Proceedings of Machine Learning Research vol 236:209–242, 2024 3rd Conference on Causal Learning and Reasoning

Bicycle: Intervention-Based Causal Discovery with Cycles

Martin Rohbeck MARTIN.ROHBECK@DKFZ.DE
German Cancer Research Center, Heidelberg University

Brian Clarke BRIAN.CLARKE@DKFZ.DE
German Cancer Research Center

Katharina Mikulik KATHARINA.MIKULIK@DKFZ.DE
German Cancer Research Center, Heidelberg University

Alexandra Pettet ALEXANDRA.X.PETTET@GSK.COM
GSK

Oliver Stegle OLIVER.STEGLE@DKFZ.DE*
German Cancer Research Center, European Molecular Biology Laboratory, Wellcome Sanger Institute

Kai Ueltzhöffer KAI.UELTZHOEFFER@EMBL.DE*
European Molecular Biology Laboratory, German Cancer Research Center, GSK

*CORRESPONDING AUTHOR

Editors: Francesco Locatello and Vanessa Didelez

Abstract

While a growing number of algorithms for causal discovery of directed acyclic graphs from observational and interventional data have been proposed, the robust identification of *cyclic causal graphs* in particular remains an open problem. Solutions to this challenge would have a considerable impact in various application domains, including single-cell genomics, where gene regulatory networks are known to contain feedback loops. Recent work has shown promise in addressing this challenge by describing the expression states in a population of cells as the steady-state solution of a stochastic dynamical system. However, the current formulation cannot account for information on interventions in the population, and consequently, it ignores the associated causal inductive biases, which are key assets to fully exploit the potential of available data for obtaining meaningful results and improving identifiability. In this work, we propose Bicycle, a method which (i) infers cyclic causal relationships from i.i.d. data, (ii) explicitly accounts for information on the perturbation state of cells by a realization of the independent causal mechanism principle and (iii) models causal effects in a latent space rather than on observed data. We benchmark Bicycle in the context of existing approaches, demonstrating improved recovery of simulated causal graphs and improved out-of-distribution prediction performance on unseen perturbations in real single-cell datasets.

Keywords: Cyclic Graphs, Stochastic Dynamical Systems, Interventions, Gene Regulatory Network Inference

1. Introduction

The prediction of the effects of unseen interventions on a system is a crucial application across many fields, such as climate modeling (Adachi et al., 2017), economics (Romer and Romer, 2010) and drug discovery in biomedicine (Michoel and Zhang, 2023). In computational single-cell biology, a central question is to infer how genes regulate each others' expression (so-called gene regulatory networks (GRNs)), and to predict the effects of perturbations – either gene knock-downs or activations. Accurate GRN inference models for these tasks would open vast opportunities for cell-fate engineering, identifying complex disease mechanisms and prioritizing targets for potential therapeutic interventions (Nelson et al., 2015). However, current methods face critical limitations. In particular, there is a lack of models that simultaneously account for cyclic dependencies and can account for information on known perturbations.

Modern CRISPR-based perturbation screens allow for probing GRNs by perturbing thousands of genes in millions of cells, followed by measuring the effects on the expression of mRNA from all genes in each cell (Replogle et al., 2022; Datlinger et al., 2021; Norman et al., 2019). Nevertheless, since known protein-coding human genes number around 20,000, and since GRNs can differ among the roughly 200 cell types in humans, exhaustively perturbing all genes in all cell types is infeasible with current technology. This is especially true when considering combinatorial interventions, i.e., those targeting more than one gene at a time. Therefore, it is of great interest to develop computational methods that can learn from a subset of all possible interventions to predict the response of a cell to unseen interventions on one or more target genes. While black-box methods are achieving increasing proficiency at this task (Roohani et al., 2023), a higher degree of explainability is desired, especially in clinical settings. Thus, methods which identify a causal generative model of the data-generating process, i.e., the underlying gene regulatory logic, offer a strong combination of interpretability and robust predictive power on unseen interventions.

There exists a growing number of application-agnostic methods for causal discovery. Critically, many methods make strong assumptions on the acyclicity of the underlying causal graph (Brouillard et al., 2020; Lopez et al., 2022; Zheng et al., 2018; Lippe et al., 2022), an assumption that is violated in biological GRNs due the presence of feedback loops (Paige et al., 2015; Dunn et al., 2014; Freimer et al., 2022). The lack of robust, general-purpose inference schemes for cyclic GRNs has resulted in a large number of biology-driven methods that seek to unravel GRNs from single-cell omics data (Bravo González-Blas et al., 2023; Jiang et al., 2022; Kamal et al., 2023). While these methods have been demonstrated to yield relevant biological insights, they either rely on additional data modalities and prior information or make strong assumptions about the underlying mechanisms, which can result in biases towards well-studied pathways. Furthermore, these approaches are typically not based on a well-defined generative model of the data and instead employ heuristics to predict a cell's response to unseen perturbations. Attempts to apply formal models for causal inference that capture cyclic structures are currently scarce, primarily because it has proven difficult to adapt existing concepts to single-cell genomics. A notable exception is the recently proposed Dictys model (Wang et al., 2023), which models the distributions of cells as draws from the equilibrium distribution of an unobserved continuous stochastic dynamical system, which is coupled to observations via a technology-tailored noise model. However, a key limitation of the method is that it does not account for perturbations, which precludes meaningful applications in the context of high-throughput CRISPR perturbation screens.

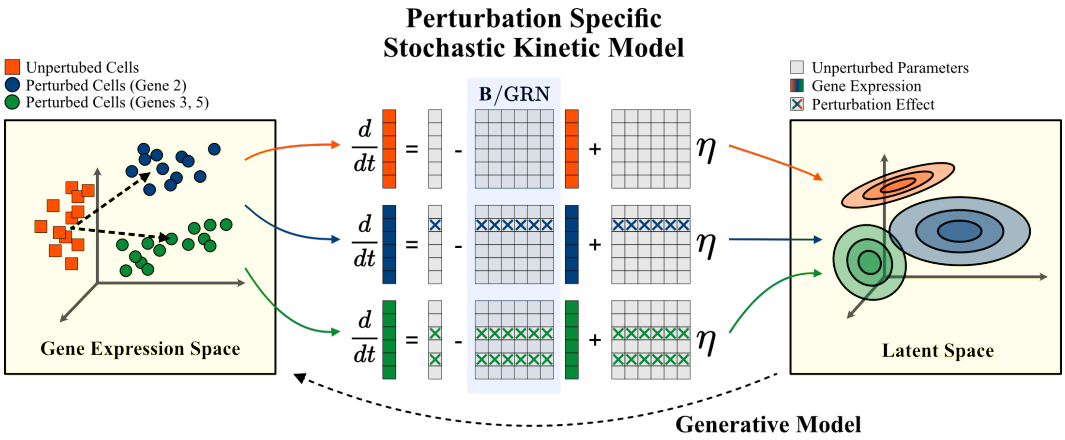

Figure 2: Given a set of observed single-cell samples in observational and interventional conditions (left), we aim to learn the causal effects between genes (the GRN matrix, light blue background, middle). Bicycle employs an intervention-specific stochastic kinetic model that can incorporate sparse changes in the model parameters based on the intervention target (middle). The steady-state solution of the model, a set of multivariate normal distributions, is modelled in a latent space (right).

Here, we propose *Bicycle*, a new model to infer causal graphs with cyclic structures from observational and interventional data (Fig. 1). Bicycle models the true (free of technical/measurement noise) causal relationships using latent variables described by the steady-state distribution of a dynamical system with unknown governing equations. A key innovation of Bicycle is that its stochastic differential equations (SDEs) are parameterized in a hierarchical fashion for each interventional condition. That is, the parameters of the SDE across conditions are identical to those of an unperturbed system, except for those genes that govern the evolution of direct interventional target variables (Fig. 2). This approach can be interpreted as an instance of the independent causal mechanisms principle (c.f., Scholkopf et al. (2021)). The model can unravel causal relationships and predict the effect of unknown interventions while providing a directly interpretable representation of the system. We assess Bicycle using both simulated and real data benchmarks, demonstrating improved causal structure recovery and prediction of perturbation effects of unseen interventions. The Bicycle package will be made publicly available upon publication.

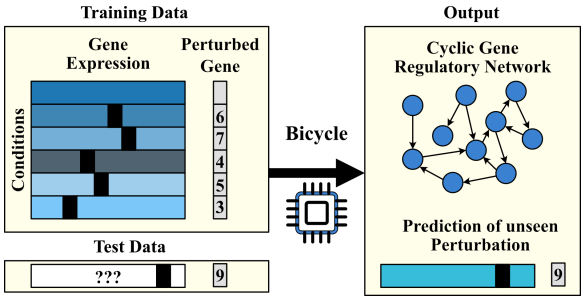

Figure 1: Bicycle uses data from multiple interventions to estimate the underlying causal graph and make predictions for unseen interventions.

2. Related Work and Background

The goal of causal discovery is to infer the underlying causal graph and the functional relationships between variables from observational or interventional data. In structural causal models (SCMs) or structural equation models (SEMs), one assumes a deterministic relationship between variables under the presence of noise. Due to their favorable statistical properties, acyclic causal models are well-studied, but they cannot model feedback loops. Therefore, methods that can account for and model cyclic causal graphs are of interest. Finally, many application-tailored methods for single-cell data eschew the causal discovery problem, instead relying on known biological knowledge to construct a graph of gene-gene interactions. In the following, we review existing approaches and provide examples of representative methods.

Acyclic Models Acyclic causal models can be identified up to their Markov equivalence class (MEC), which can be deduced given appropriate interventional data. Commonly, acyclic models can be classified as constraint-based, score-based or a combination thereof. Constraint-based algorithms, such as FCI, PC (Spirtes et al., 1993) or nonlinear ICP (Heinze-Deml et al., 2018) provide undirected graphs and make heavy use of conditional independence testing, leading to a large search space because the conditioning sets explode quickly with the number of independent variables. Scoring-based algorithms apply changes to the adjacency matrix to optimize a scoring function. A major breakthrough in this was the NOTEARS algorithm (Zheng et al., 2018), which led to multiple follow-up works (Lee et al., 2019; Brouillard et al., 2020; Lopez et al., 2022; Xue et al., 2023). However, NOTEARS recovers acyclic graphs characterized by a high level of *varsortability*, which quantifies the degree of overlap between the sequence of increasing marginal variance and the topological arrangement (Reisach et al., 2021). DCDI (Brouillard et al., 2020) and its more scalable successor DCD-FG (Lopez et al., 2022) model observational and interventional distributions by jointly learning a directed causal graph and deep neural networks to flexibly approximate the functional form of the invariant causal relationships between downstream genes and their regulators. However, these models rely on an explicit acyclicity constraint and therefore cannot model cyclic graphs. Lippe et al. (2022) present a highly scalable method, ENCO, to infer directed, acyclic causal graphs leveraging observational and interventional data. In contrast to NOTEARS, ENCO does not include an explicit acyclicity constraint (instead relying on the empirical observation that its learned sparse graphs rarely contain cycles) and learns the probability of edge existence and direction parameterized with two independent parameters.

Neural Network Models for Perturbation Data GEARS (Roohani et al., 2023) is a deep learning model that predicts the cellular response to genetic perturbations (both single and multi-gene interventions). It requires a knowledge graph of gene-gene relationships to prune the potential interactions between genes. In contrast to our method, GEARS does not provide an interpretable GRN but relies on a black-box, deep-learning-based architecture. Further methods to predict change in gene expression without uncovering the underlying GRN are PerturbNet (Yu and Welch, 2022) and CaML (Nilforoshan et al., 2023). PerturbNet is based on a conditional invertible neural network (cINN), i.e., conditional normalizing flows, which learns the mapping between noise and gene expression data. The model conditions the mapping on encoded perturbation data to allow for generating data for unseen perturbations. In contrast, CaML is a zero-shot learning approach that trains a meta-model to estimate the conditional average treatment effect (CATE) on individuals and interventions and can be applied to unseen individuals.

Simulation-based, Amortized Causal Discovery Recent work has shown promise in transferring the causal discovery problem to the setting of supervised deep learning by simulating known ground-truth data from generative models and then training flexible transformer architectures in the standard supervised setting to predict the causal graph (Ke et al., 2023a,b). While these simulation-based amortized inference approaches show promise in recovering causal structure which fits the distribution of the synthetic training data, in many fields, including molecular biology, there are no broadly accepted simulators to generate the synthetic data required to train the deep-learning model, as many aspects of the data-generating processes are still poorly understood.

Cyclic Causal Models There is a rich theoretical literature on cyclic causal models for observational steady-state data (Bongers et al., 2021; Mooij et al., 2013), which can be conceptualized as sampled from the equilibrium distribution of an underlying continuous or discrete-time stochastic dynamical system. LLC (Hyttinen et al., 2012) is a time-discrete model that infers cyclic dependencies given linear relations between the variables. LLC assumes a time-independent, constant noise covariance matrix that might contain nonzero off-diagonal elements (unobserved confounding) and can work with interventional, but not with purely observational data. NODAGS-Flow (Sethuraman et al., 2023) learns nonlinear cyclic causal graphical models from interventional data by making use of ideas from residual normalizing flows, but does not (i) account for different data likelihoods, (ii) model technical/measurement noise explicitly, (iii) allow confounders or (iv) model imperfect interventions. The closest work to ours, a steady-state, continuous-time linear dynamical model, and the one we build upon, is Dictys (Wang et al., 2023). Dictys is a dynamic GRN inference method that models latent, cyclic regulatory dynamics using a multivariate Ornstein–Uhlenbeck process, which it couples to a modality-specific technical noise model. However, it does not make use of perturbations, which are the workhorse of causal discovery. Instead, it requires prior knowledge to establish the presence and direction of causal effects in the GRN. Dictys represents a class of models called "continuous Lyapunov models" in the statistical literature (Varando and Hansen, 2020). While there are algorithms and identifiability results for estimating the causal graph from purely observational data (Dettling et al., 2022, 2023), to the best of our knowledge, we are the first to explicitly leverage perturbation data via the independent causal mechanism principle to discover the causal structure underlying a set of observational and perturbational data distributions using a continuous Lyapunov model, particularly when modeling the true expression state coupled to a modality-specific technical noise model.

3. Methods

In the following section, we first introduce the relevant notation and kinetic model in Sec. 3.1, before we describe our core methodological extensions: (i) the integration of interventional conditions into the kinetic model in Sec. 3.2, followed by (ii) an inference approach in Sec. 3.3.

3.1. Steady-State Model of Gene Regulation

We begin by following the notation used in Dictys (Wang et al., 2023) and use the GRN inference task as motivation. Nevertheless, our method is more generally applicable, with the terms "cells" and "genes" simply replaced with "samples" and "features," together with using an appropriate noise model to generate predicted observations from the underlying latent causal model. We note the terminology *transcription rate* of a gene to refer to the production rate of mRNA from that gene.

Assume a matrix $x = (x_{ng}) \in \mathbb{N}_0^{N \times G}$ of observed mRNA counts from G genes in N cells, i.e., x_{ng} corresponds to the number of mRNA molecules from gene g in cell n. The observation x_{ng} includes both biological and technical noise. We model the true expression state (free of technical noise) of an unperturbed cell, $z = (z_{ng}) \in \mathbb{R}^{N \times G}$, as a draw from the steady-state distribution of an Ornstein-Uhlenbeck process of the form

$$dz_g(t) = \left(\alpha_g + \sum_{g \neq h} \beta_{hg} z_h(t) - z_g(t) \right) dt + \sum_h \sigma_{gh} dW_h(t), \tag{1}$$

where we have omitted the cell index n for clarity. This represents the transcriptomic kinetics of a cell, with $\alpha_g \in \mathbb{R}^G$ denoting the net basal transcription rate for gene g, $\beta \in \mathbb{R}^{G \times G}$ the gene regulatory matrix, $W_h(t)$ a Wiener process modeling biological noise, and $\sigma = (\sigma_{gh}) \in \mathbb{R}^{G \times G}$ its effect on gene g.

Defining $B = \mathbb{I}_G - \beta^T$, it holds that iff the real parts of the eigenvalues of B are positive and $\sigma\sigma^T$ is positive definite, a non-singular analytical steady-state solution $p_{SS}(z) = \lim_{t \to \infty} p(z(t))$ to (1) exists. It follows a multivariate normal distribution with mean $\bar{z} = B^{-1}\alpha$ and covariance matrix ω, where B and ω are given by the solution to the continuous Lyapunov equation

$$B\omega + \omega B^T = \sigma\sigma^T, \tag{2}$$

i.e., $p_{SS}(z) = \mathcal{N}(\bar{z}, \omega)$ (Wang and Uhlenbeck, 1945). We assume the latent, unobserved gene expression state of the data corresponds to this equilibrium distribution, and we give the matrix β causal semantics consistent with recent advances in cyclic causal discovery (Bongers et al., 2021). That is, we interpret the parameter β_{hg} as the direct causal effect of the expression of gene h on the transcription rate of gene g.

This causal model of the latent expression state z_{ng} is linked to the measured expression count values x_{ng} for each cell n using a Multinomial likelihood function:

$$p(x_{n\cdot}|z_{n\cdot}) = \text{Multinomial}(x_{n\cdot}; p_{n\cdot}), \tag{3}$$

where $p_{n\cdot} = \text{Softmax}(z_{n\cdot})$. While we choose this likelihood function to accurately model count data, our model is not restricted to a specific likelihood and could also leverage other common distributions to model mRNA counts or other data modalities.

3.2. Implementing the Invariant Causal Mechanism Principle to Leverage Perturbation Data

Our main extension to the latent continuous Lyapunov model of Dictys comprises explicitly modeling multiple perturbation conditions, i.e., multiple interventional distributions. In each interventional condition $i \in 1, \ldots, I$, we assume a strict known subset $P^i \subset 1, \ldots, G$ of genes was targeted. To leverage the perturbation information to identify causal mechanisms and generate predictions for unseen perturbation conditions, we provide our model with a strong, causal inductive bias in the form of the invariant causal mechanisms principle (Scholkopf et al., 2021). Specifically, we assume sharing of most elements of α, β, σ between the interventional and unperturbed contexts – namely, all parameters governing the evolution of a gene which was not the direct target of a given perturbation, as illustrated in Figure 2. This hierarchy and parameter sharing, combined with the assumption

that the number of perturbations executed in a condition i is sparse, i.e., $|P^i| \ll G$, implement the idea of a sparse shift of independent causal mechanism as an inductive bias of our model.

The Bicycle model can be derived by constructing *condition-specific* dynamical parameters $\alpha^i \in \mathbb{R}^G$, $\beta^i \in \mathbb{R}^{G\times G}$, and $\sigma^i \in \mathbb{R}^{G\times G}$ in the following way:

$$\alpha_g^i = \begin{cases} \alpha_g & \text{if } g \notin P^i, \\ \hat{\alpha}_g^i & \text{else} \end{cases}, \qquad \beta_{hg}^i = \begin{cases} \beta_{hg} & \text{if } g \notin P^i, \\ \hat{\beta}_{hg}^i & \text{else} \end{cases}, \qquad \sigma_{gh}^i = \begin{cases} \sigma_{gh} & \text{if } g \notin P^i, \\ \hat{\sigma}_{gh}^i & \text{else} \end{cases},$$

where $\hat{\alpha}_g^i \in \mathbb{R}^G$ is the perturbed base transcription rate of gene g in condition i, and similarly for $\hat{\beta}_{hg}^i$ and $\hat{\sigma}_{gh}^i$. Now we can calculate the mean $\bar{z}^i = (B^i)^{-1}\alpha^i$ and covariance matrix ω^i of the multivariate normal distribution $p_{SS}(z^i) = \mathcal{N}(\bar{z}^i, \omega^i)$ for each intervention context i, based on the parameters of the SDE

$$dz_g^i(t) = \left(\alpha_g^i + \sum_{g\neq h} \beta_{hg}^i z_h^i(t) - z_g^i(t)\right) dt + \sum_k \sigma_{gh}^i dW_h(t), \tag{4}$$

analogously to in the unperturbed case. Denoting by i_n the perturbation carried out on cell n, we model the latent expression state of a perturbed cell n as a sample from this distribution, $z_n^{i_n} \sim p_{SS}(z^{i_n})$, as illustrated in Fig. 2. In our experiments, we assume perfect interventions and hence set $\hat{\beta}_{hg}^i = 0\ \forall h, g$. Additionally, we assume independent noise processes per gene, i.e., $\sigma_{gg}, \hat{\sigma}_{gg}^i > 0\ \forall g$ and $\sigma_{gh} = \hat{\sigma}_{gh}^i = 0\ \forall g \neq h$, which corresponds to assuming no latent confounders. Both of these assumptions can, in principle, be relaxed by additionally learning $\hat{\beta}_{hg}^i$, σ_{gh} and $\hat{\sigma}_{gh}^i$ from the data.

3.3. Inference

In the following, we index interventional conditions by i, with the unperturbed context denoted by $i = 0$. Thus, $\alpha^0 = \alpha$, $\bar{z}^0 = \bar{z}$, etc.

To infer the model parameters $\alpha^i, \beta^i, \sigma^i$, we maximize a loss function consisting of variational evidence lower bound (ELBO) with additional penalties. We perform maximum likelihood estimation of the parameters α^i, β^i and σ^i, and stochastic variational inference over the latent states z, using independent normal distributions as the variational family. Thus, $q(z \mid \mu_z, \sigma_z) = \prod_{N,G} \mathcal{N}(z_{ng}; \mu_{z_{ng}}, \sigma_{z_{ng}})$, where we optimize a mean vector μ_z and a vector of standard deviations σ_z for each data point. The complete loss function is:

Steady-State distribution of latent expression

$$\mathcal{L}^{\gamma,\xi,\lambda}(x, \alpha, \beta, \sigma, \hat{\alpha}, \hat{\sigma}, \mu_z, \sigma_z) = \frac{1}{N}\sum_{n=1}^{N}\left[\langle \log p(x_n \mid z_n)\rangle_{q(z_n\mid\mu_{z_n},\sigma_{z_n})} - \gamma \mathrm{KL}(q(z \mid \mu_{z_n}, \sigma_{z_n}) \,||\, \mathcal{N}(z; \bar{z}^{i_n}, \omega^{i_n}))\right]$$

$$- \frac{\xi}{I}\sum_{i=0}^{I} \|B^i\omega^i + \omega^i(B^i)^T - \sigma^i(\sigma^i)^T\|_2^2 - \lambda\,\|\beta\|_1 \tag{5}$$

Log-Likelihood of observed counts

Δ between LHS and RHS of Lyapunov Eq.

Sparsity

where $\gamma, \xi, \lambda > 0$ are scaling parameters and i_n denotes the intervention context of cell n. Since

the causal effect matrix can be assumed to be sparse in many applications, we employ a sparsity penalty. Depending on the dataset size (and GPU memory), we employ different approaches to infer ω^i: (i) for small datasets, we set $\xi = 0$ and solve the Lyapunov equation directly; (ii) for larger datasets, this is not feasible anymore, so we learn ω^i by minimizing the difference between the LHS and RHS of the Lyapunov equation. To ensure we infer a valid covariance matrix in the second case, we follow Wang et al. (2023) and parametrize ω as the sum of a positive diagonal matrix and a low-rank positive semidefinite matrix, i.e., $\omega = \omega_{\text{diag}} + \omega_{\text{fac}}\omega_{\text{fac}}^T$. where we select the column rank of ω_{fac} to be much smaller than the number of genes. We use Adam (Kingma and Ba, 2015) as our optimizer. We directly optimize a mean μ_{z_n} and variance parameter σ_{z_n} for each latent variable $z_n \in \mathbb{R}^G$ individually.

4. Experiments

We assessed Bicycle using synthetic data and real single-cell RNA-seq datasets with interventions. In each case, we carried out a hyperparameter search using the evidence lower-bound metric on a validation set. Additional information regarding these experiments, including the hyperparameter search grids and supplementary experiments, can be found in Appendix A.

4.1. Synthetic Data

First, to assess Bicycle's ability to infer causal relationships from perturbation data, we constructed synthetic datasets where the underlying causal graph is known. For comparison, we also considered State-of-the-Art (SOTA) methods for causal graph inference, LLC (which models cyclic graphs with linear SEMs) (Hyttinen et al., 2012), NOTEARS (which models acyclic graphs with linear SEMs) (Zheng et al., 2018) and NODAGS-Flow (which models acyclic graphs with nonlinear SEMs) (Sethuraman et al., 2023). As these approaches do not explicitly model count-based data, we considered a reduced variant of Bicycle for these experiments without its technical noise model, thereby allowing for a like-with-like comparison assuming normally distributed data in all models.

In the first set of experiments, we assessed the model's ability to (i) reconstruct the causal effect graph and its parameters β, and (ii) predict the response to unseen interventional conditions, which correspond to pairs of interventions. We sampled 20 intervention pairs at random for testing, and hence, either 0, 1, or 2 interventions had been seen during training in the form of individual interventions. We assumed perfect interventions and known target nodes. Dataset sizes scaled linearly with the number of interventions (250 samples per intervention; single node intervention), drawn from a graph with 10 nodes. To study the effect of additional observational data, we considered an augmented training dataset with 500 further observational (unperturbed) samples included for selected models and experiments, i.e., these models were trained on $k \cdot 250 + 500$ samples for k interventions. Simulated graphs were derived from an Erdős–Rényi $G(n, p)$ model, with an expected edge density of 2. Error bars are standard deviations from five repeated experiments.

We considered synthetic datasets derived from one of two alternative SEMs: an Ornstein-Uhlenbeck process (SEM1) as in Eq. (4) and a time-discrete linear SEM (SEM2), of the form

$$x(t) := U_k \beta x(t-1) + U_k \epsilon + c, \tag{6}$$

where $U_k \in \mathbb{R}^{G \times G}$ is an indicator matrix with $(U_k)_{gg} = 1$ if variable g was not perturbed and all other entries 0. The interventional distribution, $c \sim \eta \cdot \mathcal{N}(0, \mathbb{I}_G)$, representing the perturbation effect

(with different effect sizes η), is only nonzero for the perturbed gene g, and $\epsilon \sim 0.5 \cdot \mathcal{N}(0, \mathbb{I}_G)$, i.e., no latent confounding is modeled. More experimental details can be found in Appendix A. Compared to SEM1, in SEM2 we have a time-invariant effect of noise and perturbation throughout the convergence to the equilibrium.

If hyperparameters needed to be tuned for a specific model, we retained 20% of the data as holdout samples. Owing to the imbalanced number of existing vs. missing edges (approx. 10 vs 90), we quantified the graph reconstruction quality using the area under the precision-recall curve (AUPRC) (Davis and Goadrich, 2006). The predictive performance of hold-out interventions was quantified using the negative log-likelihood (NLL).

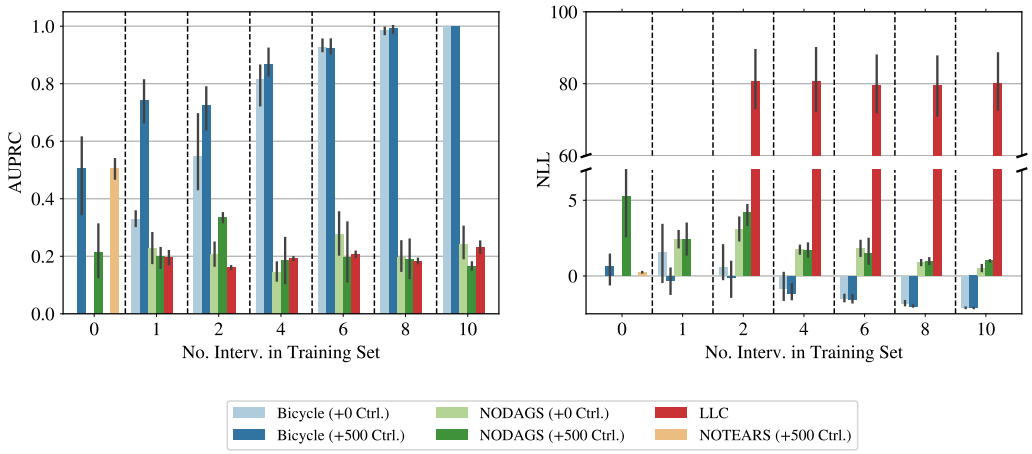

Figure 3: Benchmarking of alternative methods, considering synthetic data sampled from a 10-node Erdős–Rényi graph with cycles for a varying number of intervened variables (up to 10) seen during training created using SEM1 (Ornstein-Uhlenbeck process). Left: AUPRC of the recovered graph (higher is better); Right: NLL for unseen pairwise perturbations (lower is better). Bars denote average values across 5 repeated experiments with error bars corresponding to one standard deviation. The label "+500 Ctrl." indicates the model receives 500 additional control samples independent of the number of interventions.

The result for SEM1 (Fig. 3) clearly demonstrates the advantages of Bicycle on this data generation process (DGP). First, we note that Bicycle exploited the interventional data more effectively than SOTA methods. Furthermore, a Bicycle model trained with an additional 500 control samples on top of the interventional data ("Bicycle + 500 Ctrl.") is on par with or outperforms the model trained on interventional data only. For 0 interventional samples and 500 control samples (leftmost blue bar), the model approximately corresponds to the generative model of Dictys, albeit without the additional inductive bias of masking possible causal relationships based on additional data modalities. Such a model, exclusively trained on unperturbed control samples, is already outperformed by Bicycle with one additional intervention by +0.23 on AUPRC. As expected, the inclusion of additional control samples generally helps, especially in the regime of smaller numbers of interventions. The high standard deviations for both AUPRC and NLL in low-sample regimes indicate that Bicycle's inferences become more robust as data sizes increase. NOTEARS cannot model interventional data but, on observational data, performed comparably to Bicycle on average, with a smaller standard

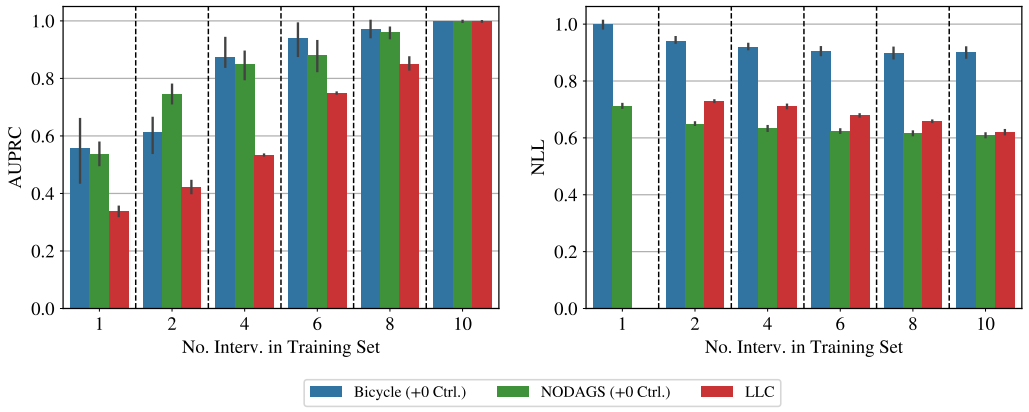

Figure 4: Benchmarking of alternative models (5 repeated experiments), considering synthetic data sampled from a 10-node Erdős–Rényi graph with cycles created using SEM2 (a time-discrete linear). Left: AUPRC of the recovered graph, right: NLL for unseen pairwise perturbations.

deviation. NODAGS-Flow and LLC show markedly lower performance compared to Bicycle under this DGP and do not improve with more interventional data.

In the SEM2 scenario (Fig. 4), NODAGS-Flow shows a rapid improvement of AUPRC with increasing numbers of interventional contexts, while LLC requires more interventional samples to reach similar performance. These behaviours are in line with the results reported in the primary manuscript of the NODAGS-Flow method. Although Bicycle was not designed for this DGP, it achieved competitive performance in AUPRC, even with NODAGS-Flow. It performed slightly worse in terms of the NLL, however, for biological applications, the AUPRC on the reconstructed GRN is often the metric of greater interest. For complete details on numerical performance indicators, we refer to Tab. 2 and Tab. 3.

4.2. Perturb-seq Data

Next, we set out to assess Bicycle on real-world data, leveraging three published single-cell perturbation screen datasets (Frangieh et al., 2021). Given that the true causal graph is unknown in this setting, we focused on evaluating generalization to unknown perturbations. We evaluate the interventional mean average error (I-MAE, more details below) based on a single-cell RNA sequencing perturbation screen (Perturb-seq) dataset, which contains targeted CRISPR knock-out perturbations of 249 target genes. The perturbations were performed in tumor-infiltrating lymphocytes (TILs) of melanoma patients in three different settings, which in the following were treated as separate datasets: a baseline culture of only TILs in a neutral medium ("Control"), a culture of TILs with the immune stimulatory cytokine interferon-γ added ("IFN-γ"), and a co-culture of TILs with patient-derived melanoma cells ("Co-Culture").

The evaluation metric we use, I-MAE on the normalized count data, was previously considered to benchmark causal discovery methods (Sethuraman et al., 2023). The I-MAE quantifies the reconstruction error of a model, in an unseen interventional condition, when predicting expression levels

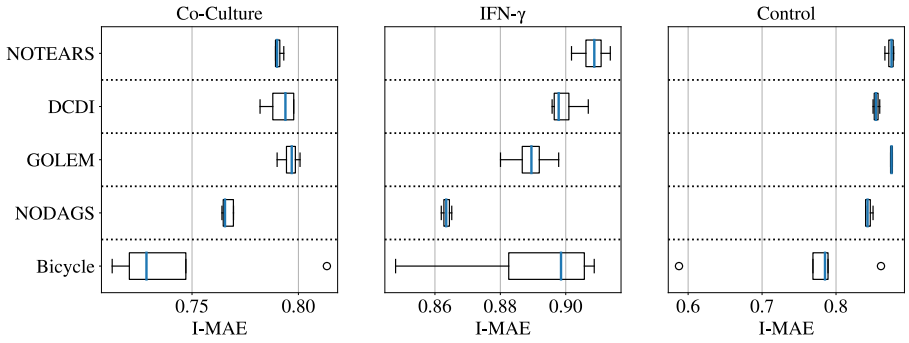

Figure 5: Comparison of out-of-distribution generalisation of learned regulatory relationships in terms of the interventional mean absolute error. I-MAE: Interventional mean absolute error, Co-culture: Lymphocytes co-cultured with melanoma cells, IFN-γ: Lymphocytes stimulated by interferon-γ, Control: Lymphocytes cultured in neutral medium. Box plots show distribution (minimum, lower quartile, median, upper quartile, maximum, and outliers) over I-MAE for 5 runs with the best hyperparameter set.

of genes that were not directly targeted, based on the expression values of the inferred causal parents. Thus, the I-MAE indirectly quantifies the accuracy of the reconstructed causal graph: A lower I-MAE implies that the inferred causal relationships hold better in novel contexts. It is similar to the interventional NLL (Lopez et al., 2022), but it allows to compare our method, which uses a discrete likelihood function, to SOTA methods using continuous likelihood functions.

We compared Bicycle to four SOTA methods for causal graph discovery: NOTEARS, DCDI, GOLEM, and NODAGS. We took the same subset of 61 genes and the exact training and test split as reported in the NODAGS-Flow paper (Sethuraman et al., 2023): We used 90% of the interventions for training and evaluated the I-MAE on the remaining 10% of held-out interventions, performing separate analyses for each of the three datasets.

To train the comparison partners, the mRNA count data was normalized for library size to remove technical noise and $\log(x+1)$-transformed to accommodate the uniform variance of the Gaussian likelihood functions of the comparison partners (see Appendix B for details). Due to its count-specific technical noise model, Bicycle was trained directly on the unnormalized count data. We performed five runs of training Bicycle using the following, fixed hyperparameters (following the nomenclature of Eq. 8): $\gamma = 1.0, \xi = 1.0$. To choose the optimal λ, we performed a search over $\lambda \in \{0.01, 0.1, 1.0\}$ on the training data.

While SOTA methods directly predict normalized counts, we used the following approach for Bicycle: For a given set of perturbed genes P_i, the I-MAE is calculated by evaluating the expectation value of the relative composition $p_{nd} = \frac{x_{nd}}{\sum_g x_{ng}}$ under the conditional probability $p(x_{nd} \mid x_{n\neg d})$ over observed count-values from our trained model for each out-of-distribution cell for all genes $d \in G \backslash P_i$, i.e., those which were not directly targeted by an intervention in the given context i. Under the I-MAE, models capturing non-causal relationships in their conditional distribution $p(x_{nd} \mid x_{n\neg d})$ are expected to generalize worse to novel interventional conditions, where these non-causal correlations may no longer be predictive. Then we normalized the relative composition

values to a total number of 10^4 counts per cell, added one "pseudo-count" per gene and took the element-wise logarithm. More details are given in Appendix B.

The results shown in Fig. 5 compare Bicycle to SOTA methods in terms of out-of-distribution generalization of the learned causal graph to novel perturbation conditions. Bicycle compares favourably to all of the SOTA methods on two of the three real-world datasets. On the IFN-γ dataset it is comparable to all methods except for NODAGS-Flow, which is the only nonlinear method leveraging perturbation data to discover cyclic causal structure in this comparison. Nevertheless, we note the scale of the x-axis in the plot, and that the difference between all methods was least pronounced on this dataset. Note that some of the runs produced outliers in the resulting I-MAE values. To achieve the lowest variance in results, we recommend the user to train an ensemble of Bicycle models (e.g., five models) and use the median, which is more robust to outliers.

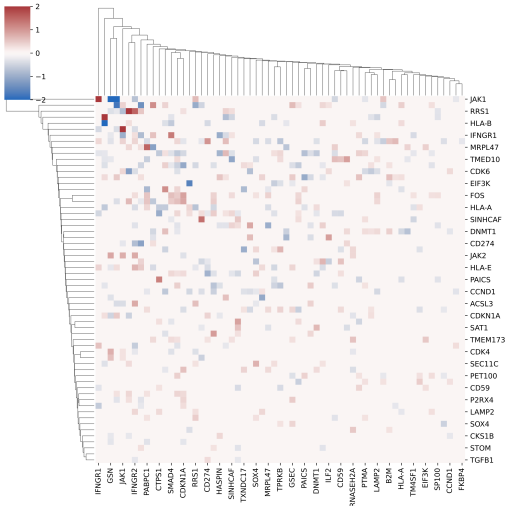

Figure 6: Average estimate for the GRN matrix β on the interferon-γ dataset.

To explore the biological relevance of an interpretable representation of the underlying inferred causal structure, Fig. 6 shows the average estimate of β over the runs with the best performing hyperparameters on the interferon-γ dataset. Note that the strongest effect is from one member of the interferon-γ receptor family (IFNGR1) to a member of the JAK/STAT signalling pathway (JAK1), which is in line with known biology (Horvath, 2004).

5. Discussion

Here, we have presented Bicycle, a new approach that can leverage data from both observational and interventional conditions to estimate cyclic causal graphs by integrating the idea of sparse interventional causal mechanisms into linear dynamical systems. Bicycle models the true expression values in a latent space, which is particularly interesting in domains such as single-cell biology, where relevant regulatory interactions have to be inferred from observations suffering from strong technical noise. The model can be equipped with various observation noise models, e.g., to model count data or other modalities. Synthetic experiments show that Bicycle performs on par with or better than SOTA methods, including on data generation processes for which it was not specifically designed. Experiments on real-world data show that Bicycle can successfully discover invariant causal mechanisms, which generalize well to novel contexts, improving over SOTA causal discovery methods in two out of three datasets studied. Furthermore, in an important advantage over methods that use priors on relationships between genes rather than inferring them, these experiments show that the interpretable estimate of the causal graph estimated by Bicycle has the potential to provide important novel insights into gene regulatory networks and thereby enhance our current understanding of biological systems. Bicycle lays the foundation for several future extensions and research directions:

Causal Modeling of Transient Dynamical Systems: Bicycle approximates a dynamical process by assuming that cells are in a steady state. However, in many datasets of interest, cells are undergoing a developmental process. Beginning from stem cells they progress through progenitor cell types to fully differentiated cell types over time. During this process, the causal GRN is constantly involved in governing cell fate decisions and maintenance. Extensions to Bicycle will allow the modeling of multiple cell types and explicitly account for the temporal causal dynamics of development and disease at the single-cell level.

Imperfect and Unknown Interventions: Currently, we assume perfect, known interventions. By introducing additional parameters, it will be possible to extend Bicycle to imperfect interventions, as well as unknown intervention strengths and targets, both critical issues when trying to extract insights from real-world single-cell perturbation screens.

Scalability and Large Datasets: Based on empirical observations from our study, it seems to be beneficial to solve the Lyapunov equation directly (as we did for the smaller, synthetic datasets) instead of using a loss component penalizing the difference between the left- and right-hand side of the equation (as we did for larger, real-world datasets). A promising future direction will be to exploit low-rank representations in Bicycle (e.g. (Kressner et al., 2019)) to more efficiently compute direct solutions even for a large number of genes. Additionally, it is also possible to implement low-rank representations of β to further improve the scalability and interpretability of large-scale causal interaction matrices in terms of latent factors (Lopez et al., 2022). Finally, instead of optimizing the sufficient statistics of latent variables for each gene and sample individually, it would be an interesting direction of future work to build on amortized variational inference, using an encoder model that directly predicts the latent variables given the input features and intervention context, thereby allowing scalability to datasets with millions of cells.

Acknowledgments

This publication was supported through state funds approved by the State Parliament of Baden-Württemberg for the Innovation Campus Health + Life Science alliance Heidelberg Mannheim as well as EMBL IT Services HPC resources. We would like to thank Alexander Aivazidis, Jana Braunger, Elizaveta Chernova, Jan Gleixner, Abdul Moeed and Marc-Jan Bonder for fruitful discussions.

References

Sachiho A. Adachi, Seiya Nishizawa, Ryuji Yoshida, Tsuyoshi Yamaura, Kazuto Ando, Hisashi Yashiro, Yoshiyuki Kajikawa, and Hirofumi Tomita. Contributions of changes in climatology and perturbation and the resulting nonlinearity to regional climate change. *Nature Communications*, 8(1):2224, December 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-02360-z.

Constantin Ahlmann-Eltze and Wolfgang Huber. Comparison of transformations for single-cell RNA-seq data. *Nature Methods*, 20(5):665–672, May 2023. ISSN 1548-7105. doi: 10.1038/s41592-023-01814-1.

Andrew V. Anzalone, Luke W. Koblan, and David R. Liu. Genome editing with CRISPR–Cas nucleases, base editors, transposases and prime editors. *Nature Biotechnology*, 38(7):824–844, July 2020. ISSN 1546-1696. doi: 10.1038/s41587-020-0561-9.

Christoph Bock, Paul Datlinger, Florence Chardon, Matthew A. Coelho, Matthew B. Dong, Keith A. Lawson, Tian Lu, Laetitia Maroc, Thomas M. Norman, Bicna Song, Geoff Stanley, Sidi Chen, Mathew Garnett, Wei Li, Jason Moffat, Lei S. Qi, Rebecca S. Shapiro, Jay Shendure, Jonathan S. Weissman, and Xiaowei Zhuang. High-content CRISPR screening. *Nature Reviews Methods Primers*, 2(1):1–23, February 2022. ISSN 2662-8449. doi: 10.1038/s43586-021-00093-4.

Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M. Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885–2915, October 2021. doi: 10.1214/21-AOS2064.

Carmen Bravo González-Blas, Seppe De Winter, Gert Hulselmans, Nikolai Hecker, Irina Matetovici, Valerie Christiaens, Suresh Poovathingal, Jasper Wouters, Sara Aibar, and Stein Aerts. SCENIC+: Single-cell multiomic inference of enhancers and gene regulatory networks. *Nature Methods*, 20(9):1355–1367, September 2023. ISSN 1548-7105. doi: 10.1038/s41592-023-01938-4.

Philippe Brouillard, Sébastien Lachapelle, Alexandre Lacoste, Simon Lacoste-Julien, and Alexandre Drouin. Differentiable Causal Discovery from Interventional Data. *Advances in Neural Information Processing Systems*, pages 21865–21877, 2020.

Paul Datlinger, André F. Rendeiro, Thorina Boenke, Martin Senekowitsch, Thomas Krausgruber, Daniele Barreca, and Christoph Bock. Ultra-high-throughput single-cell RNA sequencing and perturbation screening with combinatorial fluidic indexing. *Nature Methods*, 18(6):635–642, June 2021. ISSN 1548-7105. doi: 10.1038/s41592-021-01153-z.

Jesse Davis and Mark Goadrich. The relationship between Precision-Recall and ROC curves. In *Proceedings of the 23rd International Conference on Machine Learning - ICML '06*, pages 233–240, Pittsburgh, Pennsylvania, 2006. ACM Press. ISBN 978-1-59593-383-6. doi: 10.1145/1143844.1143874.

Philipp Dettling, Mathias Drton, and Mladen Kolar. On the lasso for graphical continuous lyapunov models. *arXiv e-prints*, August 2022. doi: 10.48550/arXiv.2208.13572.

Philipp Dettling, Roser Homs, Carlos Améndola, Mathias Drton, and Niels Richard Hansen. Identifiability in Continuous Lyapunov Models. *SIAM Journal on Matrix Analysis and Applications*, pages 1799–1821, December 2023. ISSN 0895-4798. doi: 10.1137/22M1520311.

S.-J. Dunn, G. Martello, B. Yordanov, S. Emmott, and A. G. Smith. Defining an essential transcription factor program for naïve pluripotency. *Science*, 344(6188):1156–1160, June 2014. doi: 10.1126/science.1248882.

Franklin M. Fisher. A Correspondence Principle for Simultaneous Equation Models. *Econometrica*, 38(1):73–92, 1970. ISSN 0012-9682. doi: 10.2307/1909242.

Chris J. Frangieh, Johannes C. Melms, Pratiksha I. Thakore, Kathryn R. Geiger-Schuller, Patricia Ho, Adrienne M. Luoma, Brian Cleary, Livnat Jerby-Arnon, Shruti Malu, Michael S. Cuoco, Maryann Zhao, Casey R. Ager, Meri Rogava, Lila Hovey, Asaf Rotem, Chantale Bernatchez, Kai W. Wucherpfennig, Bruce E. Johnson, Orit Rozenblatt-Rosen, Dirk Schadendorf, Aviv Regev, and Benjamin Izar. Multimodal pooled Perturb-CITE-seq screens in patient models define

mechanisms of cancer immune evasion. *Nature genetics*, 53(3):332–341, March 2021. ISSN 1546-1718 1061-4036. doi: 10.1038/s41588-021-00779-1.

Jacob W. Freimer, Oren Shaked, Sahin Naqvi, Nasa Sinnott-Armstrong, Arwa Kathiria, Christian M. Garrido, Amy F. Chen, Jessica T. Cortez, William J. Greenleaf, Jonathan K. Pritchard, and Alexander Marson. Systematic discovery and perturbation of regulatory genes in human T cells reveals the architecture of immune networks. *Nature Genetics*, 54(8):1133–1144, August 2022. ISSN 1546-1718. doi: 10.1038/s41588-022-01106-y.

Christoph Hafemeister and Rahul Satija. Normalization and variance stabilization of single-cell RNA-seq data using regularized negative binomial regression. *Genome Biology*, 20(1):296, December 2019. ISSN 1474-760X. doi: 10.1186/s13059-019-1874-1.

Christina Heinze-Deml, Jonas Peters, and Nicolai Meinshausen. Invariant Causal Prediction for Nonlinear Models. *Journal of Causal Inference*, 6(2), September 2018. ISSN 2193-3685. doi: 10.1515/jci-2017-0016.

Curt M. Horvath. The Jak-STAT pathway stimulated by interferon gamma. *Science's STKE : signal transduction knowledge environment*, 2004(260):tr8, November 2004. ISSN 1525-8882. doi: 10.1126/stke.2602004tr8.

Antti Hyttinen, Frederick Eberhardt, and Patrik O. Hoyer. Learning Linear Cyclic Causal Models with Latent Variables. *Journal of Machine Learning Research*, 13(109):3387–3439, 2012. ISSN 1533-7928.

Junyao Jiang, Pin Lyu, Jinlian Li, Sunan Huang, Jiawang Tao, Seth Blackshaw, Jiang Qian, and Jie Wang. IReNA: Integrated regulatory network analysis of single-cell transcriptomes and chromatin accessibility profiles. *iScience*, 25(11):105359, November 2022. ISSN 2589-0042. doi: 10.1016/j.isci.2022.105359.

Aryan Kamal, Christian Arnold, Annique Claringbould, Rim Moussa, Nila H Servaas, Maksim Kholmatov, Neha Daga, Daria Nogina, Sophia Mueller-Dott, Armando Reyes-Palomares, Giovanni Palla, Olga Sigalova, Daria Bunina, Caroline Pabst, and Judith B Zaugg. GRANIE and GRANPA : Inference and evaluation of enhancer-mediated gene regulatory networks. *Molecular Systems Biology*, 19(6):e11627, June 2023. ISSN 1744-4292, 1744-4292. doi: 10.15252/msb.202311627.

Nan Rosemary Ke, Silvia Chiappa, Jane X Wang, Jorg Bornschein, Anirudh Goyal, Melanie Rey, Theophane Weber, Matthew Botvinick, Michael Curtis Mozer, and Danilo Jimenez Rezende. Learning to induce causal structure. In *International Conference on Learning Representations*, 2023a.

Nan Rosemary Ke, Sara-Jane Dunn, Jorg Bornschein, Silvia Chiappa, Melanie Rey, Jean-Baptiste Lespiau, Albin Cassirer, Jane Wang, Theophane Weber, David Barrett, Matthew Botvinick, Anirudh Goyal, Mike Mozer, and Danilo Rezende. DiscoGen: Learning to discover gene regulatory networks. *arXiv e-prints*, (arXiv:2304.05823), April 2023b. doi: 10.48550/arXiv.2304.05823.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*. arXiv, 2015. doi: 10.48550/arXiv.1412.6980.

Daniel Kressner, Stefano Massei, and Leonardo Robol. Low-Rank Updates and a Divide-And-Conquer Method for Linear Matrix Equations. *SIAM Journal on Scientific Computing*, 41(2): A848–A876, January 2019. ISSN 1064-8275. doi: 10.1137/17M1161038.

Hao-Chih Lee, Matteo Danieletto, Riccardo Miotto, Sarah T. Cherng, and Joel T. Dudley. Scaling structural learning with NO-BEARS to infer causal transcriptome networks. In *Biocomputing 2020*, pages 391–402, Kohala Coast, Hawaii, USA, 2019. WORLD SCIENTIFIC. ISBN 9789811215629 9789811215636. doi: 10.1142/9789811215636_0035.

Phillip Lippe, Taco Cohen, and Efstratios Gavves. Efficient Neural Causal Discovery without Acyclicity Constraints. In *International Conference on Learning Representations*. arXiv, 2022. doi: 10.48550/arXiv.2107.10483.

Romain Lopez, Jan-Christian Huetter, Jonathan Pritchard, and Aviv Regev. Large-Scale Differentiable Causal Discovery of Factor Graphs. *Advances in Neural Information Processing Systems*, 35:19290–19303, December 2022.

Malte D Luecken and Fabian J Theis. Current best practices in single-cell RNA-seq analysis: A tutorial. *Molecular Systems Biology*, 15(6):e8746, June 2019. ISSN 1744-4292. doi: 10.15252/msb.20188746.

Tom Michoel and Jitao David Zhang. Causal inference in drug discovery and development. *Drug Discovery Today*, 28(10):103737, October 2023. ISSN 1359-6446. doi: 10.1016/j.drudis.2023.103737.

Joris M. Mooij, Dominik Janzing, and Bernhard Schölkopf. From Ordinary Differential Equations to Structural Causal Models: The deterministic case. *AUAI Press*, Proceedings of the Twenty-Ninth Conference Annual Conference on Uncertainty in Artificial Intelligence:440–448, 2013.

Matthew R. Nelson, Hannah Tipney, Jeffery L. Painter, Judong Shen, Paola Nicoletti, Yufeng Shen, Aris Floratos, Pak Chung Sham, Mulin Jun Li, Junwen Wang, Lon R. Cardon, John C. Whittaker, and Philippe Sanseau. The support of human genetic evidence for approved drug indications. *Nature Genetics*, 47(8):856–860, August 2015. ISSN 1546-1718. doi: 10.1038/ng.3314.

Hamed Nilforoshan, Michael Moor, Yusuf Roohani, Yining Chen, Anja Šurina, Michihiro Yasunaga, Sara Oblak, and Jure Leskovec. Zero-shot causal learning. *arXiv*, page 2301.12292, August 2023.

Thomas M. Norman, Max A. Horlbeck, Joseph M. Replogle, Alex Y. Ge, Albert Xu, Marco Jost, Luke A. Gilbert, and Jonathan S. Weissman. Exploring genetic interaction manifolds constructed from rich single-cell phenotypes. *Science*, 365(6455):786–793, August 2019. doi: 10.1126/science.aax4438.

Sharon L. Paige, Karolina Plonowska, Adele Xu, and Sean M. Wu. Molecular Regulation of Cardiomyocyte Differentiation. *Circulation research*, 116(2):341–353, January 2015. ISSN 0009-7330. doi: 10.1161/CIRCRESAHA.116.302752.

Alexander Gilbert Reisach, Christof Seiler, and Sebastian Weichwald. Beware of the Simulated DAG! Causal Discovery Benchmarks May Be Easy to Game. In *Advances in Neural Information Processing Systems*, November 2021.

Joseph M. Replogle, Reuben A. Saunders, Angela N. Pogson, Jeffrey A. Hussmann, Alexander Lenail, Alina Guna, Lauren Mascibroda, Eric J. Wagner, Karen Adelman, Gila Lithwick-Yanai, Nika Iremadze, Florian Oberstrass, Doron Lipson, Jessica L. Bonnar, Marco Jost, Thomas M. Norman, and Jonathan S. Weissman. Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. *Cell*, 185(14):2559–2575.e28, July 2022. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2022.05.013.

Christina D Romer and David H Romer. The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks. *American Economic Review*, 100(3):763–801, June 2010. ISSN 0002-8282. doi: 10.1257/aer.100.3.763.

Yusuf Roohani, Kexin Huang, and Jure Leskovec. Predicting transcriptional outcomes of novel multigene perturbations with GEARS. *Nature Biotechnology*, pages 1–9, August 2023. ISSN 1546-1696. doi: 10.1038/s41587-023-01905-6.

Bernhard Scholkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward Causal Representation Learning. *Proceedings of the IEEE*, 109(5):612–634, May 2021. ISSN 0018-9219, 1558-2256. doi: 10.1109/JPROC.2021.3058954.

Daniel Schraivogel, Lars M. Steinmetz, and Leopold Parts. Pooled Genome-Scale CRISPR Screens in Single Cells. *Annual Review of Genetics*, 57(1):223–244, 2023. doi: 10.1146/annurev-genet-072920-013842.

Muralikrishnna G. Sethuraman, Romain Lopez, Rahul Mohan, Faramarz Fekri, Tommaso Biancalani, and Jan-Christian Hütter. NODAGS-Flow: Nonlinear Cyclic Causal Structure Learning. *arXiv e-prints*, (arXiv:2301.01849), January 2023. doi: 10.48550/arXiv.2301.01849.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*, volume 81 of *Lecture Notes in Statistics*. Springer, New York, NY, 1993. ISBN 978-1-4612-7650-0 978-1-4612-2748-9. doi: 10.1007/978-1-4612-2748-9.

Alejandro Tejada-Lapuerta, Paul Bertin, Stefan Bauer, Hananeh Aliee, and Fabian J Theis. Causal machine learning for single-cell genomics. *arXiv.org*, 2023. doi: 10.48550/arXiv.2310.14935.

Gherardo Varando and Niels Hansen. Graphical continuous lyapunov models. In Jonas Peters and David Sontag, editors, *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 124 of *Proceedings of Machine Learning Research*, pages 989–998. PMLR, August 2020.

Lingfei Wang, Nikolaos Trasanidis, Ting Wu, Guanlan Dong, Michael Hu, Daniel E. Bauer, and Luca Pinello. Dictys: Dynamic gene regulatory network dissects developmental continuum with single-cell multiomics. *Nature Methods*, 20(9):1368–1378, September 2023. ISSN 1548-7105. doi: 10.1038/s41592-023-01971-3.

Ming Chen Wang and G. E. Uhlenbeck. On the Theory of the Brownian Motion II. *Reviews of Modern Physics*, 17(2-3):323–342, April 1945. ISSN 0034-6861. doi: 10.1103/RevModPhys.17.323.

Albert Xue, Jingyou Rao, Sriram Sankararaman, and Harold Pimentel. Dotears: Scalable, consistent DAG estimation using observational and interventional data. *arXiv*, page 2305.19215, May 2023.

Hengshi Yu and Joshua D. Welch. PerturbNet predicts single-cell responses to unseen chemical and genetic perturbations. *bioRxiv : the preprint server for biology*, 2022. doi: 10.1101/2022.07.20.500854.

Xun Zheng, Bryon Aragam, Pradeep Ravikumar, and Eric P. Xing. DAGs with NO TEARS: Continuous Optimization for Structure Learning. In *32nd Conference on Neural Information Processing Systems*, Montréal, Canada, November 2018. doi: 10.48550/arXiv.1803.01422.

Appendix A. Synthetic Experiments

A.1. Data Generation

A.1.1. GRAPH GENERATION

The generated graphs were constructed using an Erdős–Rényi $G(n, p)$ model, with an expected edge density of 2 and $n = 10$ nodes. Edge weights were sampled uniformly from $[-0.95, -0.25] \cup [0.25, 0.95]$.

A.1.2. STRUCTURAL EQUATION MODELS

We used two SEMs to generate data, which we elaborate in more detail here.

Ornstein–Uhlenbeck process (SEM1) The results shown in Fig. 3 use an Ornstein–Uhlenbeck (OU) process to generate synthetic data.

The OU process is a stochastic model developed to describe the velocity of a massive Brownian particle subject to friction. This process is a stationary Gauss–Markov process, meeting the criteria of being Gaussian (characterized by a normal distribution), Markovian (future states depend only on the present state, not on the sequence of events that preceded it), and temporally homogeneous (its statistical properties do not change over time). It can be viewed as a modification of the continuous-time random walk or Wiener process. Unlike a standard random walk, it tends to move back towards a central location, reflecting the mean-reverting nature of the process. In one dimension, we can write the OU process as

$$\mathrm{d}z(t) = \theta \cdot (\mu - z(t))\mathrm{d}t + \sigma \mathrm{d}W(t) \qquad \text{with} \quad z_0 = a, \tag{7}$$

with $\mu, a \in \mathbb{R}$ and $\sigma, \theta > 0$. We make several remarks about this definition. (i) σ describes the influence of stochasticity on the process. Without it (i.e., supposing $\sigma = 0$), z converges exponentially quickly towards μ. (ii) μ is the mean reversion level, i.e., if $z(t)$ is above this level, the drift will drag it down, and vice versa. (iii) θ is the strength of this drift effect. The one-dimensional process is a special case of the multi-dimensional process of Eq. (1).

This process is chosen as it is a common model of the transcriptional kinetics of gene expression, i.e., the production of mRNA molecules. Hence, dependencies between nodes in the causal graph (i.e., genes) are generated according to Eq. (4). We set the following values for our parameters:

$$\begin{aligned} \alpha_g &= 1.0, \quad \hat{\alpha}_g^i = 0.1 \\ \hat{\beta}_{kg}^i &= 0 \\ \sigma_{gk} &= 0.1, \quad \hat{\sigma}_{gk}^i = 0.001 \end{aligned}$$

We compute ω according to Eq. (2) for different interventions and sample our observed data from the steady-state distribution $p_{SS}(x) = \mathcal{N}(\bar{x}, \omega)$ with $\bar{x} = B^{-1}\alpha$.

General Linear SEM (SEM2) In addition to the data generation process presented above, we ran experiments on a linear SEM as used in (Sethuraman et al., 2023; Hyttinen et al., 2012). For a perturbation experiment k, we have:

$$\begin{aligned} x(t) &= f_k(x(t-1)) + \epsilon \\ &= U_k \beta x(t-1) + U_k \epsilon + c, \end{aligned}$$

Then, we generated data specific for each sampling of noise from the corresponding equilibrium state,

$$x = (\mathbb{I}_G - U_k\beta)^{-1}(U_k\epsilon + c).$$

As described in the main text, we set $c \sim \eta \cdot \mathcal{N}(0, \mathbb{I}_G)$, representing the perturbation effect (with different effect sizes η) and $\epsilon \sim 0.5 \cdot \mathcal{N}(0, \mathbb{I}_G)$. To guarantee asymptotic stability, we ensure that all eigenvalues λ_i of $U_k\beta$ satisfy $|\lambda_i| < 1$ (Fisher, 1970) across all interventional conditions k. For further details and explanations, we refer the reader to (Hyttinen et al., 2012).

A.1.3. DATASET CONSTRUCTION AND HYPERPARAMETER SELECTION

To allow a fair comparison across all models, we used the same dataset for all models and repeated each experiment at least three times with different seeds and different graphs. In case the model required tuning hyperparameters, we used an 80%/20% train-validation split of the samples and selected the model that performed best on the validation set. For LLC, where no hyperparameter tuning was necessary, we used all data for training. The reported values are predictions on a hold-out test set that contains pairwise perturbations, containing perturbations where 0, 1 or 2 interventions have been seen before – but only individually. Tab. 6 shows an overview of hyperparameters tuned for each model.

The choice of 500 observational samples and 250 samples for each interventional context was chosen to mirror the statistics of common single-cell perturbation datasets. These usually contain on the order of 100-250 cells per perturbation context and a larger number of unperturbed (observational) cells.

We use the implementations of LLC, NOTEARS and NODAGS-Flow as given in the NODAGS-Flow paper (Sethuraman et al., 2023). The remaining parameters of the NODAGS-Flow model that remained untuned were set according to the synthetic experiments in that paper – except for the number of epochs, which we increased to 500 to ensure convergence.

Table 1: Tuned Hyperparamters

Model	**Hyperparameters**	**Values/Range**
LLC	-	-
NOTEARS	λ (Sparsity Penalty)	[1e-5, 1e-4, 1e-3, 1e-2, 1e-1, 1, 10, 100]
NODAGS-Flow	λ_c	[1e-4, 1e-3, 1e-2, 1e-1, 1]
	LR	[1e-4, 1e-3, 1e-2, 1e-1, 1]
	n-hidden	[0, 1, 2, 3]
	Func. Relation	[gst-mlp, lin-mlp]

A.2. Computational complexity

Considering our objective function shown in Eq. (5), we can parallelize the training across interventions and can iterate over batches of samples to update the latent variables. However, our method is bounded by the complexity of inverting $B = \mathbb{I}_G - \beta^T$ to compute $\bar{z} = B^{-1}\alpha$, which is of complexity $\mathcal{O}(n^3)$ for n genes. The current memory footprint scales with $\mathcal{O}(\max(m, n) \cdot n)$, given m samples.

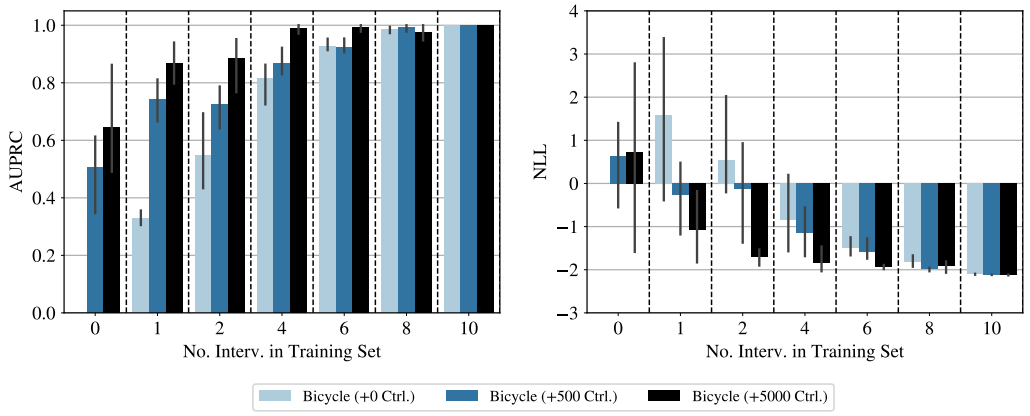

Figure 7: Model comparison for different Bicycle models on SEM1 when the number of additional control samples is varied. The number of perturbations per intervention is set to 200 per perturbation.

Empirically, we can scale the current method to approximately 1,000 features/genes, several hundred interventions and 100,000 samples/cells across multiple GPUs. Future work (cf. Sec. 5) will evaluate ways to improve the scalability of Bicycle.

Empirical results show that we can use the analytical solution to the Lyapunov Equation (see Eq. (2)) for datasets of up to 50 features given a 24GB GPU. Note that this is independent of the number of samples and only depends on the number of features. If the number of features per sample is higher, we use the objective as shown in Eq. (5) to approximate the covariance matrix ω.

A.3. Details and Additional Results

Tab. 2 reports detailed results corresponding to Fig. 3.

A.4. Additional Experiments

A.4.1. Influence of Observational Samples on Bicycle

We extended the experiment in Fig. 3 by adding one more set of models using 5000 observational samples. The results presented in Fig. 7 and corresponding Tab. 4 show that the model benefits greatly from additional observational samples in the regimes where few interventional samples are available, but these become less relevant when more perturbation data is provided. Especially in low-sample regimes, the standard deviations are high but shrink with more evidence about the process. Hence, especially in single-cell studies, where large amounts of unperturbed data are available, these should be included during training. However, this experiment confirms that a small amount of interventional samples can already improve causal discovery significantly even when a large amount of control samples are given.

A.4.2. Impact of Precise Perturbation Effects

To deepen our understanding of the performance of Bicycle and other methods on SEM2, we varied the perturbation strength η in Eq. (6). By decreasing the value of η towards zero, we mimicked

Table 2: Results for synthetic data generated from SEM1.

Int. Contexts	Model	AUPRC	NLL
0	Bicycle (+500 Ctrl.)	0.51 ± 0.14	0.65 ± 1.05
0	NODAGS (+500 Ctrl.)	0.21 ± 0.12	5.24 ± 3.6
0	NOTEARS (+500 Ctrl.)	0.51 ± 0.13	0.24 ± 0.06
1	Bicycle (+0 Ctrl.)	0.33 ± 0.03	1.58 ± 1.88
1	Bicycle (+500 Ctrl.)	0.74 ± 0.07	-0.28 ± 0.84
1	LLC	0.2 ± 0.07	
1	NODAGS (+0 Ctrl.)	0.23 ± 0.11	2.44 ± 1.09
1	NODAGS (+500 Ctrl.)	0.2 ± 0.05	2.44 ± 1.45
2	Bicycle (+0 Ctrl.)	0.55 ± 0.13	0.54 ± 1.28
2	Bicycle (+500 Ctrl.)	0.72 ± 0.07	-0.13 ± 1.16
2	LLC	0.16 ± 0.01	80.81 ± 24.52
2	NODAGS (+0 Ctrl.)	0.21 ± 0.09	3.08 ± 1.82
2	NODAGS (+500 Ctrl.)	0.33 ± 0.02	4.2 ± 0.82
4	Bicycle (+0 Ctrl.)	0.82 ± 0.08	-0.86 ± 0.93
4	Bicycle (+500 Ctrl.)	0.87 ± 0.05	-1.15 ± 0.57
4	LLC	0.19 ± 0.01	80.63 ± 25.91
4	NODAGS (+0 Ctrl.)	0.14 ± 0.06	1.72 ± 0.51
4	NODAGS (+500 Ctrl.)	0.19 ± 0.11	1.71 ± 0.6
6	Bicycle (+0 Ctrl.)	0.93 ± 0.02	-1.5 ± 0.22
6	Bicycle (+500 Ctrl.)	0.92 ± 0.03	-1.58 ± 0.27
6	LLC	0.21 ± 0.03	79.57 ± 24.01
6	NODAGS (+0 Ctrl.)	0.28 ± 0.14	1.82 ± 0.95
6	NODAGS (+500 Ctrl.)	0.2 ± 0.14	1.51 ± 1.1
8	Bicycle (+0 Ctrl.)	0.99 ± 0.01	-1.82 ± 0.14
8	Bicycle (+500 Ctrl.)	0.99 ± 0.01	-1.99 ± 0.04
8	LLC	0.18 ± 0.02	79.56 ± 24.66
8	NODAGS (+0 Ctrl.)	0.2 ± 0.11	0.88 ± 0.34
8	NODAGS (+500 Ctrl.)	0.19 ± 0.09	0.99 ± 0.28
10	Bicycle (+0 Ctrl.)	1.0 ± 0.0	-2.11 ± 0.01
10	Bicycle (+500 Ctrl.)	1.0 ± 0.0	-2.12 ± 0.0
10	LLC	0.23 ± 0.06	80.22 ± 24.58
10	NODAGS (+0 Ctrl.)	0.24 ± 0.12	0.51 ± 0.44
10	NODAGS (+500 Ctrl.)	0.17 ± 0.01	1.0 ± 0.04

a perfect CRISPRi experiment, where genes were *knocked down*, or inactivated, i.e., do($x_g = 0$). Fig. 8 show results for $\eta = 0.5$, and Fig. 9 show results for $\eta = 0.1$. We observed empirically that LLC suffered more the smaller the variation of the effects, while Bicycle and NODAGS lost

Table 3: Results for Synthetic Data generated from SEM2.

Int. Contexts	Model	AUPRC	NLL
1	Bicycle (+0 Ctrl.)	0.56 ± 0.15	1.0 ± 0.01
1	LLC	0.34 ± 0.07	
1	NODAGS (+0 Ctrl.)	0.54 ± 0.13	0.71 ± 0.02
2	Bicycle (+0 Ctrl.)	0.61 ± 0.06	0.94 ± 0.01
2	LLC	0.42 ± 0.08	0.73 ± 0.01
2	NODAGS (+0 Ctrl.)	0.75 ± 0.11	0.65 ± 0.01
4	Bicycle (+0 Ctrl.)	0.88 ± 0.06	0.92 ± 0.01
4	LLC	0.53 ± 0.01	0.71 ± 0.02
4	NODAGS (+0 Ctrl.)	0.85 ± 0.15	0.63 ± 0.02
6	Bicycle (+0 Ctrl.)	0.94 ± 0.06	0.91 ± 0.01
6	LLC	0.75 ± 0.01	0.68 ± 0.01
6	NODAGS (+0 Ctrl.)	0.88 ± 0.15	0.62 ± 0.02
8	Bicycle (+0 Ctrl.)	0.97 ± 0.03	0.9 ± 0.02
8	LLC	0.85 ± 0.08	0.66 ± 0.01
8	NODAGS (+0 Ctrl.)	0.96 ± 0.06	0.62 ± 0.02
10	Bicycle (+0 Ctrl.)	1.0 ± 0.0	0.9 ± 0.02
10	LLC	1.0 ± 0.0	0.62 ± 0.02
10	NODAGS (+0 Ctrl.)	1.0 ± 0.01	0.61 ± 0.02

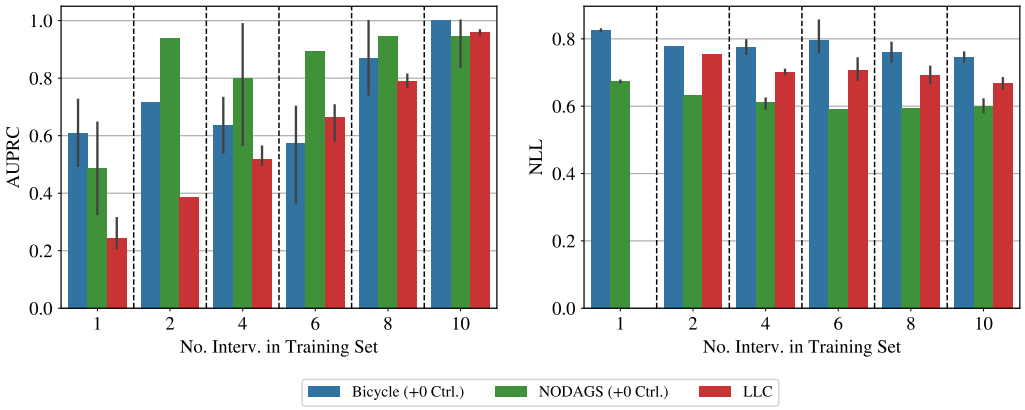

Figure 8: AUPRC and NLL on SEM2 for $\eta = 0.5$ in Eq. (6).

reconstruction and predictive power compared to more noisy knock-down, but still maintained good performance.

Table 4: Results for Synthetic Data generated from SEM1

Int. Contexts	Model	AUPRC	NLL
0	Bicycle (+500 Ctrl.)	0.51 ± 0.14	0.65 ± 1.05
0	Bicycle (+5000 Ctrl.)	0.65 ± 0.19	0.74 ± 2.2
1	Bicycle (+0 Ctrl.)	0.33 ± 0.03	1.58 ± 1.88
1	Bicycle (+500 Ctrl.)	0.74 ± 0.07	-0.28 ± 0.84
1	Bicycle (+5000 Ctrl.)	0.87 ± 0.07	-1.07 ± 0.83
2	Bicycle (+0 Ctrl.)	0.55 ± 0.13	0.54 ± 1.28
2	Bicycle (+500 Ctrl.)	0.72 ± 0.07	-0.13 ± 1.16
2	Bicycle (+5000 Ctrl.)	0.88 ± 0.1	-1.71 ± 0.19
4	Bicycle (+0 Ctrl.)	0.82 ± 0.08	-0.86 ± 0.93
4	Bicycle (+500 Ctrl.)	0.87 ± 0.05	-1.15 ± 0.57
4	Bicycle (+5000 Ctrl.)	0.99 ± 0.02	-1.84 ± 0.32
6	Bicycle (+0 Ctrl.)	0.93 ± 0.02	-1.5 ± 0.22
6	Bicycle (+500 Ctrl.)	0.92 ± 0.03	-1.58 ± 0.27
6	Bicycle (+5000 Ctrl.)	0.99 ± 0.01	-1.94 ± 0.05
8	Bicycle (+0 Ctrl.)	0.99 ± 0.01	-1.82 ± 0.14
8	Bicycle (+500 Ctrl.)	0.99 ± 0.01	-1.99 ± 0.04
8	Bicycle (+5000 Ctrl.)	0.98 ± 0.03	-1.92 ± 0.14
10	Bicycle (+0 Ctrl.)	1.0 ± 0.0	-2.11 ± 0.01
10	Bicycle (+500 Ctrl.)	1.0 ± 0.0	-2.12 ± 0.0
10	Bicycle (+5000 Ctrl.)	1.0 ± 0.0	-2.12 ± 0.01

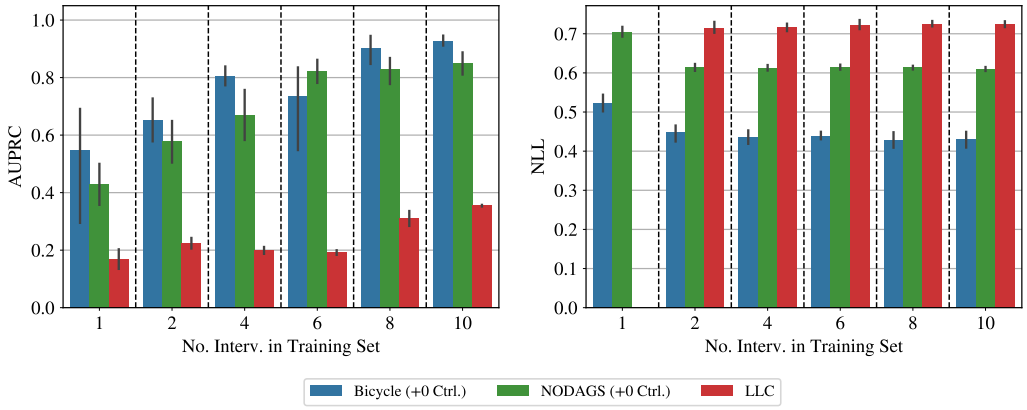

Figure 9: AUPRC and NLL on SEM2 for $\eta = 0.1$ in Eq. (6).

A.4.3. INFLUENCE OF PERTURBATION STRENGTH

The corresponding mean and standard deviations can be found in Tabs. 5 and 6.

Table 5: Results for Synthetic Data generated from SEM2 using $\eta = 0.5$

Int. Contexts	Model	AUPRC	NLL
1	Bicycle (+0 Ctrl.)	0.61 ± 0.16	0.83 ± 0.0
1	LLC	0.24 ± 0.06	
1	NODAGS (+0 Ctrl.)	0.49 ± 0.22	0.67 ± 0.0
2	Bicycle (+0 Ctrl.)	0.72	0.78
2	LLC	0.39 ± 0.0	0.75 ± 0.0
2	NODAGS (+0 Ctrl.)	0.94	0.63
4	Bicycle (+0 Ctrl.)	0.64 ± 0.13	0.78 ± 0.03
4	LLC	0.52 ± 0.04	0.7 ± 0.01
4	NODAGS (+0 Ctrl.)	0.8 ± 0.21	0.61 ± 0.02
6	Bicycle (+0 Ctrl.)	0.57 ± 0.18	0.8 ± 0.05
6	LLC	0.66 ± 0.07	0.71 ± 0.03
6	NODAGS (+0 Ctrl.)	0.89	0.59
8	Bicycle (+0 Ctrl.)	0.87 ± 0.18	0.76 ± 0.04
8	LLC	0.79 ± 0.03	0.69 ± 0.03
8	NODAGS (+0 Ctrl.)	0.95	0.59
10	Bicycle (+0 Ctrl.)	1.0 ± 0.0	0.75 ± 0.02
10	LLC	0.96 ± 0.01	0.67 ± 0.02
10	NODAGS (+0 Ctrl.)	0.95 ± 0.11	0.6 ± 0.02

Table 6: Results for Synthetic Data generated from SEM2 using $\eta = 0.1$

Int. Contexts	Model	AUPRC	NLL
1	Bicycle (+0 Ctrl.)	0.55 ± 0.22	0.52 ± 0.02
1	LLC	0.17 ± 0.08	
1	NODAGS (+0 Ctrl.)	0.43 ± 0.17	0.7 ± 0.03
2	Bicycle (+0 Ctrl.)	0.65 ± 0.07	0.45 ± 0.02
2	LLC	0.22 ± 0.04	0.71 ± 0.03
2	NODAGS (+0 Ctrl.)	0.58 ± 0.13	0.61 ± 0.02
4	Bicycle (+0 Ctrl.)	0.81 ± 0.03	0.44 ± 0.02
4	LLC	0.2 ± 0.04	0.72 ± 0.03
4	NODAGS (+0 Ctrl.)	0.67 ± 0.22	0.61 ± 0.02
6	Bicycle (+0 Ctrl.)	0.74 ± 0.16	0.44 ± 0.01
6	LLC	0.19 ± 0.02	0.72 ± 0.03
6	NODAGS (+0 Ctrl.)	0.82 ± 0.09	0.62 ± 0.01
8	Bicycle (+0 Ctrl.)	0.9 ± 0.05	0.43 ± 0.02
8	LLC	0.31 ± 0.09	0.73 ± 0.02
8	NODAGS (+0 Ctrl.)	0.83 ± 0.14	0.61 ± 0.01
10	Bicycle (+0 Ctrl.)	0.93 ± 0.02	0.43 ± 0.02
10	LLC	0.35 ± 0.01	0.72 ± 0.03
10	NODAGS (+0 Ctrl.)	0.85 ± 0.13	0.61 ± 0.02

Appendix B. Biological Experiments with Gene Knock-Outs

B.1. Preprocessing for single-cell RNA-seq data

Here we employ data processing strategies commonly used in the scRNA-seq community. In scRNA-seq experiments different cells might have different library sizes (number of unique molecular identifiers (UMI)), depending on the fraction of molecules captured in an experiment or the size differences between cells (Luecken and Theis, 2019; Hafemeister and Satija, 2019). To normalize library sizes between single cells, the number of raw counts for each gene is commonly divided by the total number of counts in a cell $\tilde{c}_{ng} = \frac{c_{ng}}{\sum_g c_{ng}}$, where c_{ng} corresponds to raw counts for gene g in cell n. The fractions are then scaled with a size factor of 10^4, which corresponds to the order of magnitude of average total counts found per cell in a scRNA-seq experiment $\tilde{c}'_{ng} = 10^4 \cdot \tilde{c}_{ng}$.

In addition to library size normalization we also account for the heteroskedasticity of the data. The variance of the expression of a given gene scales with the mean expression of that gene (i.e. the gene expression follows a Poisson/Gamma-Poisson distribution). Since most methods assume normal distribution and, consequently, uniform variance, several variance-stabilizing transformations have been proposed in the field. The one most commonly used is the logarithm with a pseudo-count of one: $\log(x+1)$ (Luecken and Theis, 2019; Ahlmann-Eltze and Huber, 2023). According to a recent benchmark it also shows the most reliable performance (Ahlmann-Eltze and Huber, 2023). For a more comprehensive overview of different normalization strategies we refer the interested reader to Luecken and Theis (2019); Ahlmann-Eltze and Huber (2023); Hafemeister and Satija (2019).

B.2. CRISPR interventions

CRISPR (Clustered Regularly Interspaced Short Palindromic Repeats) in combination with Cas9 (CRISPR-associated protein 9) is a commonly used tool for genome editing. In CRISPR-Cas9 *knock-out* experiments, the target gene's DNA sequence is cut and altered by the Cas9 nucleases in such a way that the protein encoded by the target gene becomes non-functional (Bock et al., 2022; Anzalone et al., 2020). The three data sets used in the experiments above employ a pooled CRISPR-Cas9 experiment called Perturb-seq (Frangieh et al., 2021).

Interventions in causal modeling are usually categorized into perfect interventions, which remove the dependencies on their causal parents, and imperfect interventions, which only alter the causal dependencies on their parents (Brouillard et al., 2020). Here, we assume CRISPR-Cas9 knock-outs to be perfect interventions. These methods usually target the translated region of genes. Consequently, a knock-out might not be directly reflected in the mRNA expression levels, even though it affects the functionality of the translated protein. Assuming perfect interventions is justified by the fact that the protein will be non-functional, independent of upstream regulators of the gene's mRNA levels.

However, it should be mentioned that CRISPR-Cas9 experiments are limited by off-target effects, the uncertainty of success of the perturbation and the stress the experiments induce on cells. (Schraivogel et al., 2023; Tejada-Lapuerta et al., 2023). Additionally, compensatory pathways can be activated, which might mitigate the perturbation effect (Bock et al., 2022). In such cases, assuming perfect interventions might not be justified anymore.

Conversely, CRISPR intervention (CRISPRi), as used in (Replogle et al., 2022), is a tool to *knock-down* (reduce the expression of) target genes using deactivated Cas9 nucleases (dCas9) (Bock et al.,

2022; Anzalone et al., 2020). dCas9 proteins can still recognise their target DNA sequences, but cannot cut the DNA sequence. Instead, they are combined with other proteins, so-called transcriptional repressors. CRISPRi experiments do not target the translated region of a gene, but rather its regulatory regions, thereby repressing mRNA expression of the target gene without altering the DNA sequence (Bock et al., 2022).

The effect of genetic perturbations in pooled CRISPR screens can be read out using single-cell RNA sequencing (scRNA-seq), which captures a noisy image each cell's transcriptomic state by counting the number of mRNA transcripts for each gene in each cell. For a recent, comprehensive overview of single-cell perturbation screen approaches, interested readers are refererred to Schraivogel et al. (2023)

B.3. Evaluation of I-MAE

We evaluate the I-MAE on hold-out cells of an unseen interventional condition P_i in the following way: We assume the interventions completely cut the incoming causal edges of the perturbed gene and therefore set $\hat{\beta}^i_{gt} = 0$ for each direct target gene $t \in P_i$ and all genes $g \in G$. As the aim of our model is to make good predictions for the effects of *given* interventions on target genes $t \in P_i$ on all other genes $d \in G \backslash P^i$, we optimize $\hat{\alpha}^i_t$ and $\hat{\sigma}^i_{tt}$ for each direct target gene $t \in P_i$ to fit the corresponding marginal distribution.

To see how well the learned causal relationships hold in the condition P_i, we then evaluate the I-MAE on the genes $d \in G \backslash P^i$, i.e., those which were not directly perturbed. To this end, we evaluate $p(x_{nd}|x_{n\neg d})$, i.e., the distribution over the measured expression-count value of each gene and cell, conditioned on the measured count-values of all the other genes for that specific cell. If our model succeeds in identifying invariant causal mechanisms, this distribution should also work in unseen contexts. In contrast, the breaking of non-causal correlations in the training conditions due to novel interventions in the test conditions would severely impair predictions based on such correlations.

The calculation of the conditional probability $p(x_{nd}|x_{n\neg d})$ in a latent variable model requires a marginalization over the latent variables z_{ng}:

$$\begin{aligned} p(x_{nd}|x_{n\neg d}) &= \frac{p(x_{nd}, x_{n\neg d})}{p(x_{n\neg d})} \\ &= \frac{\int p(x_{nd}, x_{n\neg d}|z_{n\cdot})p(z_{n\cdot})\mathrm{d}z}{p(x_{n\neg d})} \\ &= \frac{\int p(x_{nd}|x_{n\neg d}, z_{n\cdot})p(x_{n\neg d}|z_{n\cdot})p(z_{n\cdot})\mathrm{d}z}{p(x_{n\neg d})} \\ &= \int p(x_{nd}|x_{n\neg d}, z_{n\cdot})p(z_{n\cdot}|x_{n\neg d})\mathrm{d}z \end{aligned} \tag{8}$$

As the marginalization over the high-dimensional latent variables z_{ng} is infeasible, we approximate the exact conditional probability by a variational approximation, where we replace the exact posterior $p(z_{n\cdot}|x_{n\neg d})$ with a variational distribution $q(z_{n\cdot}|x_{n\neg d}) = \prod_{g\in G}\mathcal{N}(z_{ng}; \mu_{ng}, \sigma_{ng})$ under a mean-field assumption. We optimize the mean and standard deviation parameters μ_{ng}, σ_{ng} using the following ELBO:

$$\mathrm{ELBO}(\mu_n, \sigma_n) = \langle \log p(x_{n\neg g}|z_{n\neg g})\rangle_{q(z_{n\cdot}|\mu_{z_n}, \sigma_{z_n})} - \mathrm{KL}(q(z_{n\cdot}|\mu_{z_n\cdot}, \sigma_{z_n\cdot}) \,||\, \mathcal{N}(z_{n\cdot}; \bar{x}^i, \omega^i))$$

where $\mathcal{N}(z_{n\cdot}; \bar{x}^i, \omega^i)$ is the steady-state distribution derived from the dynamical parameters α^i, β^i, and σ^i for intervention condition P^i, and $p(x_{n\neg g}|z_{n\neg g}) = \text{Multinomial}(x_{n\neg g}|p_{x_{n\neg g}})$ for $p_{n\neg g} = \text{Softmax}(z_{n\neg g})$.

Now, we can calculate the expected relative composition $p_{nd} = \frac{x_{nd}}{\sum_g x_{nd}}$ under the approximate conditional probability distribution

$$\hat{p}(x_{nd}|x_{n\neg d}) = p(x_{nd}|\mu_{n\cdot})$$

over observed count-values from the trained model for each out-of-distribution cell for all genes $d \in G \backslash P_i$.

Data is then normalized and $\log(x+1)$-transformed, as mentioned in Sec. 4.2. In detail, we scale count values to a total sum of 10^4 counts per cell, add one "pseudo-count" per gene and take the element-wise logarithm. This procedure is standard in the single-cell field to make count-data amenable to fit by models with Gaussian likelihood functions (Ahlmann-Eltze and Huber, 2023), see also (Sethuraman et al., 2023). This allows us to compare our prediction of the normalized expression value for each cell and gene, based on our model and the expression values of all other genes, to the corresponding predictions of SOTA methods, as shown in Fig. 5.

Appendix C. Identifiability results for Bicycle

The upcoming section provides an overview and establishes the conditions necessary for Bicycle to converge to the true causal model, i.e. recovering the true underlying parameters $\alpha, \hat{\alpha}, \beta, \sigma$, and $\hat{\sigma}$.

C.1. Assumptions

In each of the following theorems, we consider samples from different distributions: (i) In the first theorem, we consider N samples from the observational distribution. (ii) In the second theorem, we consider samples from the G interventional conditions $P^{\neg i} = G \setminus \{i\}\ \forall i \in G$, i.e., from all conditions in which all but one variable was directly perturbed. (iii) In the last theorem, we consider samples from the G interventional conditions $P^i = \{i\}, i \in G$, i.e., from all conditions in which a single variable was directly perturbed.

Moreover, we assume that we can get a perfect estimate of $\bar{z}_i$ and the positive-definite covariance matrices ω_i of the multivariate normal distributions representing the steady-state distributions for all observed conditions.

Additionally, we again assume perfect interventions, i.e., $\hat{beta}^i = 0$, and we assume independent noise processes per gene, i.e., $\sigma_{gg}, \hat{\sigma}^i_{gg} > 0\ \forall g$ and $\sigma_{gh} = \hat{\sigma}^i_{gh} = 0\ \forall g \neq h$, which corresponds to assuming no latent confounders.

Theorem 1 *Purely observational datasets are not sufficient for recovering ground-truth parameters.*

In general, there is a many-to-one mapping from underlying dynamical parameters to the parameters (mean and covariance matrix) of the steady state distribution for $t \to \infty$. E.g., all changes to the dynamical parameters which rescale the time-axis or add a divergence-free flux to the final steady state leave the steady-state distribution invariant. The identifiability of the dynamical parameters from the observational steady-state distribution alone is discussed in detail in Dettling et al. (2023).

Theorem 2 *Assume we observe data generated from a bicycle model with unknown parameters for all conditions, in which all genes but gene i were directly perturbed, i.e., $P^{\neg i} = G \setminus \{i\}\ \forall i \in G$. If we the data suffice to estimate the corresponding covariance matrices $\omega^{\neg i}\ \forall i \in G$, the GRN β and the noise parameters $\sigma, \hat{\sigma}$ of the data generating model can be identified. If additionally the corresponding mean vectors $\bar{z}^{\neg i}$ can be estimated, the baseline transcription rate parameters $\alpha, \hat{\alpha}$ can be identified.*

Considering the causal effect/GRN matrix β, where β_{hg} corresponds to the effect from $h \to g$, we additionally define $\neg i$ as the intervention in which all genes except for gene i were knocked down. Now we have to show that given these data we can identify the true generative dynamical system in terms of its parameters $\alpha, \hat{\alpha}, \beta, \sigma$, and $\hat{\sigma}$. In this (rather unrealistic real-world) setting, we have G interventional distributions leading to G coupled Lyapunov equations:

$$B^{\neg i}\omega^{\neg i} + \omega^{\neg i}\left(B^{\neg i}\right)^T = \begin{pmatrix} \hat{\sigma}_1^2 & & & & & \dots & 0 \\ & \ddots & & & & & \vdots \\ & & \hat{\sigma}_{i-1}^2 & & & & \\ & & & \sigma_i^2 & & & \\ & & & & \hat{\sigma}_{i+1}^2 & & \\ \vdots & & & & & \ddots & \\ 0 & \dots & & & & & \hat{\sigma}_G^2 \end{pmatrix}.$$

If all genes but gene i are perturbed, the resulting modified causal graph can only contain edges from other genes to i. Or, in other words, gene i is the only remaining gene which can have causal parents in this intervention graph. Therefore, we can define $B^{\neg i}$ as $B^{\neg i} = \mathbb{I} - (S^{\neg i})^T$, where $S^{\neg i}$ is the modified matrix of causal effects in which only column i is non-zero:

$$S^{\neg i} = \begin{pmatrix} 0 & & \beta_{1,i} & & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & & \beta_{i-1,i} & & 0 \\ 0 & \dots \quad \dots & 0 & \dots \quad \dots & 0 \\ 0 & & \beta_{i+1,i} & & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & & \beta_{G,i} & & 0 \end{pmatrix} = \begin{pmatrix} 0 \dots 0 & b^i & 0 \dots 0 \end{pmatrix}.$$

Here the column vector b^i denotes the vector of dynamical coefficients which determine the influences of the current transcriptomic state on the transcription rate of gene i. We use the notation Thus, if we observe nonzero entries of the corresponding covariance matrix $\omega^{\neg i}$, they have to correspond to direct causal influences from other genes on i and we can determine all entries of the

vector b^i from this context. Using the definitions above, we observe the following identity:

$$\begin{aligned} B^{\neg i}\omega^{\neg i} + \omega^{\neg i}\left(B^{\neg i}\right)^T &= (\mathbb{I} - (S^{\neg i})^T)\omega^{\neg i} + \omega^{\neg i}(\mathbb{I} - S^{\neg i}) \\ &= 2\omega^{\neg i} - (S^{\neg i})^T\omega^{\neg i} - \omega^{\neg i}S^{\neg i} \\ &= 2\omega^{\neg i} - \begin{pmatrix} 0 \\ \vdots \\ 0 \\ (b^i)^T\omega^{\neg i} \\ 0 \\ \vdots \\ 0 \end{pmatrix} - \begin{pmatrix} 0 & \dots & 0 & \omega^{\neg i}b^i & 0 & \dots & 0 \end{pmatrix} \\ &= \begin{pmatrix} \hat{\sigma}_1^2 & & & & & & \\ & \ddots & & & & & \\ & & \hat{\sigma}_{i-1}^2 & & & & \\ & & & \sigma_i^2 & & & \\ & & & & \hat{\sigma}_{i+1}^2 & & \\ & & & & & \ddots & \\ & & & & & & \hat{\sigma}_G^2 \end{pmatrix} = \Sigma^{\neg i} \end{aligned}$$

Element-wise inspection of this equation yields the following identities:

1. $2\omega_{j,j}^{\neg i} = \hat{\sigma}_j^2 \quad \forall\, j \neq i$
2. $2\omega_{i,j}^{\neg i} = (\omega^{\neg i}b^i)_j \quad \forall\, j \neq i$
3. $2\omega_{i,i}^{\neg i} = \sigma_i^2 + 2(\omega^{\neg i}b^i)_i.$

Introducing the notation of a lower index $\neg i$ as dropping the i-th row or column of a matrix, we observe in identity (2) that $2\omega_{i,j}^{\neg i} = (\omega^{\neg i}b^i)_j$ for $j \neq i$ is equivalent to $2\omega_{\neg i,i}^{\neg i} = \omega_{\neg i,\neg i}^{\neg i}b_{\neg i}^i$, because $b_i^i = 0$ by definition. As the $(G-1)\times(G-1)$ submatrix $\omega_{\neg i,\neg i}^{\neg i}$ of the positive-definite matrix $\omega^{\neg i}$ is also positive definite, we can uniquely solve this system for the non-zero entries of b^i, which we can use to compute σ_i using identity (3) afterwards:

$$b_{\neg i}^i = 2(\omega_{\neg i,\neg i}^{\neg i})^{-1}\omega_{\neg i,i}^{\neg i}$$

Furthermore, using $B^{\neg i} = \mathbb{I} - (S^{\neg i})^T$, we can calculate the parameters α and $\hat{\alpha}$ via $\alpha^{\neg i} = B^{\neg i}\bar{z}^{\neg i}$. In summary, given the coupled Lyapunov equations $B^{\neg i}\omega^{\neg i} + \omega^{\neg i}\left(B^{\neg i}\right)^T = \Sigma^{\neg i}$ for all i, corresponding to perfectly observing the multivariate normal distributions of the latent expression states in all conditions in which all but one gene was perturbed, we can uniquely recover the generating parameters.

Theorem 3

Assume we observe data generated from a bicycle model with unknown parameters for all conditions, in which a single gene i was directly targeted, i.e., $P^i = \{i\}\ \forall i \in G$. If the data suffice to

estimate the corresponding covariance matrices $\omega^i,\ \forall i \in G$ *and if the submatrices* $\omega^I_{\neg j,\neg j}\ \forall j \in G$ *of the matrix* $\omega^I = \begin{pmatrix} \omega^1_{1,:} \\ \vdots \\ \omega^G_{G,:} \end{pmatrix}$ *are invertible, the GRN* β *and the noise parameters* $\sigma, \hat{\sigma}$ *of the data generating model can be identified. If additionally the corresponding mean vectors* $\bar{z}^i$ *can be estimated, the baseline transcription rate parameters* $\alpha, \hat{\alpha}$ *can be identified.*

Now, we consider the scenario where we are given interventions on each node/gene individually. This is a more realistic real-world case, found in most of the existing Perturb-seq datasets (e.g., Replogle et al. (2022)). If we assume we have access to the true resulting covariance matrices ω^i and means $\bar{z}^i$ of the steady-state distributions of the latent expression variables for all contexts in which a single gene i was perturbed, we can give a criterion on these covariance matrices which is sufficient to identify the true underlying dynamical system in terms of its parameters $\alpha, \hat{\alpha}, \beta, \sigma$, and $\hat{\sigma}$.
Given G coupled Lyapunov equations for each gene $i = 1, \ldots, G$:

$$B^i\omega^i + \omega^i \left(B^i\right)^T = \begin{pmatrix} \sigma_1^2 & & & & & \cdots & 0 \\ & \ddots & & & & & \vdots \\ & & \sigma_{i-1}^2 & & & & \\ & & & \hat{\sigma}_i^2 & & & \\ & & & & \sigma_{i+1}^2 & & \\ \vdots & & & & & \ddots & \\ 0 & \cdots & & & & & \sigma_G^2 \end{pmatrix},$$

B^i can be written as $B^i = \mathbb{I} - \sum_{j \neq i} (S^j)^T$, where S^j is the is the modified matrix of causal effects in which only column j is non-zero:

$$S^j = \begin{pmatrix} 0 & & & \beta_{1,j} & & & 0 \\ \vdots & & & \vdots & & & \vdots \\ 0 & & & \beta_{j-1,j} & & & 0 \\ 0 & \cdots & \cdots & 0 & \cdots & \cdots & 0 \\ 0 & & & \beta_{j+1,j} & & & 0 \\ \vdots & & & \vdots & & & \vdots \\ 0 & & & \beta_{G,j} & & & 0 \end{pmatrix} = \begin{pmatrix} 0 \ldots 0 & b^j & 0 \ldots 0 \end{pmatrix}.$$

Here the column vector b^j denotes the vector of dynamical coefficients which determine the influences of the current transcriptomic state on the transcription rate of gene j. Plugging these

definitions into the Lyapunov equation, we obtain:

$$
\begin{aligned}
B^i\omega^i + \omega^i \left(B^i\right)^T &= (\mathbb{I} - \sum_{j\neq i}(S^j)^T)\omega^i + \omega^i(\mathbb{I} - \sum_{j\neq i} S^j) \\
&= 2\omega^i - \sum_{j\neq i}(S^j)^T\omega^i - \omega^i \sum_{j\neq i} S^j \\
&= 2\omega^i - \sum_{j\neq i} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ (b^j)^T\omega^i \\ 0 \\ \vdots \\ 0 \end{pmatrix} - \sum_{j\neq i} \begin{pmatrix} 0 & \dots & 0 & \omega^i b^j & 0 & \dots & 0 \end{pmatrix} \\
&= \begin{pmatrix} \sigma_1^2 & & & & & & \\ & \ddots & & & & & \\ & & \sigma_{i-1}^2 & & & & \\ & & & \hat{\sigma}_i^2 & & & \\ & & & & \sigma_{i+1}^2 & & \\ & & & & & \ddots & \\ & & & & & & \sigma_G^2 \end{pmatrix} = \Sigma^{\neg i}
\end{aligned}
$$

Again, element-wise inspection of this equation yields the following identities:

1. $2\omega^i_{i,i} = \hat{\sigma}_i^2 \quad \forall\, i$
2. $2\omega^i_{j,j} = 2(\omega^i b^j)_j + \sigma_j^2 \quad \forall\, i \neq j$
3. $2\omega^i_{i,j} = (\omega^i b^j)_i \quad \forall\, i \neq j$
4. $2\omega^i_{j,k} = (\omega^i b^j)_k + (\omega^i b^k)_j \quad \forall$ pairwise distinct i, j, k

If we make use of identity (3) for all $i \neq j$ for a given j, we obtain:

$$
\begin{pmatrix} \omega^1_{1,:} \\ \vdots \\ \omega^{j-1}_{j-1,:} \\ \omega^{j+1}_{j+1,:} \\ \vdots \\ \omega^n_{G,:} \end{pmatrix} b^j_{\neg j} = 2 \begin{pmatrix} \omega^1_{1,j} \\ \vdots \\ \omega^{j-1}_{j-1,j} \\ \omega^{j+1}_{j+1,j} \\ \vdots \\ \omega^G_{G,j} \end{pmatrix},
$$

which we can rewrite using the notation introduced in Theorem 2 as $\omega^I_{\neg j,\neg j} b^j_{\neg j} = 2\omega^I_{\neg j,j}$, where

$$\omega^I = \begin{pmatrix} \omega^1_{1,:} \\ \vdots \\ \omega^G_{G,:} \end{pmatrix}.$$

Thus, if all $\omega^I_{\neg j,\neg j}$ are invertible, we can find a unique solution of b^j (again leveraging $b^j_j = 0$), which allows us to determine σ^j using identity (1) and α^j via $\alpha^j = B^j \bar{z}^j$. In conclusion, we have shown that if all the sub-matrices $\omega^i_{\neg j,\neg j}$ of the matrix ω^I, which is composed of the i-th rows of the covariance matrices ω^i of the conditions in which the individual gene i was perturbed, is invertible, we can uniquely identify the true parameters of the generative dynamics.