Equations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- 367 [10] Cundy, C., Grover, A., and Ermon, S. BCD Nets: Scalable Variational Approaches for Bayesian  
368 Causal Discovery. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- 369 [11] Deleu, T., Góis, A., Emezue, C., Rankawat, M., Lacoste-Julien, S., Bauer, S., and Bengio,  
370 Y. Bayesian Structure Learning with Generative Flow Networks. *Uncertainty in Artificial  
371 Intelligence (UAI)*, 2022.
- 372 [12] Friedman, N., Murphy, K., and Russell, S. Learning the structure of dynamic probabilistic  
373 networks. *Uncertainty in Artificial Intelligence (UAI)*, 1998.
- 374 [13] Glymour, C., Zhang, K., and Spirtes, P. Review of Causal Discovery Methods Based on  
375 Graphical Models. *Frontiers in Genetics*, 2019.
- 376 [14] Granger, C. W. Investigating causal relations by econometric models and cross-spectral methods.  
377 *Econometrica: journal of the Econometric Society*, 1969.
- 378 [15] Ha, D., Dai, A. M., and Le, Q. V. Hypernetworks. *International Conference on Learning  
379 Representations (ICLR)*, 2017.
- 380 [16] Hashimoto, T. B., Gifford, D. K., and Jaakkola, T. S. Learning Population-Level Diffusions  
381 with Generative Recurrent Networks. *International Conference on Machine Learning (ICML)*,  
382 2016.
- 383 [17] Hauser, A. and Bühlmann, P. Characterization and greedy learning of interventional markov  
384 equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research (JMLR)*,  
385 13(1), 2012.
- 386 [18] He, Y.-B. and Geng, Z. Active learning of causal networks with intervention experiments and  
387 optimal designs. *Journal of Machine Learning Research (JMLR)*, 9, 2008.
- 388 [19] Huang, B., Zhang, K., Gong, M., and Glymour, C. Causal discovery and forecasting in  
389 nonstationary environments with state-space models. *International Conference on Machine  
390 Learning (ICML)*, 2019.

- 391 [20] Huang, B., Zhang, K., Zhang, J., Ramsey, J., Sanchez-Romero, R., Glymour, C., and Schölkopf,  
392 B. Causal discovery from heterogeneous/nonstationary data. *The Journal of Machine Learning*  
393 *Research (JMLR)*, 21, 2020.
- 394 [21] Huguet, G., Magruder, D. S., Tong, A., Fasina, O., Kuchroo, M., Wolf, G., and Krishnaswamy,  
395 S. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in Neural*  
396 *Information Processing Systems (NeurIPS)*, 2022.
- 397 [22] Huguet, G., Tong, A., Zapatero, M. R., Wolf, G., and Krishnaswamy, S. Geodesic Sinkhorn:  
398 Optimal transport for high-dimensional datasets. *arXiv preprint 2211.00805*, 2022.
- 399 [23] Itani, S., Ohannessian, M., Sachs, K., Nolan, G. P., and Dahleh, M. A. Structure learning in  
400 causal cyclic networks. *Proceedings of Workshop on Causality: Objectives and Assessment at*  
401 *NIPS 2008*, 2010.
- 402 [24] Kanehisa, M., Furumichi, M., Sato, Y., Ishiguro-Watanabe, M., and Tanabe, M. KEGG:  
403 Integrating viruses and cellular organisms. *Nucleic Acids Research*, 49, 2021.
- 404 [25] Karlebach, G. and Shamir, R. Modelling and analysis of gene regulatory networks. *Nature*  
405 *reviews Molecular cell biology*, 9, 2008.
- 406 [26] Ke, N. R., Bilaniuk, O., Goyal, A., Bauer, S., Larochelle, H., Pal, C., and Bengio, Y. Learning  
407 neural causal models from unknown interventions. *arXiv preprint 1910.01075*, 2019.
- 408 [27] La Manno, G., Soldatov, R., Zeisel, A., Braun, E., Hochgerner, H., Petukhov, V., Lidschreiber,  
409 K., Kastrioti, M. E., Lönnerberg, P., Furlan, A., Fan, J., Borm, L. E., Liu, Z., van Bruggen,  
410 D., Guo, J., He, X., Barker, R., Sundström, E., Castelo-Branco, G., Cramer, P., Adameyko, I.,  
411 Linnarsson, S., and Kharchenko, P. V. RNA velocity of single cells. *Nature*, 560, 2018.
- 412 [28] Lacerda, G., Spirtes, P. L., Ramsey, J., and Hoyer, P. O. Discovering cyclic causal models by  
413 independent components analysis. *arXiv preprint 1206.3273*, 2012.
- 414 [29] Lachapelle, S., Brouillard, P., Deleu, T., and Lacoste-Julien, S. Gradient-Based Neural DAG  
415 Learning. *International Conference on Learning Representations (ICLR)*, 2020.
- 416 [30] Liu, Q. and Wang, D. Stein variational gradient descent: A general purpose bayesian inference  
417 algorithm. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- 418 [31] Lopez, R., Hütter, J.-C., Pritchard, J. K., and Regev, A. Large-scale differentiable causal  
419 discovery of factor graphs. *Advances in Neural Information Processing Systems (NeurIPS)*,  
420 2022.
- 421 [32] Lorch, L., Rothfuss, J., Schölkopf, B., and Krause, A. DiBS: Differentiable Bayesian Structure  
422 Learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- 423 [33] Madan, K., Rector-Brooks, J., Korablyov, M., Bengio, E., Jain, M., Nica, A., Bosc, T., Bengio,  
424 Y., and Malkin, N. Learning gflownets from partial episodes for improved convergence and  
425 stability. *arXiv preprint 2209.12782*, 2022.
- 426 [34] Malkin, N., Jain, M., Bengio, E., Sun, C., and Bengio, Y. Trajectory Balance: Improved Credit  
427 Assignment in GFlowNets. *Advances in Neural Information Processing Systems (NeurIPS)*,  
428 2022.
- 429 [35] Monti, R. P., Zhang, K., and Hyvärinen, A. Causal discovery with general non-linear relation-  
430 ships using non-linear ica. *Uncertainty in Artificial Intelligence (UAI)*, 2020.
- 431 [36] Mooij, J. M., Janzing, D., Heskes, T., and Schölkopf, B. On causal discovery with cyclic  
432 additive noise models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2011.
- 433 [37] Mooij, J. M., Janzing, D., and Schölkopf, B. From Ordinary Differential Equations to Structural  
434 Causal Models: The deterministic case. *Uncertainty in Artificial Intelligence (UAI)*, 2013.
- 435 [38] Mooij, J. M., Magliacane, S., and Claassen, T. Joint causal inference from multiple contexts.  
436 *The Journal of Machine Learning Research (JMLR)*, 21, 2020.

- 437 [39] Murphy, K. P. Active learning of causal bayes net structure. Technical report, UC Berkeley,  
438 2001.
- 439 [40] Pamfil, R., Sriwattanaworachai, N., Desai, S., Pilgerstorfer, P., Georgatzis, K., Beaumont, P.,  
440 and Aragam, B. Dynotears: Structure learning from time-series data. *International Conference*  
441 *on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- 442 [41] Pearl, J. *Causality*. Caimbridge University Press, second edition, 2009.
- 443 [42] Peters, J., Janzing, D., and Schölkopf, B. *Elements of causal inference: foundations and*  
444 *learning algorithms*. The MIT Press, 2017.
- 445 [43] Peters, J., Bauer, S., and Pfister, N. Causal models for dynamical systems. *Probabilistic and*  
446 *Causal Inference: The Works of Judea Pearl*, 2022.
- 447 [44] Pratapa, A., Jalihal, A. P., Law, J. N., Bharadwaj, A., and Murali, T. Benchmarking algorithms  
448 for gene regulatory network inference from single-cell transcriptomic data. *Nature methods*, 17,  
449 2020.
- 450 [45] Qiu, X., Zhang, Y., Martin-Rufino, J. D., Weng, C., Hosseinzadeh, S., Yang, D., Pogson, A. N.,  
451 Hein, M. Y., Hoi (Joseph) Min, K., Wang, L., Grody, E. I., Shurtleff, M. J., Yuan, R., Xu, S.,  
452 Ma, Y., Replogle, J. M., Lander, E. S., Darmanis, S., Bahar, I., Sankaran, V. G., Xing, J., and  
453 Weissman, J. S. Mapping transcriptomic vector fields of single cells. *Cell*, 185, 2022.
- 454 [46] Riba, A., Oravec, A., Durik, M., Jiménez, S., Alunni, V., Cerciat, M., Jung, M., Keime, C.,  
455 Keyes, W. M., and Molina, N. Cell cycle gene regulation dynamics revealed by RNA velocity  
456 and deep-learning. *Nature Communications*, 13, 2022.
- 457 [47] Saelens, W., Cannoodt, R., Todorov, H., and Saeys, Y. A comparison of single-cell trajectory  
458 inference methods. *Nature Biotechnology*, 37, 2019.
- 459 [48] Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Gould, J., Liu,  
460 S., Lin, S., Berube, P., Lee, L., Chen, J., Brumbaugh, J., Rigollet, P., Hochedlinger, K., Jaenisch,  
461 R., Regev, A., and Lander, E. S. Optimal-Transport Analysis of Single-Cell Gene Expression  
462 Identifies Developmental Trajectories in Reprogramming. *Cell*, 176, 2019.
- 463 [49] Spirtes, P. and Glymour, C. An algorithm for fast recovery of sparse causal graphs. *Social*  
464 *Science Computer Review*, 9, 1991.
- 465 [50] Tank, A., Covert, I., Foti, N., Shojaie, A., and Fox, E. Neural Granger Causality. *IEEE*  
466 *Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- 467 [51] Tong, A., Huang, J., Wolf, G., van Dijk, D., and Krishnaswamy, S. TrajectoryNet: A Dynamic  
468 Optimal Transport Network for Modeling Cellular Dynamics. *International Conference on*  
469 *Machine Learning (ICML)*, 2020.
- 470 [52] Tong, S. and Koller, D. Active learning for structure in bayesian networks. *International Joint*  
471 *Conference on Artificial Intelligence (IJCAI)*, 17, 2001.
- 472 [53] Zhang, K., Huang, B., Zhang, J., Glymour, C., and Schölkopf, B. Causal discovery from nonsta-  
473 tionary/heterogeneous data: Skeleton estimation and orientation determination. *International*  
474 *Joint Conference on Artificial Intelligence (IJCAI)*, 2017.
- 475 [54] Zheng, X., Aragam, B., Ravikumar, P., and Xing, E. P. DAGs with NO TEARS: Continuous  
476 Optimization for Structure Learning. *Advances in Neural Information Processing Systems*  
477 *(NeurIPS)*, 2018.

## 478 Supplementary Material

### 479 A Proof of Proposition 1

480 Proposition 1 calculates the number of admissible structure graphs for a linear ODE system with  
481 correlated variables. We will first show the general case this is  $\prod_{i \in d} (2^{m_i} - 1)^{\mathcal{G}_i^1}$ , then analyze the  
482 case of an  $L^0$  penalty on the edges of  $G$ , which reduces the size of the set of admissible graphs to  
483  $\prod_i (m_i)^{\mathcal{G}_i^1}$ .

484 *Proof.* Consider an identifiable linear system  $\frac{dx}{dt} = Ax$  where we directly observe  $(x, \frac{dx}{dt})$  with  $\mathcal{G}^*$   
485 identifiable. Then the system with  $m = \mathbf{1}^d$  has exactly one admissible graph by definition. For each  
486 node, we analyze its set of child nodes in  $\mathcal{G}$ , i.e.  $c(u) = \{v \in \mathcal{V} \text{ s.t. } u \rightarrow v \in \mathcal{G}\}$ . For an identifiable  
487 system, each child  $v$  must have an incoming edge from its parent.

488 Next, we consider the process of adding a correlated variable, i.e. consider the situation of w.l.o.g.  
489 consider  $m = (s, 1, 1, \dots, 1)$  for some  $s > 1$ . Then for each child of  $c_j(v_1)$ , there are now  $s$  possible  
490 parents. This has multiplied the number of possible graphs by  $2^s - 1$ . Since each element of  $m$  is  
491 independent, this leads to the first statement, i.e.  $|\mathcal{G}'| = \prod_{i \in d} (2^{m_i} - 1)^{\mathcal{G}_i^1}$ .

492 Under an  $L^0$  penalty, then we constrain the possible graphs to  $s$  different graphs, where each  
493 child node picks exactly one of the  $s$  possible parents. This leads to the second statement,  $|\mathcal{G}'| =$   
494  $\prod_i (m_i)^{\mathcal{G}_i^1}$   $\square$

### 495 B Experimental Details

#### 496 B.1 Single Cell Dataset Preprocessing

497 We start with the processed data from [46]. We first filter it applying steps from the ScVelo tutorial.  
498 We then sub-select the genes of interest and use the ‘‘Ms’’ and ‘‘velocity’’ layers, which we normalize  
499 to mean zero standard deviation one for the states and scale the  $dx$  with the same parameters.

#### 500 B.2 Baselines for Bayesian Dynamic Structure Learning

501 Existing Bayesian structure learning methods are typically constrained to learning DAGs. Tem-  
502 poral information about the the dynamic relationships amongst variables in a system can help  
503 alleviate this constraint. DiBS and BCD-Nets are two state-of-the-art Bayesian structure learn-  
504 ing approaches for static systems. We apply versions of DiBS and BCD-Nets such that they  
505 are applicable in our Bayesian dynamic structure learning framework for learning cyclic graph  
506 structure from dynamic data. We use the approach taken in DiBS and parameterize the distribu-  
507 tion over graphs as  $P_{\alpha_t}(G|Z) = \prod_i \prod_j P_{\alpha_t}(G_{ij}|Z_{ij})$ , where  $Z = U^T V$ ,  $U, V \in \mathbb{R}^{k \times d}$ . Here  
508  $P_{\alpha_t}(G_{ij} = 1|Z_{ij}) = \sigma(\alpha_t Z_{ij})$ ,  $\sigma(x) = 1/(1 + e^{-\alpha_t x})$ , and  $\alpha_t = \alpha c(t)$  ( $t$  denotes the training  
509 iteration. We use  $c(t) = \sqrt{t}$ ). As  $t \rightarrow \infty$ ,  $P_{\alpha_t}(G|Z) \rightarrow \delta(G|Z)$ . In DiBS, Stein variational  
510 gradient decent (SVGD) [30] is used to iteratively transport particles  $Z$  to learn the target distribu-  
511 tion. Following from the above parameterization, we implement a version of BCD-Nets by treating  
512  $U \sim \mathcal{N}(\mu_u, \sigma_u^2)$ ,  $V \sim \mathcal{N}(\mu_v, \sigma_v^2)$ , and learning  $\mu_u, \mu_v, \sigma_u$ , and  $\sigma_v$ . Since our framework uses  
513 dynamic data, we incorporate DiBS and BCD-Nets within the framework (labeled DynDiBS and  
514 DynBCD, respectively) to leverage dynamic information for Bayesian structure learning of cyclic  
515 graphs.

#### 516 B.3 Hyper-parameters for Baselines

517 For both DynBCD and DynDiBS we use  $k = d$  across datasets. Since DynDiBS is an ensemble  
518 based method, we use 1024 samples for the linear and non-linear synthetic systems and 1000 samples  
519 for the cellular system (both training and evaluation). Since DynBCD is a variational approach  
520 and doesn’t require parallelized model ensembles, we use a large quantity of samples for training  
521 and evaluation. In the case of the cellular system, since there is a significantly smaller quantity of  
522 admissible graphs, we use less samples for DynBCD. We use graph sparsity regularization denoted by  
523  $\lambda_0$  and a temperature parameter  $T$  that scales the magnitude of the likelihood (e.g.  $\frac{1}{T^2} \text{MSE}(dx, \widehat{dx})$ ).

524 In Table 5 and Table 6 we outline the hyper-parameters we found to yield the most competitive results.  
 525 We use grid search to tune DynBCD and DynDiBS. All baselines are trained for 1000 epochs.

Table 5: Hyper-parameters for DynBCD. We define learning rate as  $\epsilon$ .

<b>Linear System</b>					
<b>Model</b>	$\epsilon$	$\lambda_0$	$T$	$\alpha$	samples
$\ell$ -DynBCD	0.0001	0.001	0.01	0.1	5000
$h$ -DynBCD	0.0001	0.0025	0.01	2	2000
<b>Non-linear System</b>					
<b>Model</b>	$\epsilon$	$\lambda_0$	$T$	$\alpha$	samples
$\ell$ -DynBCD	0.00005	0.01	0.01	2	5000
$h$ -DynBCD	0.0001	0.001	0.01	1	2000
<b>Cellular System</b>					
<b>Model</b>	$\epsilon$	$\lambda_0$	$T$	$\alpha$	samples
$\ell$ -DynBCD	0.0001	0.001	0.05	0.05	1000
$h$ -DynBCD	0.00001	0.0005	0.1	2	1000

Table 6: Hyper-parameters for DynDiBS. We define learning rate as  $\epsilon$ .

<b>Linear System</b>					
<b>Model</b>	$\epsilon$	$\lambda_0$	$T$	$\alpha$	$\gamma$
$\ell$ -DynDiBS	0.0025	500	0.01	0.0001	3000
$h$ -DynDiBS	0.0001	3	0.01	0.0001	10000
<b>Non-linear System</b>					
<b>Model</b>	$\epsilon$	$\lambda_0$	$T$	$\alpha$	$\gamma$
$\ell$ -DynDiBS	0.001	10	0.01	0.0001	3000
$h$ -DynDiBS	0.0001	0.1	0.01	0.0001	10000
<b>Cellular System</b>					
<b>Model</b>	$\epsilon$	$\lambda_0$	$T$	$\alpha$	$\gamma$
$\ell$ -DynDiBS	0.0025	1	0.05	0.0001	3000
$h$ -DynDiBS	0.00001	0.1	0.01	0.01	3000

526 We note that when evaluation on validation and test data for Bayes-SHD and AUC metrics, we hard  
 527 threshold  $P_{\alpha_t}(G|Z)$ . We find that through training this the final  $\alpha_t$  is typically small enough in  
 528 magnitude such that  $P_{\alpha_t}(G|Z)$  does not yield a full threshold of  $Z$ . To this end, we select large  $\alpha_t$   
 529 when computing the KL metric to mimic hard threshold behaviour as experienced during training.  
 530 We use  $\alpha_t = 1 \times 10^4$  for DynBCD and  $\alpha_t = 1 \times 10^8$  for DynDiBS methods, respectively. In  
 531 DynDiBS The parameter  $\gamma$  helps control separation of particles  $Z$  during training. In general, we  
 532 found DynBCD and DynDiBS baselines are challenging to train and to find hyper-parameter settings  
 533 with good performance. In part, we believe this is due to the numerous hyper-parameters required to  
 534 tune as well as the general difficulty of the objective.

#### 535 B.4 Neural Network Architectures and Hyper-parameters

536 We parameterize  $P_F(s_i|s_{i-1}; \psi)$  and  $h_\phi$  with MLP architectures.  $P_F(s_i|s_{i-1}; \psi)$  takes the current  
 537 state as input and firsts computes common representations using a 3 layer MLP. Then a 2 layer  
 538 MLP with a softmax output activation takes the representations as input and outputs a distribution  
 539 over possible actions. The latter MLP is used to parameterize one head for each distribution  
 540  $P_F(s_i|s_{i-1}; \psi)$ . We use a hidden unit dimension of 128 and leaky rectified linear unit (Leaky ReLU)  
 541 activation functions for the  $P_F(s_i|s_{i-1}; \psi)$  MLP architecture. We use a uniform backward policy  
 542 for  $P_B(s_{i-1}|s_i; \psi)$ . To parameterize  $h_\phi$ , we use a 3 layer MLP with hidden layer dimensions of  
 543  $\{64, 64, 64\}$  and exponential linear unit activations (ELU). We consider two parametrizations for  $f_\theta$ :

Table 7: Hyper-parameters for DynGFN. We define learning rate as  $\epsilon$ ,  $m_{train}$  as number of training samples, and  $m_{eval}$  the number of evaluation samples.

Linear System					
Model	$\epsilon$	$\lambda_0$	$T$	$m_{train}$	$m_{eval}$
$\ell$ -DynGFN	0.0001	100	0.01	1024	5000
$h$ -DynGFN	0.00001	100	0.005	256	3000

Non-linear System					
Model	$\epsilon$	$\lambda_0$	$T$	$m_{train}$	$m_{eval}$
$\ell$ -DynGFN	0.0001	150	0.01	1024	5000
$h$ -DynGFN	0.00001	150	0.005	256	3000

Cellular System					
Model	$\epsilon$	$\lambda_0$	$T$	$m_{train}$	$m_{eval}$
$\ell$ -DynGFN	0.00005	45	0.01	1024	1000
$h$ -DynGFN	0.0001	10	0.1	1024	1000

544 single linear parameters, i.e  $dx = \theta x$ , and a single hidden layer neural network  $dx = f_\theta(x)$ . We use  
 545 these parametrizations to model linear and non-linear node-wise parent-child structural equations,  
 546 where  $x \in \mathbb{R}^d$  are the node-wise input observations.

## 547 B.5 Hyper-parameters for DynGFN

548 DynGFN requires setting a variety of hyper-parameters that lead to different trade offs in model  
 549 performance. In particular,  $\lambda_0$  (sparsity encouragement for identified graphs), a temperature parameter  
 550  $T$  that scales the magnitude of the reward likelihood (e.g.  $\frac{1}{T^2} \text{MSE}(dx, \widehat{dx})$ ), learning rate  $\epsilon$ , softmax  
 551 tempering  $c$  (we always use a cosine schedule for  $c$ , with a discrete period of 5 epochs), and number  
 552 of training epochs. In our experiments we select hyper-parameter values that lead to competitive  
 553 performance (this pertains to  $\ell$ -DynGFN and  $h$ -DynGFN models) by observing performance over a  
 554 few values. We outline the selected hyper-parameters for each respective model in Table 7. Due to  
 555 computational limits, we use less training samples than evaluation samples for DynGFN. We train  
 556 DynGFN for 1000 epochs.

## 557 B.6 GFlowNet Exploration vs. Exploitation

558 The general procedure for training GFlowNets is inspired from reinforcement learning where the  
 559 primary objective is to learn a stochastic policy  $\pi(a|s)$  to sample actions from an action space given  
 560 a current state. In our setting, the action space represents possible locations where an additional edge  
 561 can be placed to an existing graph and each state is represented by a current graph. Since under this  
 562 training procedure we are sampling from the GFlowNet policy  $P_F(s_i|s_{i-1}; \psi)$  at every iteration then  
 563 attributing a reward associated to the sampled state/graph, the policy is susceptible to exploitation: if  
 564  $P_F(s_i|s_{i-1}; \psi)$  samples a graph(s) with a high reward, it becomes easy for the policy to focus on  
 565 sampling said graphs since they yield high reward. To alleviate this we encourage exploration using  
 566 softmax tempering on our stochastic policy, by multiplying the logits of our forward policy by  $1/c$   
 567 before applying the softmax function. A larger  $c$  flattens the stochastic policy such that exploration  
 568 within the action space is encouraged. However, setting the parameters  $c$  is challenging and there  
 569 exists a trade-off between exploring and exploiting the stochastic policy during optimization. We  
 570 address this by using a cosine schedule for  $c$  such that  $1 \leq c \leq 1000$ . We treat the period of the  
 571 cosine schedule as a hyper-parameter.

## 572 B.7 Discussion of Training Dynamics

573 GFlowNets are a relatively recent class of models that can be challenging to optimize. We discuss  
 574 some of the challenges with training them especially in the context of a learned energy function.  
 575 We observed that in settings where the energy reward is fixed and we could proportionally penalize  
 576 missing edges as well as the addition of incorrect edges (e.g.  $\ell$ -DynGFN), we were able to better learn

577 posteriors over admissible graphs over models that require sparse priors and/or trainable energies.  
578 This suggests that DynGFN may be limited by an inadequate energy reward. However, we found  
579 training DynGFN with a trainable energy function challenging since the GFlowNet stochastic policy  
580 depends on the rewards, and vice versa. Further investigation and experimentation into this alternating  
581 optimization procedure is required.

## 582 **B.8 Evaluating Quality of Learned Posteriors**

583 Using our indeterminacy model defined in section 5, we can determine  $P(G^*)$  for a given set of  
584 correlated variables, where  $G^*$  denotes the set of true equally admissible structures. Here we assume  
585  $P(G^*)$  is a uniform distribution over  $G^*$  and determine the KL between  $Q(G)$  and  $P(G^*)$ . We  
586 compute the KL considering only the probabilities of a trained models to generate all structures in  
587  $G^*$ , i.e.  $Q(G^*)$ . This is due to the computational constraints for calculating the KL for DynGFN  
588 since even for sparse graphs of moderate size this is a combinatorial computation. Nonetheless,  
589 this approach allows us to directly compare the learned posteriors  $Q(G)$  to a ground truth discrete  
590 distribution over structure  $G$  to evaluate the effectiveness of Bayesian structure learning approaches.

## 591 **B.9 Implementation Details**

592 Our model is implemented in Pytorch and Pytorch Lightning and is available at <https://github.com/anonymous/anonymous>. Models were trained on a heterogeneous mix of HPC clusters for a  
593 total of ~1,000 GPU hours primarily on NVIDIA RTX8000 GPUs.  
594

## 595 **C Additional Details**

### 596 **C.1 Single-cell Biology and Gene Regulatory Network Inference**

#### 597 **C.1.1 Gene Regulatory Networks and Cell Dynamics**

598 One dynamical system of interest is that of cells. Cellular response to environmental stimuli or genetic  
599 perturbations can be modelled as a complex time-varying dynamical system [16, 1]. In general,  
600 dynamical system models are a useful tool for downstream scientific reasoning. In this work we are  
601 primarily interested in identifying the underlying cell dynamics from data. A reasonable model for cell  
602 dynamics is as a stochastic dynamical system with many, possibly unobserved, components. There  
603 are many data collection models for gaining insight into this system from single-cell RNA-sequencing  
604 data. We will primarily focus on RNA velocity type methods, where both  $x$  and an estimate of  $dx$  are  
605 available in each cell, but note that there are other assumptions to infer dynamics and regulation such as  
606 pseudotime-based methods [47, 1], and optimal transport methods [16, 48, 51, 21, 22]. After learning  
607 a possible explanatory regulation, this is used in downstream tasks, but the resulting conclusions  
608 drawn from these models are necessarily conditional on the inferred regulation. Motivated by gene  
609 regulatory networks, we explicitly model uncertainty over graphs which allows us to propagate the  
610 resulting uncertainty to downstream conclusions.

#### 611 **C.1.2 Learning Gene Regulatory Networks From Single-cell Data**

612 Single-cell transcriptomics has an interesting property in that from a single measurement we can  
613 estimate both the current state  $x$  and the current velocity  $dx$ . Because mechanistically RNA undergoes  
614 a splicing process, we can measure the quantities of both the unspliced (early) and spliced (late) RNA  
615 in the cell. From these two quantities we can estimate the current RNA content for each gene and the  
616 current transcription rate. There exist many models for denoising and interpreting this data [27, 8, 45].  
617 Furthermore, there exist more elaborate measurement techniques to extract more accurate velocity  
618 estimates [45]. The fact that we have an estimate of the current velocity is exceptionally useful for  
619 continuous time structure discovery because it allows us to avoid explicitly unrolling the dynamical  
620 system.

621 Learning the underlying causal structure from data is one of the open problems in biology. There are  
622 many works that attempt to learn the effect of a change in a gene, or the addition of a drug. These  
623 works often build models that directly predict the outcome of an intervention. This may be useful for  
624 certain applications, but often does not generalize well out of distribution. We would like to learn a

625 model of the underlying instantaneous dynamics that give rise to effects at longer time scales. This  
 626 approach has a number of advantages. (1) it is closer to the mechanistic model; it may be easier to  
 627 learn a model of the instantaneous dynamics rather than the dynamics over long time scales (details  
 628 in Appendix C.2). (2) One model can be trained and applied to data from many sources including  
 629 RNA-velocity, Pseudotime, Single-cell time series, and steady state perturbational data. (3) The  
 630 instantaneous graph may be significantly sparser (and therefore easier to learn) than the summary  
 631 graph or the equilibrium graph.

### 632 C.1.3 Further Details on Experiments with Single-cell RNA-velocity Data

633 The process of Eukaryotic cell division can be divided into four well regulated stages based on the  
 634 phenotype, Gap 1 ( $G_1$ ), Synthesis ( $S$ ), Gap 2 ( $G_2$ ), and Mitosis ( $M$ ). This process is a good starting  
 635 point for GRN discovery as it is (1) relatively well understood, (2) deterministic, and (3) well studied  
 636 with plentiful data. While there is an underlying control loop controlling the progression of the cell  
 637 cycle, there are many other genes that also change during this cycle. To rediscover the true control  
 638 process from data we must disentangle the true causal genes from the downstream correlated genes.  
 639 This may become very difficult when we only observe dynamics at longer time scales. We hide a cell  
 640 cycle regulator among two downstream genes that are highly correlated (Spearman  $\rho > 0.75$ ) and  
 641 test whether we can model the Bayesian posterior – namely that we are uncertain about which of the  
 642 three genes (Cdc25A, Mcm2, or Mcm5) is the true causal parent of Cdk1.

## 643 C.2 Instantaneous Graph and Long-horizon Graph

644 The graph recovered depends on the time scale considered. We make a distinction between the  
 645 conditional structure of the graph based on the time scale. Consider a system governed by  $\frac{dx(t)}{dt} =$   
 646  $f(x(t))$ . We define the instantaneous graph as:

$$647 \quad g_{i,j} := \cup_x \mathbf{1} \left( \frac{df(x)_j}{dx_i} \right), \quad i = 1, \dots, d, j = 1, \dots, d. \quad (6)$$

647 Here,  $\mathbf{1}(z) = 1$  if  $z \neq 0$ , otherwise  $\mathbf{1}(z) = 0$ . We can define the temporal summary of the system  
 648 drift  $f$  as:

$$649 \quad F(T, x) = x_0 + \int_0^T f(x(t)) dt. \quad (7)$$

649 Then, we can define a long-horizon graph over long time-scales as:

$$650 \quad G_{i,j} := \cup_x \mathbf{1} \left( \frac{\partial F(T, x)_j}{\partial x_i} \right), \quad i = 1, \dots, d, j = 1, \dots, d. \quad (8)$$

650 If we are integrating the drift  $f$  over long time-scales, the long-horizon graph may be less sparse  
 651 than the instantaneous graph. In cellular systems, this equates to observing cell dynamics over long  
 652 time-scales, in turn observing increase quantities of correlations between variables of the system.  
 653 Thus, trying to delineate the instantaneous dynamics from long time-scales may be difficult depending  
 654 on the underlying system dynamics.

## 655 C.3 Further Discussion on Future Work

656 Although DynGFN is currently limited to smaller systems, we foresee approaches that would enable  
 657 some degree of scaling DynGFN to larger systems. One approach is to leverage biological information  
 658 of known gene-gene connections as a more informative prior for DynGFN. Currently, DynGFN learns  
 659 a forward stochastic policy for  $Q(G|D)$  starting from an initialized state  $s_0$  of all zeros. Instead, we  
 660 can define  $\tilde{s}_0$  using a prior of high confidence biological connections and sequentially add edges  
 661 starting from this new initial state. This would reduce the number of possible structures DynGFN  
 662 would need to search over, thus improving the potentially scalability of DynGFN. Another approach  
 663 is to learn structure between sets of genes (variables) rather than single genes. Since GRNs are  
 664 generally very sparse, it makes sense to group genes in sets. Consequently, we can then learn structure  
 665 between these grouped genes, rather than just individual genes. In turn, DynGFN can explore/learn  
 666 the structure over a smaller space while effectively capturing structure between a significantly larger  
 667 set of genes. To group genes, we would use prior biological information, either from existing literature  
 668 or expert domain knowledge.

669 In this work we exploit the use of a minimal prior, i.e.  $L^0$  sparsity prior, for learning Bayesian  
670 dynamic structure between variables. In general, the aforementioned approaches for scaling DynGFN  
671 to larger systems involve the use of more informative priors on  $G$ . Although we mention two ways  
672 we foresee approaching this in the biological context of GRN inference, the general approach of  
673 using more informative priors can help scale DynGFN to larger systems across applications.

#### 674 **C.4 Broader Impacts**

675 While it is important to acknowledge the potential risks of drawing incorrect scientific conclusions  
676 due to incorrect assumptions, our work embraces a Bayesian perspective for structure learning. A  
677 key component of our work is to account for uncertainty within our method, aiming to minimize the  
678 chances of incorrect conclusions. It is important to note that the accuracy of conclusions relies on  
679 applying the method in settings that align with the underlying assumptions, such as causal sufficiency  
680 and the use of dynamic observational data. By adhering to these guidelines, our approach holds  
681 promise for producing robust and reliable scientific outcomes.