LEARNING EARLY TREATMENT STRATEGY FROM SNAPSHOTS FOR MRNA-PROTEIN REGULATORY NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Early disease detection with snapshot data has been effectively addressed by the Dynamical Network Biomarkers (DNBs) theory. After early disease detection, it is crucial to consider early medical treatment to prevent it. This paper presents a novel framework for identifying mRNA-protein regulatory systems from snapshot data and designing interventions. We estimate the state covariance of mRNA-protein expression using multi-episode snapshot samples and identify the underlying continuous-time dynamics by solving a Lyapunov-based regression problem. We provide finite-sample guarantees on the estimation accuracy of the system matrix and its dominant eigenvectors, which are essential for downstream treatment design. Building on these estimates, we formulate an optimal re-stabilization strategy that minimizes input energy with desired spectral shifts. To ensure practical feasibility, we further propose a diagonal re-stabilization scheme that identifies key regulatory nodes using a first-order eigenvalue sensitivity analysis. Validations on synthetic mRNA-protein network and real dataset demonstrate that our method accurately identifies regulatory nodes under snapshot data conditions and significantly outperforms existing baselines.

1 INTRODUCTION

In biological systems, mRNA regulation and protein interactions evolve through complex networks, and many diseases can be interpreted as abrupt deterioration of these networks via bifurcation phenomena (Briat et al., 2016; Chen et al., 2012; Sadria & Bury, 2024). Using snapshot data, the Dynamical Network Biomarker (DNB) theory detects the pre-disease stage immediately before a critical transition by identifying amplified fluctuations in specific nodes (Liu et al., 2015; Aihara et al., 2022). This early detection is important because intervention before full disease onset is often more effective and less invasive. Indeed, both empirical and theoretical studies have begun to investigate early intervention based on snapshot data, including directed-network settings, where the left eigenvector of the system matrix associated with the dominant eigenvalue plays a key role in input placement (Koizumi et al., 2020; Chen et al., 2022a; Yasukata et al., 2023; Shen et al., 2024a). However, this information is not directly available from standard covariance-based methods such as PCA, which makes system identification from snapshot data essential. More broadly, learning stochastic dynamics from snapshot data has become an active topic in machine learning (Song et al., 2021; Neklyudov et al., 2023; Tong et al., 2024). A dominant line of work first reconstructs temporal trajectories and then performs system identification, with important applications in single-cell analysis through trajectory inference, optimal transport, and Schrödinger bridge formulations (Bunne et al., 2024; Schiebinger et al., 2019; Zhang et al., 2025). These approaches are powerful for recovering likely state transitions, but they often require large sample sizes and substantial computation, and they are not designed to recover the control-relevant regulatory parameters needed for stabilization design. Gene regulatory systems are also intrinsically multiscale, with relatively slow transcriptional processes coupled to faster protein-level dynamics (Fletcher & Osborne, 2022). Classical mechanistic differential-equation models can capture such structure (Sadria & Layton, 2021a;b; Ingalls, 2013), but they usually require dense time-series data or carefully designed experiments, which are often unavailable in high-throughput single-cell settings. Recent representation-learning and sparse-discovery approaches improve fate prediction and dynamical model extraction from snapshot data (Sadria et al., 2022; Sadria & Swaroop, 2025), yet they do

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

not directly provide the identifiable, control-ready parameterization targeted in this work. A more detailed background and related work discussion is provided in Appendix A.

In contrast, in this paper, we leverage the structural properties of the mRNA–protein regulatory network to achieve reliable and computationally efficient system identification. This identified system enables the design of effective early treatment strategies. Our main contributions are summarized as follows: (a) **System identification from snapshot data:** We develop a framework that uses structural constraints of the mRNA–protein regulatory network to identify the system matrix from finite snapshot data via Lyapunov-based regression. (b) **Theoretical guarantees:** We establish finite-sample confidence bounds for both the system matrix estimation and the associated eigenvectors, ensuring reliability even under High-Dimension Low-Sample-Size (HDLSS) conditions. (c) **Early intervention design:** Building on the estimated system, we design both optimal and diagonal re-stabilization strategies, providing a practical approach for early treatment at the pre-disease stage. This is crucial as early (pre-disease) interventions are significantly more effective and less invasive than treatments applied after full disease onset.

2 Preliminaries and Backgrounds

Dynamic system for gene regulation. Gene transcription is regulated by transcription factors that bind to DNA, with transcription rates modulated by their concentrations. Translated proteins can further regulate gene expression. Translation lacks feedback to mRNA, and both mRNAs and proteins degrade stochastically. This gene-mRNA-protein feedback is modeled by (Liu et al., 2016; Passemiers et al., 2022; Sanders et al., 2020; Weidmann et al., 2021):

$$\dot{\mathbf{z}} = \mathbf{F}_{\boldsymbol{\theta}}(\mathbf{z}) + \mathbf{w}, \quad \mathbf{F}_{\boldsymbol{\theta}}(\mathbf{z}) = \begin{bmatrix} -\Delta_{\mathsf{r},\boldsymbol{\theta}}\mathbf{z}_{\mathsf{r}} + \mathbf{f}_{\boldsymbol{\theta}}(\mathbf{z}_{\mathsf{p}}) \\ \Gamma_{\mathsf{r},\boldsymbol{\theta}}\mathbf{z}_{\mathsf{r}} - \Delta_{\mathsf{p},\boldsymbol{\theta}}\mathbf{z}_{\mathsf{p}} \end{bmatrix}. \tag{1}$$

where $\mathbf{z} := [\mathbf{z}_{\mathsf{r}}, \mathbf{z}_{\mathsf{p}}]^\top \in \mathbb{R}^{2n}$ denotes mRNA and protein concentrations, and $\mathbf{w}$ is isotropic Gaussian white noise with known covariance $\mathbf{D} = \sigma^2\mathbf{I}_{2n}$, $\sigma > 0$, matching the noise model used in our simulations. Transcription function $\mathbf{f}_{\boldsymbol{\theta}}(\mathbf{z}_{\mathsf{p}})$ is a nonlinear n-dimensional vector encoding protein-mediated regulation. Diagonal matrices $\Gamma_{\mathsf{r},\boldsymbol{\theta}}$, $\Delta_{\mathsf{r},\boldsymbol{\theta}}$, and $\Delta_{\mathsf{p},\boldsymbol{\theta}}$ represent translation and degradation rates and are all non-degenerate. These quantities are parameterized by $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^m$. Let $\mathbf{z}^{\mathsf{e}}$ be the equilibrium point such that $\mathbf{F}_{\boldsymbol{\theta}}(\mathbf{z}) = 0$, where $\mathbf{z}^{\mathsf{e}} = [\mathbf{z}_{\mathsf{r}}^{\mathsf{e}}, \mathbf{z}_{\mathsf{p}}^{\mathsf{e}}]^\top$. A linearized approximation around $\mathbf{z}^{\mathsf{e}}$ can be written by (Chen et al., 1999). When only mRNA-level measurements are available (as in Section 6 biological data), a standard quasi-stationarity argument for fast protein degradation ($\gamma_{p,i} \gg \gamma_{m,i}$) reduces the system to an n-dimensional mRNA-only SDE that preserves the sparsity structure of $\mathbf{A}_{\boldsymbol{\theta}}$; details are given in Appendix K.1.

$$\dot{\mathbf{x}} = \mathbf{A}_{\boldsymbol{\theta}}\mathbf{x} + \mathbf{w}, \ \mathbf{x} = \mathbf{z} - \mathbf{z}^{\mathsf{e}}, \ \mathbf{A}_{\boldsymbol{\theta}} = \begin{bmatrix} -\Delta_{\mathsf{r},\boldsymbol{\theta}} & \Gamma_{\mathsf{p},\boldsymbol{\theta}} \\ \Gamma_{\mathsf{r},\boldsymbol{\theta}} & -\Delta_{\mathsf{p},\boldsymbol{\theta}} \end{bmatrix}, \ \Gamma_{\mathsf{p},\boldsymbol{\theta}} = \left.\frac{\partial \mathbf{f}_{\boldsymbol{\theta}}(\mathbf{z}_{\mathsf{p}})}{\partial \mathbf{z}_{\mathsf{p}}}\right|_{\mathbf{z}=\mathbf{z}^{\mathsf{e}}}. \tag{2}$$

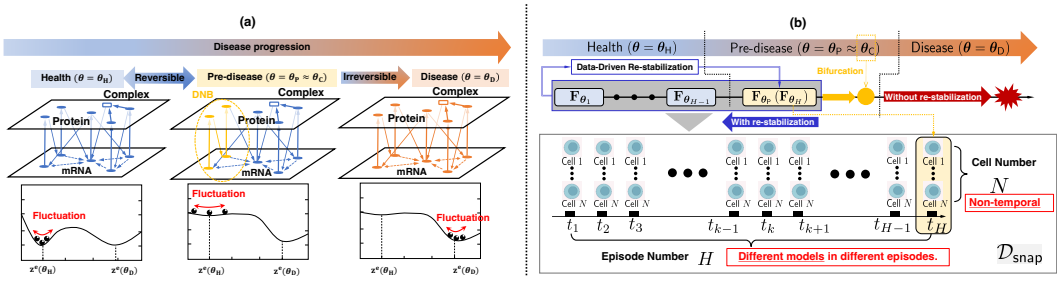

Figure 1: Conceptual illustration of the disease progression and snapshot dataset: (a) Illustration of the disease progression from the health stage to the disease stage through the pre-disease stage. (b) Illustration of the snapshot data-driven re-stabilization and the issue for system identification.

Disease progress and snapshot data. Figure 1 (a) shows chronic disease progression driven by the parameter $\boldsymbol{\theta}$, where a tipping point $\boldsymbol{\theta}_{\mathsf{C}}$ separates the healthy stage Θ_{H} from the disease stage

Θ_{D}. This transition corresponds to a bifurcation in the nonlinear system (Chen et al., 2012). Let $\overline{\Theta}_{\mathsf{H}} := \Theta_{\mathsf{H}} \cup \partial\Theta_{\mathsf{H}}$ denote the closure of the healthy region. As $\boldsymbol{\theta}$ evolves within $\Theta := \overline{\Theta}_{\mathsf{H}} \cup \Theta_{\mathsf{D}}$, the equilibrium $\mathbf{z}^{\mathsf{e}}_{\boldsymbol{\theta}}$ changes accordingly. In $\overline{\Theta}_{\mathsf{H}}$, the equilibrium is stable and approximated by $\mathbf{z}^{\mathsf{e},\mathsf{H}}$, while in Θ_{D}, it shifts to $\mathbf{z}^{\mathsf{e},\mathsf{D}}$, far from the healthy state (bottom of Figure 1 (a)). As shown in Figure 1 (a), the system is stable in both Θ_{H} and Θ_{D}, since the maximum eigenvalue $\lambda_{\mathsf{d},\boldsymbol{\theta}}$ of $\mathbf{A}_{\boldsymbol{\theta}}$ has a significantly negative real part. At the tipping point $\boldsymbol{\theta} \in \partial\Theta_{\mathsf{H}}$, this eigenvalue becomes zero. In the pre-disease stage $\boldsymbol{\theta}_{\mathsf{P}} \approx \boldsymbol{\theta}_{\mathsf{C}}$, we have the real part $\mathsf{Re}\{\lambda_{\mathsf{d},\boldsymbol{\theta}}\} \approx 0^-$, indicating low stability and high sensitivity to perturbations. As illustrated in Figure 1(b), measurements of the system state $\mathbf{x}$ (mRNA and protein levels) are collected over H episodes. Each episode k corresponds to a biological sampling time t_k, where N individual cells are measured. These single-cell observations form a snapshot of the internal state: $\mathcal{D}^{(k)}_{\mathsf{snap}} := \left\{\mathbf{x}^{(k)}_m\right\}_{m=1}^{N}, \quad \mathbf{x}^{(k)}_m \in \mathbb{R}^{2n}$. Such data are not time-series but population-level samples reflecting heterogeneity at t_k. The full dataset is denoted by $\mathcal{D}_{\mathsf{snap}} := \left\{\mathcal{D}^{(k)}_{\mathsf{snap}}\right\}_{k=1}^{H}$.

Assumption 1. *For each $k = 1, \ldots, H$, the following holds: (1) The samples $\left\{\mathbf{x}^{(k)}_m\right\}_{m=1}^{N}$ are i.i.d. from a continuous distribution μ_k: $\mathbf{x}^{(k)}_m \sim \mu_k$. (2) $\boldsymbol{\theta}$ evolves smoothly across time t_k*[1].

This setting reflects realistic biological measurement conditions, where mRNA-protein regulatory systems are quasi-stationary during sampling, and population-level expression distributions shift gradually due to slow parameter changes. Note that the i.i.d. assumption applies only to the sample index m, and no independence across time step k is assumed.

3 Addressed Problem and Challenging Issues

The pre-disease stage can be efficiently detected from snapshot data. Details are summarized in Appendix B. Once this stage is identified, we consider interventions that prevent further progression and restore the gene regulatory network toward the healthy regime. Such early intervention is not only clinically desirable for reducing patient burden, but is also generally more effective and less invasive than treatment after full disease onset. Recovery from the disease stage requires steering the system from a diseased equilibrium $\mathbf{z}^{\mathsf{e}}_{\boldsymbol{\theta}}$ with $\boldsymbol{\theta} \in \Theta_{\mathsf{D}}$ back to a healthy equilibrium with $\boldsymbol{\theta} \in \Theta_{\mathsf{H}}$, which is a nonlinear control problem due to the large deviation between equilibria and the complexity of the dynamics. By contrast, intervention at the pre-disease stage is more tractable. Since the equilibrium $\mathbf{z}^{\mathsf{e}}_{\boldsymbol{\theta}}$ changes only slightly for $\boldsymbol{\theta} \in \overline{\Theta}_{\mathsf{H}}$, the problem can be reasonably approximated as a linear control task. In this setting, the goal is not to shift the equilibrium itself, but to improve robustness against perturbations by enlarging the local stability margin. To distinguish this objective from full recovery, we refer to it as re-stabilization. Specifically, re-stabilization shifts the dominant eigenvalue $\lambda_{\mathsf{d},\boldsymbol{\theta}}$ further into the left-half complex plane, thereby increasing resilience while the system remains near the healthy regime. This is achieved by introducing a feedback loop that modifies the system matrix to $\mathbf{A}_{\boldsymbol{\theta}} + \mathbf{B}_{\boldsymbol{\theta}}\mathbf{K}_{\boldsymbol{\theta}}$, where $\mathbf{B}_{\boldsymbol{\theta}} = [\mathbf{b}_{1,\boldsymbol{\theta}}, \ldots, \mathbf{b}_{l,\boldsymbol{\theta}}] \in \mathbb{R}^{2n\times l}$ specifies the input placement, i.e., the selected genes for intervention. We now formulate the re-stabilization problem with snapshot dataset $\mathcal{D}_{\mathsf{snap}}$:

Problem 1 (Re-stabilization problem)**.** *At each episode $k = 1, \ldots, H$, the parameter of the transcription function takes the value $\boldsymbol{\theta}_k$, and the true system matrix $\mathbf{A}_{\boldsymbol{\theta}_k}$ in the linearized dynamics* (2) *is unknown. Given a snapshot dataset $\mathcal{D}_{\mathsf{snap}}$, the objective is to design a feedback intervention $\mathbf{B}_{\boldsymbol{\theta}_H}\mathbf{K}_{\boldsymbol{\theta}_H}$ such that the dominant eigenvalue of the closed-loop system $\tilde{\lambda}_{\mathsf{d},\boldsymbol{\theta}_H} := \max \mathsf{Re}\left(\mathrm{eig}\left(\mathbf{A}_{\boldsymbol{\theta}_H} + \mathbf{B}_{\boldsymbol{\theta}_H}\mathbf{K}_{\boldsymbol{\theta}_H}\right)\right)$ satisfies $\mathsf{Re}(\tilde{\lambda}_{\mathsf{d},\boldsymbol{\theta}_H}) < \mathsf{Re}(\lambda_{\mathsf{d},\boldsymbol{\theta}_H})$, realizing re-stabilization.*

Enlarging this margin makes the state distribution concentrate around the healthy equilibrium, so that even close to a critical transition the system is less likely to drift into the disease regime; the biological reading of this effect is expanded in Appendix D. Identifying regulatory nodes therefore yields not only an early-warning signal but actionable intervention targets.

Challenging issues. Solving Problem 1 faces several key challenges. First, the snapshot dataset $\mathcal{D}_{\mathsf{snap}}$ collected at the pre-disease stage ($k = H$), $\mathcal{D}^{(H)}_{\mathsf{snap}}$, suffers from the high-dimensional, low-sample-size (HDLSS) regime, as it typically contains expression measurements for $n > 10^4$ genes across a few thousand single cells at most. That is, the number of molecular features far exceeds the number

[1]"Smooth" means that $\boldsymbol{\theta}$ is a smooth function of time t, indicating a smooth parameter evolution in time.

of observable cells, creating severe challenges for statistical inference on the system matrix $\mathbf{A}_{\boldsymbol{\theta}_H}$ using only the available data at t_H. Second, the dataset $\mathcal{D}_{\mathsf{snap}}$ consists of non-temporal (snapshot) observations, which precludes the use of system identification techniques that rely on time-series trajectories. Furthermore, although snapshot data from earlier episodes are available, they originate from distinct underlying models due to variation in $\boldsymbol{\theta}$ across time and therefore cannot be directly pooled with data from the pre-disease stage. As a result, classical control techniques such as pole placement, which require full knowledge of the system matrix, are not applicable. These issues necessitate the development of a novel data-driven approach capable of designing re-stabilizing feedback using only distributional information extracted from limited snapshot data.

4 Proposed method: Generative Re-stabilization

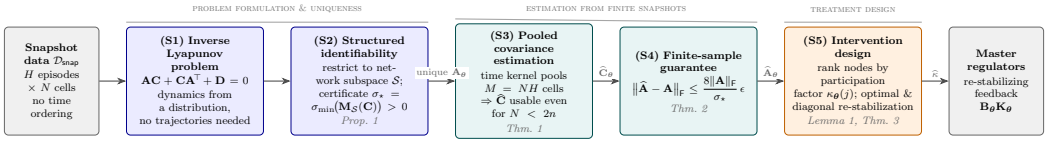

Figure 2: Pipeline of the proposed method: problem formulation and identifiability (S1–S2), estimation from finite snapshots (S3–S4), and treatment design (S5), with the supporting guarantees indicated in each step.

The proposed method is summarised in Figure 2. It turns snapshot data into intervention targets through five steps, each supplying exactly what the next one needs. **(S1)** Identification is recast as an inverse Lyapunov problem, removing the need for time-series trajectories. **(S2)** The problem is restricted to the structured subspace carrying the network's biology, where a single computable scalar certificate decides uniqueness. **(S3)** The covariance is estimated by pooling cells across episodes, which makes the regime $N < 2n$ tractable. **(S4)** Finite-sample covariance accuracy becomes a guarantee on the identified system matrix, and **(S5)** its spectrum ranks intervention targets by their participation in the critical mode, closing the loop from data to treatment design.

(S1) Snapshot identification as an inverse Lyapunov problem. At the steady state, the covariance matrix of $\mathbf{x}$, denoted by $\mathbf{C}_{\boldsymbol{\theta}} \in \mathbb{R}^{2n\times 2n}$, satisfies the Lyapunov equation (Chen et al., 2012; Oku & Aihara, 2018):

$$\mathbf{A}_{\boldsymbol{\theta}}\mathbf{C}_{\boldsymbol{\theta}} + \mathbf{C}_{\boldsymbol{\theta}}\mathbf{A}_{\boldsymbol{\theta}}^\top + \mathbf{D} = 0. \tag{3}$$

Equation (3) is the pivot of the whole approach: it links the covariance of the stationary distribution algebraically to the system matrix, so identification needs only a distribution over cells and never a temporal ordering. This is what converts an apparently impossible task — recovering dynamics from non-temporal data — into a regression problem on covariance matrices, which single-cell measurements estimate natively. Vectorising (3) with $\mathscr{A}_{\boldsymbol{\theta}} = \mathsf{vec}(\mathbf{A}_{\boldsymbol{\theta}})$, $\mathscr{C}_{\boldsymbol{\theta}} = \mathbf{C}_{\boldsymbol{\theta}} \otimes \mathbf{I} + (\mathbf{I} \otimes \mathbf{C}_{\boldsymbol{\theta}})\mathbf{T}$ and $\mathscr{D} = \mathsf{vec}(\mathbf{D})$, where $\mathbf{T}$ satisfies $\mathsf{vec}(\mathbf{X}) = \mathbf{T}\mathsf{vec}(\mathbf{X}^\top)$, gives the linear equation

$$\mathscr{C}_{\boldsymbol{\theta}}\mathscr{A}_{\boldsymbol{\theta}} = -\mathscr{D}. \tag{4}$$

Since $\mathbf{C}_{\boldsymbol{\theta}}$ and $\mathbf{D}$ are symmetric, (4) has rank at most $n(2n+1)$ and therefore cannot by itself determine all $4n^2$ entries of $\mathscr{A}_{\boldsymbol{\theta}}$. The structure of the model supplies the missing equations: because $\Delta_{\mathsf{r},\boldsymbol{\theta}}, \Delta_{\mathsf{p},\boldsymbol{\theta}}, \Gamma_{\mathsf{r},\boldsymbol{\theta}}$ are diagonal, $3n^2 - 3n$ entries of $\mathbf{A}_{\boldsymbol{\theta}}$ are known to vanish. Writing $\mathbf{u}^{(ij)} \in \mathbb{R}^{1\times 4n^2}$ for the row with a single 1 in position $i + 2n(j-1)$, so that $\mathbf{u}^{(ij)}\mathscr{A}_{\boldsymbol{\theta}} = \mathbf{A}_{\boldsymbol{\theta}}(ij)$, and stacking these rows over the $3n^2 - 3n$ known-zero pairs into $\mathbf{U} \in \mathbb{R}^{(3n^2-3n)\times 4n^2}$ gives

$$\mathbf{U}\mathscr{A}_{\boldsymbol{\theta}} = \mathbf{0}. \tag{5}$$

Appending (5) to (4) yields the extended system that is solved in practice,

$$\mathscr{C}_{\boldsymbol{\theta},\mathsf{ext}}\mathscr{A}_{\boldsymbol{\theta}} = \mathscr{D}_{\mathsf{ext}}, \qquad \mathscr{C}_{\boldsymbol{\theta},\mathsf{ext}} := \begin{bmatrix}\mathscr{C}_{\boldsymbol{\theta}} \\ \mathbf{U}\end{bmatrix}, \qquad \mathscr{D}_{\mathsf{ext}} := \begin{bmatrix}-\mathscr{D} \\ \mathbf{0}\end{bmatrix}. \tag{6}$$

(S2) Structured identifiability and a computable certificate. Equation (6) has more rows than unknowns, but that alone does not settle uniqueness, which is a statement about its column rank. We therefore work directly in the subspace picked out by (5), namely

$$\mathcal{S} := \left\{ \Delta\mathbf{A} = \begin{bmatrix}\mathbf{D}_{\mathsf{r}} & \mathbf{B} \\ \mathbf{G} & \mathbf{D}_{\mathsf{p}}\end{bmatrix} : \mathbf{D}_{\mathsf{r}}, \mathbf{D}_{\mathsf{p}}, \mathbf{G} \text{ diagonal}, \ \mathbf{B} \in \mathbb{R}^{n\times n} \right\}, \tag{7}$$

of dimension $p := n^2 + 3n$, which counts the free parameters of the network: the n diagonal entries of each of $\mathbf{D}_\mathsf{r}$, $\mathbf{D}_\mathsf{p}$ and $\mathbf{G}$, together with the n^2 entries of $\mathbf{B}$. Index these free positions by $k = 1, \dots, p$, and let $\mathbf{E}_k$ be the matrix that has a 1 in the k-th free position and zeros everywhere else. The $\mathbf{E}_k$ form a basis of $\mathcal{S}$, so every $\Delta\mathbf{A} \in \mathcal{S}$ can be written $\Delta\mathbf{A} = \sum_{k=1}^{p} a_k \mathbf{E}_k$, where the coefficient vector $a \in \mathbb{R}^p$ simply lists the free entries of $\Delta\mathbf{A}$. Next, for any $\Delta\mathbf{A}$ the residual $\Delta\mathbf{A}\mathbf{C} + \mathbf{C}\Delta\mathbf{A}^\top$ is a symmetric $2n \times 2n$ matrix and is therefore already determined by its lower triangle; we write $\mathsf{vech}(\cdot)$ for the operator that stacks those $n(2n+1)$ lower-triangular entries into one column vector, which keeps each scalar equation exactly once. Collecting the image of each basis direction as a column gives the matrix of the Lyapunov map restricted to $\mathcal{S}$,

$$\mathbf{M}_{\mathcal{S}}(\mathbf{C}) := \Big[\, \mathsf{vech}\big(\mathbf{E}_1\mathbf{C} + \mathbf{C}\mathbf{E}_1^\top\big) \;\Big|\; \cdots \;\Big|\; \mathsf{vech}\big(\mathbf{E}_p\mathbf{C} + \mathbf{C}\mathbf{E}_p^\top\big) \,\Big] \;\in\; \mathbb{R}^{n(2n+1)\times p}, \tag{8}$$

so that $\mathbf{M}_{\mathcal{S}}(\mathbf{C})\, a = \mathsf{vech}\big(\Delta\mathbf{A}\mathbf{C} + \mathbf{C}\Delta\mathbf{A}^\top\big)$ for every $a \in \mathbb{R}^p$. Writing $d := -\mathsf{vech}(\mathbf{D})$, the Lyapunov equation (3) restricted to $\mathcal{S}$ is thus the linear system $\mathbf{M}_{\mathcal{S}}(\mathbf{C}_{\boldsymbol{\theta}})\, a = d$ in the p unknowns a, and $\mathbf{A}_{\boldsymbol{\theta}}$ is its unique solution in $\mathcal{S}$ iff

$$\mathsf{rank}\, \mathbf{M}_{\mathcal{S}}(\mathbf{C}_{\boldsymbol{\theta}}) = p, \quad \text{equivalently} \quad \sigma_\star := \sigma_{\min}(\mathbf{M}_{\mathcal{S}}(\mathbf{C}_{\boldsymbol{\theta}})) > 0, \tag{9}$$

where $\sigma_{\min}(\cdot)$ denotes the smallest singular value, i.e. the p-th largest here, so that $\sigma_\star > 0$ is exactly full column rank. A single computable scalar therefore carries the whole identifiability question, and the next proposition shows it is positive for almost every network in the class.

Proposition 1 (Generic structured identifiability)**.** *Fix the known isotropic noise covariance $\mathbf{D} = \sigma^2\mathbf{I}_{2n}$ with $\sigma > 0$. Let $n \geq 3$, let $\mathcal{S}_{\mathsf{stab}} := \{\mathbf{A} \in \mathcal{S} : \mathbf{A}$ Hurwitz$\}$, and let $\mathbf{C}(\mathbf{A}) \succ 0$ be the solution of* (3) *at $\mathbf{A}$. Let $\mathcal{O} \subseteq \mathcal{S}_{\mathsf{stab}}$ be a connected component that contains at least one $\mathbf{A}^\star$ with $\sigma_{\min}(\mathbf{M}_{\mathcal{S}}(\mathbf{C}(\mathbf{A}^\star))) > 0$. Then the exceptional set $\mathcal{E} := \{\mathbf{A} \in \mathcal{O} : \mathsf{rank}\, \mathbf{M}_{\mathcal{S}}(\mathbf{C}(\mathbf{A})) < p\}$ is closed in $\mathcal{O}$ and has Lebesgue measure zero in $\mathbb{R}^p$. Such a component exists: an explicit $\mathbf{A}^\star$, of the master-regulator form used throughout this paper, is exhibited in closed form in Appendix E. Consequently $\mathbf{A}_{\boldsymbol{\theta}}$ is identifiable from $\mathbf{C}_{\boldsymbol{\theta}}$ for almost every parameter in that component of the model class.*

The proof is in Appendix E: analyticity of $\mathbf{A} \mapsto \mathbf{C}(\mathbf{A})$ reduces genericity to a single explicit witness, which is exhibited in closed form. Two features make (9) usable rather than decorative. It is quantified over the model parameters, i.e. over the covariances the network can actually produce, which is the statement the estimator ranges over. It is computable: $\sigma_\star$ is obtained from the estimated covariance by (8), so identifiability is certified at run time rather than assumed.

Identification is then the structured least-squares problem in these p unknowns,

$$\min_{a \in \mathbb{R}^p} \; J\,(a, \mathbf{C}_{\boldsymbol{\theta}}) := \|\mathbf{M}_{\mathcal{S}}(\mathbf{C}_{\boldsymbol{\theta}})\, a - d\|_2^2\,, \tag{$\mathcal{P}_{\boldsymbol{\theta}}$}$$

from which the system matrix is recovered as $\mathbf{A}_{\boldsymbol{\theta}} = \sum_k a_k \mathbf{E}_k$. Solving the extended system (6) in the full $4n^2$ coordinates is the equivalent implementation form of the same problem, the structure being imposed there by the rows of $\mathbf{U}$ instead of by the choice of unknowns. In the coordinates a, J is a quadratic with constant Hessian $\nabla^2 J = 2\,\mathbf{M}_{\mathcal{S}}(\mathbf{C}_{\boldsymbol{\theta}})^\top \mathbf{M}_{\mathcal{S}}(\mathbf{C}_{\boldsymbol{\theta}})$, whose smallest eigenvalue is $\sigma_\star^2$. Hence, under (9), J is strongly convex with modulus $\mu = 2\sigma_\star^2$ and has a unique minimiser, and that minimiser is $\mathbf{A}_{\boldsymbol{\theta}}$ itself because J vanishes there and is positive elsewhere.

(S3) Covariance estimation by pooling episodes. Since $\mathbf{C}_{\boldsymbol{\theta}}$ is unknown, we replace it by an estimate $\widehat{\mathbf{C}}_{\boldsymbol{\theta}}$ and solve

$$\min_{a \in \mathbb{R}^p} \; J\left(a, \widehat{\mathbf{C}}_{\boldsymbol{\theta}}\right) := \left\|\mathbf{M}_{\mathcal{S}}(\widehat{\mathbf{C}}_{\boldsymbol{\theta}})\, a - d\right\|_2^2, \tag{$\widehat{\mathcal{P}}_{\boldsymbol{\theta}}\,(\mathcal{D}_{\mathsf{snap}})$}$$

writing $\widehat{\mathbf{A}}_{\boldsymbol{\theta}}$ for the resulting estimate of $\mathbf{A}_{\boldsymbol{\theta}}$. The estimate $\widehat{\mathbf{C}}_{\boldsymbol{\theta}}$ is built from a Nadaraya–Watson conditional density estimator $\widehat{p}^{\,\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t)$ with a product kernel in $(\mathbf{x}, t)$ and bandwidth h satisfying $h \to 0$, $NHh^{2n} \to \infty$ (Gooijer & Zerom, 2003; Hall et al., 1999); explicit formulae are in Appendix F. The time kernel is the operative ingredient: it lets neighbouring episodes contribute to the estimate at t_k, so the effective sample size is the pooled count $M = NH$ rather than the N cells of one snapshot. Since a single biological snapshot cannot supply arbitrarily many cells, this is what converts an HDLSS obstruction into a tractable estimation problem. The true conditional density is Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{C}_{\boldsymbol{\theta}})$, and with $\{\widehat{\mathbf{x}}_m^{(k)}\}_{m=1}^{M_\mathsf{s}}$ drawn i.i.d. from $\widehat{p}^{\,\mathsf{c}}_{\mathbf{X}}(\cdot \mid t_k)$ we set

$\widehat{\mathbf{C}}_{\boldsymbol{\theta}} := \frac{1}{M_{\mathsf{s}}} \sum_{m=1}^{M_{\mathsf{s}}} \widehat{\mathbf{x}}_m^{(k)} (\widehat{\mathbf{x}}_m^{(k)})^\top$. We keep M_{s}, the number of samples entering the estimator, distinct from the pooled data size $M = NH$: the former is what concentration acts on, the latter is how much data exists. For the estimator used here the two coincide, $M_{\mathsf{s}} = M = NH$, because the time kernel gives every cell of every episode a non-zero weight at t_k.

Theorem 1 (Consistency). *Assume $\sigma_\star > 0$. If $h \to 0$ and $NHh^{2n} \to \infty$, then $\widehat{\mathbf{A}}_{\boldsymbol{\theta}} \xrightarrow{p} \mathbf{A}_{\boldsymbol{\theta}}$ as the pooled sample size $M = NH \to \infty$.*

The proof is in Appendix G. The rate is set by the pooled count $M = NH$ rather than by the N cells of a single episode, which is exactly what the time kernel buys and is the regime the method is designed for. The hypothesis $\sigma_\star > 0$ is not an additional modelling restriction: by Proposition 1 it holds for almost every parameter, and it is verifiable from the data at hand.

(S4) Finite-sample accuracy of the system matrix. For the finite-sample regime we use the following standard boundedness condition.

Assumption 2. *There exists $L_\beta > 0$ such that, with probability at least $1 - \beta$,* all *of the samples $\mathbf{x}$ drawn from $\widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k)$ that are used to form $\widehat{\mathbf{C}}_{\boldsymbol{\theta}}$ satisfy $\|\mathbf{x}\|_\infty \le L_\beta$.*

Assumption 2 holds whenever the estimated density has bounded or sufficiently concentrated support, and it is what turns concentration of the sample covariance into a bound on the identified matrix.

Theorem 2 (Finite-sample identification error). *Let $\sigma_\star > 0$ and $0 < \epsilon \le \sigma_\star/4$, and suppose the bandwidth is fine enough that the kernel smoothing bias contributes at most $\epsilon/2$ to the covariance error (stated precisely as (27) in Appendix H). If the pooled sample size satisfies $M_{\mathsf{s}} = M = NH \ge 32n^2\epsilon^{-2}L_\beta^4 \log\left(16n^2/\delta\right)$, then with probability at least $1 - \beta - \delta$,*

$$\left\| \widehat{\mathbf{A}}_{\boldsymbol{\theta}} - \mathbf{A}_{\boldsymbol{\theta}} \right\|_{\mathsf{F}} \le \frac{8 \left\| \mathbf{A}_{\boldsymbol{\theta}} \right\|_{\mathsf{F}}}{\sigma_\star} \epsilon.$$

The proof is in Appendix H. The error scales as $1/\sigma_\star$, so the identifiability margin of (9) is not only the condition for uniqueness but also the exact quantity governing statistical accuracy: a single computable scalar simultaneously certifies identifiability and predicts the error, and networks approaching $\mathcal{E}$ are flagged by a small $\sigma_\star$ rather than failing silently.

(S5) From the identified spectrum to the intervention. Optimal re-stabilization design considers

$$\begin{aligned} \min_{\mathbf{B}_{\boldsymbol{\theta}},\, \mathbf{K}_{\boldsymbol{\theta}}} \quad & J_\theta(\mathbf{x}(0)) := \int_0^\infty \mathbf{x}^\top(t) \mathbf{K}_{\boldsymbol{\theta}}^\top \mathbf{K}_{\boldsymbol{\theta}} \mathbf{x}(t) \mathrm{d}t,\ \forall \mathbf{x}(0) \\ \text{s.t.} \quad & \mathsf{Re}(\lambda_{\mathsf{d},\boldsymbol{\theta}}) - \mathsf{Re}(\tilde{\lambda}_{\mathsf{d},\boldsymbol{\theta}}) = \lambda_{\mathsf{s}} > 0,\ \mathbf{B}_{\boldsymbol{\theta}} \in \mathcal{B},\ \mathbf{K}_{\boldsymbol{\theta}} \in \mathcal{K}(\mathbf{B}_{\boldsymbol{\theta}}, \lambda_{\mathsf{s}}), \end{aligned} \tag{10}$$

minimising input energy subject to achieving the re-stabilization margin λ_{s} for any initial state, where $\mathcal{B}$ and $\mathcal{K}(\mathbf{B}_{\boldsymbol{\theta}}, \lambda_{\mathsf{s}})$ are the admissible input placements and gains. By Yasukata et al. (2023, Theorem 1) the optimum is $\mathbf{B}_{\boldsymbol{\theta}}^\star = \arg\max_{\mathbf{B} \in \mathcal{B}} \|\mathbf{w}_{\mathsf{d},\boldsymbol{\theta}}^\top \mathbf{B}\|$, $\mathbf{K}_{\boldsymbol{\theta}}^\star = -\lambda_{\mathsf{s}} \mathbf{B}_{\boldsymbol{\theta}}^{\star\top} \mathbf{v}_{\mathsf{d},\boldsymbol{\theta}} \mathbf{w}_{\mathsf{d},\boldsymbol{\theta}}^\top$, where $\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}$ and $\mathbf{w}_{\mathsf{d},\boldsymbol{\theta}}$ are, respectively, the right and left eigenvectors of the open-loop system matrix $\mathbf{A}_{\boldsymbol{\theta}}$ associated with its dominant eigenvalue $\lambda_{\mathsf{d},\boldsymbol{\theta}}$. The design is therefore governed entirely by the dominant spectral data of $\mathbf{A}_{\boldsymbol{\theta}}$, which is exactly what step (S4) delivers with a guarantee.

Since many mRNAs and proteins cannot be actuated simultaneously, both the placement and the gain must be sparse. Each column is a distinct standard basis vector, $\mathbf{b}_{i,\boldsymbol{\theta}} \in \mathcal{E}_n^{\mathsf{r}} := \{\mathbf{e}_1, \ldots, \mathbf{e}_n\}$ for mRNA-level interventions or $\mathbf{b}_{i,\boldsymbol{\theta}} \in \mathcal{E}_n^{\mathsf{p}} := \{\mathbf{e}_{n+1}, \ldots, \mathbf{e}_{2n}\}$ for protein-level ones, with one modality per strategy, and $\mathbf{k}_{i,\boldsymbol{\theta}} = g_i \mathbf{b}_{i,\boldsymbol{\theta}}$ with scalar gains g_i, so the feedback acts only along the selected site; the feasible sets are written $\mathcal{B}_l$ and $\mathcal{K}_l$. The feedback then alters only the diagonal of $\mathbf{A}_{\boldsymbol{\theta}}$, which matches how interventions such as RNA interference and gene overexpression act on the self-dynamics of individual genes (Meister et al., 2013).

Definition 1. *Given an intervention budget $l < n$, diagonal re-stabilization refers to the design of $\mathbf{B}_{\boldsymbol{\theta}} \in \mathcal{B}_l$ and $\mathbf{K}_{\boldsymbol{\theta}} \in \mathcal{K}_l$ such that the dominant eigenvalue of the closed-loop system $\tilde{\lambda}_{\mathsf{d},\boldsymbol{\theta}} := \max \mathsf{Re}\left(\mathrm{eig}\left(\mathbf{A}_{\boldsymbol{\theta}} + \mathbf{B}_{\boldsymbol{\theta}} \mathbf{K}_{\boldsymbol{\theta}}\right)\right)$ satisfies $\mathsf{Re}(\tilde{\lambda}_{\mathsf{d},\boldsymbol{\theta}}) < \mathsf{Re}(\lambda_{\mathsf{d},\boldsymbol{\theta}})$.*

The design then reduces to choosing which node to actuate, and the correct figure of merit is that node's share of the critical mode. Let $\lambda_{\mathsf{d},\boldsymbol{\theta}}$ be the dominant eigenvalue of $\mathbf{A}_{\boldsymbol{\theta}}$, with right and left eigenvectors $\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}$ and $\mathbf{w}_{\mathsf{d},\boldsymbol{\theta}}$, and let

$$\delta_\lambda := \min_{\lambda \in \mathsf{spec}(\mathbf{A}_{\boldsymbol{\theta}}),\ \lambda \neq \lambda_{\mathsf{d},\boldsymbol{\theta}}} |\lambda_{\mathsf{d},\boldsymbol{\theta}} - \lambda| > 0 \tag{11}$$

be its separation from the rest of the spectrum. The participation factor of node j is

$$\kappa_{\boldsymbol{\theta}}(j) := \mathbf{w}_{\mathsf{d},\boldsymbol{\theta}}(j)\,\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}(j)\big/\mathbf{w}_{\mathsf{d},\boldsymbol{\theta}}^{\top}\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}, \tag{12}$$

which is the j-th diagonal entry of the spectral projector onto the critical mode; where the parameter is clear from context we abbreviate $\kappa_i := \kappa_{\boldsymbol{\theta}}(i)$. When several modules become critical at once the dominant eigenvalue is repeated, and (12) is then read off the projector onto the whole dominant eigenspace; this changes none of the statements below, and is treated in Appendix I. Writing $s_1 < \cdots < s_l$ for the indices selected by $\mathbf{B}_{\boldsymbol{\theta}} \in \mathcal{B}_l$, we summarise Shen et al. (2024a, Theorems 3 and 4) for our setting.

Lemma 1. *The first-order approximation, with respect to $\|\mathbf{B}_{\boldsymbol{\theta}}\mathbf{K}_{\boldsymbol{\theta}}\|_2$, of the dominant eigenvalue of $\mathbf{A}_{\boldsymbol{\theta}} + \mathbf{B}_{\boldsymbol{\theta}}\mathbf{K}_{\boldsymbol{\theta}}$ is $\tilde{\lambda}^{(1)}_{\mathsf{d},\boldsymbol{\theta}} = \lambda_{\mathsf{d},\boldsymbol{\theta}} + \sum_{i=1}^{l} g_i\,\kappa_{\boldsymbol{\theta}}(s_i)$, where $\kappa_{\boldsymbol{\theta}}(j)$ is the participation factor* (12) *of node j. If the input placement $\mathbf{B}_{\boldsymbol{\theta}}$ includes a key gene, then diagonal re-stabilization can be achieved.*

Lemma 1 makes the intervention target explicit: the leading-order eigenvalue shift is proportional to $\kappa_{\boldsymbol{\theta}}$ at the actuated node, so ranking nodes by $|\kappa_{\boldsymbol{\theta}}(j)|$ is exactly the selection rule to use, and it is the rule we evaluate in Sections 5 and 6. Its reliability from finite data is settled by the following theorem.

Theorem 3 (Finite-sample accuracy of the intervention target)**.** *Let $\sigma_\star > 0$ and $\delta_\lambda > 0$, and let $c_{\boldsymbol{\theta}}$ be the spectral conditioning constant of $\mathbf{A}_{\boldsymbol{\theta}}$ defined in Appendix I, which equals* 32 *when $\mathbf{A}_{\boldsymbol{\theta}}$ is normal. Under the sample-size condition of Theorem 2, and for ϵ small enough that the critical mode stays isolated (the explicit threshold is given in the same appendix), with probability at least $1 - \beta - \delta$,*

$$\max_j |\widehat{\kappa}_{\boldsymbol{\theta}}(j) - \kappa_{\boldsymbol{\theta}}(j)| \;\le\; \frac{c_{\boldsymbol{\theta}}\,\|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}}{\sigma_\star\,\delta_\lambda}\,\epsilon.$$

The proof is in Appendix I. The bound is governed by the two quantities that decide how sharply the critical mode is resolved, the identifiability margin $\sigma_\star$ and the spectral separation δ_λ, and both are computable from the identified model. This closes the pipeline: it bounds the error in the very quantity that ranks candidate sites, so the ordering of targets, and hence the intervention finally applied, is preserved even when $N < 2n$.

5 Numerical Examples on Synthetic mRNA-Protein Network

Simulation model. To validate our framework for system identification and early treatment design, we conduct numerical experiments on synthetic mRNA-protein regulatory networks that mimic disease progression through bifurcation.

We consider a synchronous modular bifurcation scenario: $R = 3$ structurally identical modules of 15 genes each give $n = 45$ genes and a $2n = 90$-dimensional state. In module r the master-regulator protein P_{15r} activates the transcription of all 15 genes of that module, including its own, so $\mathbf{A}_{\boldsymbol{\theta}}$ is block diagonal with one self-reinforcing loop per module. A single bifurcation parameter K_{master} — the half-activation threshold of the Hill term — is shared by all modules, so they approach their saddle-node bifurcations simultaneously. The ground truth is therefore the six-node set $\mathcal{M} = \{\mathrm{mRNA}_{15}, \mathrm{mRNA}_{30}, \mathrm{mRNA}_{45}, P_{15}, P_{30}, P_{45}\}$, whose participation factors (12) module symmetry makes exactly equal ($|\kappa_{\boldsymbol{\theta}}| = 0.437$ on the master proteins, 0.377 on the master mRNAs), fixing an unambiguous six-node target. Because the three modules become critical together, κ is here read off the projector onto the dominant eigenspace as described after (12), with the rest of the spectrum separated from it by $\delta_\lambda = 0.516$, so Theorem 3 applies in that form. The dynamics are driven by additive Gaussian noise with covariance $\mathbf{D} = \sigma^2\mathbf{I}_{90}$, $\sigma = 0.05$; snapshots record 2000 cells per episode over $H = 10$ episodes spanning $K_{\mathrm{master}} \in [K_{\mathrm{bif}} - 3,\ K_{\mathrm{bif}} - 0.4]$. This master-regulator pattern is of the same type as the identifiability witness of Proposition 1, so the simulated systems lie away from the exceptional set $\mathcal{E}$; full equations and parameters are in Appendix J.

Baselines and metric. On identical cell subsamples we compare the covariance-based identification of Shen et al. (2024b) (SOTA), the proposed method (Proposed), and scInTime (Xu et al., 2022), the leading single-cell method for this same task, which builds per-episode gene-regulatory networks along the trajectory and ranks genes by a trajectory-embedding statistic. Since a re-stabilising input is only useful when applied to a genuine regulator, we adopt the strict criterion that every one of

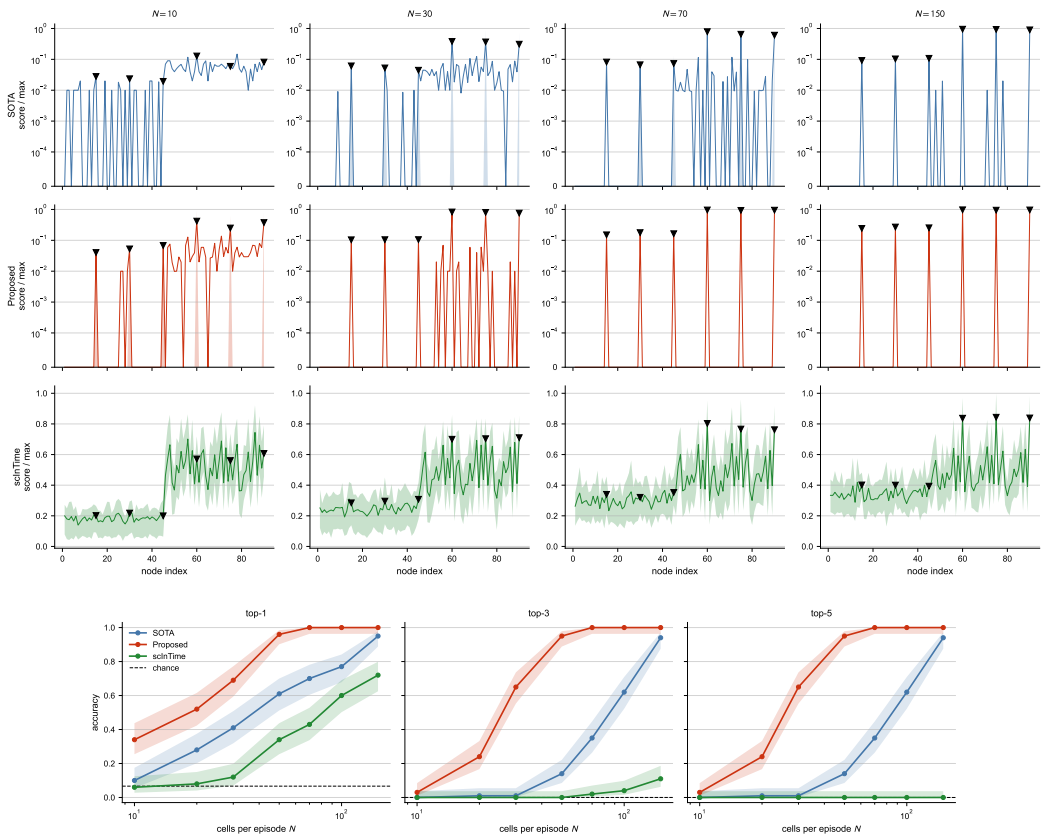

Figure 3: Synchronous modular network ($2n = 90$, six master regulators). **(a)** Per-node mean score (line) and interquartile range (band) over 100 trials, normalised by the per-trial maximum; triangles mark the masters. The proposed method places all of its mass on the masters and is numerically zero elsewhere (logarithmic axis), whereas scInTime (linear axis) leaves them inside a dense background. **(b)** Accuracy versus cells per episode N, counting a trial correct only if all k top-ranked nodes are masters; bands are Wilson 95% intervals and dashed lines are chance. The proposed method dominates both baselines at every N and is exact from $N = 70$, including under the strictest $k = 5$ criterion.

the k top-ranked nodes must lie in $\mathcal{M}$, for $k = 1, 3, 5$, with chance levels 6.7%, $1.7 \times 10^{-2}\%$ and $1.4 \times 10^{-5}\%$. We run 100 trials at each $N \in \{10, 20, 30, 50, 70, 100, 150\}$, covering the sample-starved regime $N < 2n = 90$ and $N > 2n$.

Results. Figure 3(a) shows where each method places its signal. The proposed estimator concentrates κ exactly on $\mathcal{M}$: at $N = 150$ its mean normalised score is 0.95–0.97 on the master proteins while every one of the 84 non-master nodes stays below 10^{-14}, whereas the scInTime statistic spreads across the network, its background reaching 0.65 against master values of 0.39–0.84. Figure 3(b) converts this into accuracy. At $N = 50$, barely half the cells needed to span the 90-dimensional state, the proposed method attains 96% top-1 against 61% for SOTA and 34% for scInTime, and under the strict top-5 criterion 95% versus 14% and 0%. From $N = 70$ it is exact (100% for $k = 1, 3, 5$), recovering the complete six-node master set in every trial, while SOTA reaches 95%/94% only at $N = 150$ and scInTime saturates at 72% top-1 and 0% top-5; paired McNemar tests confirm the ordering over scInTime at every N ($p \leq 1.6 \times 10^{-15}$ for $N \geq 30$) and over SOTA for all $N \leq 100$ ($p \leq 1.8 \times 10^{-4}$). The gain comes from pooling cells across episodes in the kernel covariance and confining the solution to the Lyapunov-admissible subspace of Section 4, which keeps the dominant eigenspace accurate precisely where $N < 2n$ renders per-episode covariance estimation rank-deficient. Appendix J.5 reports the complementary localized modular network, with an even larger margin at small N.

Table 1: Top-5 regulatory genes from the Paul et al. (2015) dataset by the participation factor $|\kappa_{\boldsymbol{\theta}}(i)|$ of (12).

Rank	Gene	κ_i	Lineage	Rank	Gene	κ_i	Lineage
1	*Mpo*	$+0.5081$	Granulocyte/Monocyte	2	*Klf1*	$+0.4154$	Erythroid
3	*Gata1*	$+0.3835$	Erythroid/Megakaryocyte	4	*Car2*	$+0.3446$	Erythroid
5	*Elane*	$+0.3041$	Granulocyte	-	-	-	-

6 Empirical Validation on Single-Cell Hematopoietic Data

We validate our method on Paul et al. (2015) single-cell RNA-seq dataset of mouse myeloid progenitor differentiation, with 2,730 cells across 19 cell-type clusters spanning erythroid, granulocyte, monocyte, and megakaryocyte lineages (visualized in Figure 7, Appendix K). Each cell type is treated as one episode k, with progressive transcriptional changes along the differentiation trajectory satisfying Assumption 1. As only mRNA-level measurements are available, we use the n-dimensional reduction of Appendix K.1. We constructed a 50-gene panel by combining curated hematopoietic TF genes with additional high-variance genes selected. Using this panel, we inferred a signed structural GRN prior $\widehat{\mathbf{A}}_{\mathrm{ref}}$ (illustrated in Figure 8) from pairwise Spearman correlations across all cells, retaining statistically supported edges ($|r_s| \geq 0.3$, $p < 0.05$; activation: $r_s > 0$, repression: $r_s < 0$). We then selected 15 cell-type episodes by minimizing the maximum pairwise SW_2 distance in this 50-dimensional space. Full pre-processing details are provided in Appendix K. The approximate identification problem ($\widehat{\mathcal{P}}_{\boldsymbol{\theta}}(\mathcal{D}_{\mathsf{snap}})$) was then solved with $N = 25$ cells per episode, yielding the participation factor $\kappa_{\boldsymbol{\theta}}(i)$ of (12) for each of the $n = 50$ genes; the dominant eigenvalue is simple here, so $\kappa_{\boldsymbol{\theta}}(i)$ reduces to the eigenvector product $\mathbf{w}_{\mathsf{d}}(i)\,\mathbf{v}_{\mathsf{d}}(i)$ up to normalisation. The κ signal is concentrated on a sparse connected subnetwork, while many nodes are near zero under the inferred structural prior. This shows that only a subset of genes carries strong first-order eigenvalue sensitivity for re-stabilization, consistent with the theoretical prediction of Appendix B that $\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}$ has support only on master-regulator nodes. Table 1 shows the top-5 genes by $|\kappa_i|$, all with positive sign, indicating that negative feedback gain is required for diagonal re-stabilisation per Lemma 1. All five genes are independently validated master regulators in hematopoietic differentiation (Iwasaki & Akashi, 2007; Tallack et al., 2010; Friedman, 2007; Horwitz et al., 2013); biological descriptions are in Appendix K.4. Notably, the top-5 span both the erythroid (*Klf1*, *Gata1*, *Car2*) and myeloid/granulocyte (*Mpo*, *Elane*) branches, the two major bifurcating lineages of Paul et al. (2015). This cross-lineage coverage is non-trivial: κ_i identifies genes whose perturbation most effectively shifts $\lambda_{\mathsf{d},\boldsymbol{\theta}}$ regardless of which downstream branch the system approaches, providing generalised re-stabilisation targets rather than lineage-specific markers. Given only unlabelled snapshot scRNA-seq data, our method recovers the same master regulators that decades of experimental hematopoiesis research have established as myeloid fate's primary controllers (Iwasaki & Akashi, 2007; Tallack et al., 2010; Ferreira et al., 2005).

7 Conclusions

We proposed a system identification framework for mRNA-protein regulatory networks from snapshot data, tailored to design effective intervention strategies. By exploiting the Lyapunov equation with structural constraints, our method achieves reliable estimation of the system matrix under snapshot data setting, with theoretical guarantees on finite-sample accuracy. Building on these results, we developed both optimal and approximate diagonal re-stabilization designs, offering actionable insights for early treatment at the pre-disease stage. We validate the proposed method on real single-cell datasets, showing its potential to impact broader applications in biology and beyond.

ETHICS STATEMENT

This work is computational and relies solely on publicly available, de-identified single-cell RNA-seq data (Paul et al., 2015) (NCBI GEO accession GSE72857); no new human- or animal-subject experiments were conducted. Our method outputs candidate regulatory genes for early intervention, and we emphasise that computationally identified nodes must not be treated as validated clinical targets without rigorous experimental follow-up. A fuller discussion of societal implications appears in Appendix M.

REPRODUCIBILITY STATEMENT

All modelling assumptions are stated explicitly in Assumptions 1 and 2, and complete proofs of Proposition 1 and Theorems 1, 2, and 3 are given in Appendices E–I. The identifiability condition (9) underlying Theorems 1–3 is checkable from data, and the witness of Proposition 1 is verified both analytically and numerically. The synthetic generative model, its full parameter table, and the noise settings are specified in Appendix J. For the real data, the dataset accession, gene-panel construction, GRN-prior thresholds, episode-selection procedure, and covariance-estimation hyper-parameters are documented in Appendix K and K.3, including the per-episode cell counts in Table 3. Source code for both the synthetic and real-data pipelines is included in the supplementary material as the self-contained `master_regulator_identification` package, whose `README` lists the exact commands that regenerate every figure and table reported here; it will be released publicly upon acceptance.

AI USE STATEMENT

In this work, we used generative AI tools as a general-purpose coding and writing assistant. Concretely, AI assistance was used to help implement and refactor the released experimental code—the synthetic snapshot generators, the benchmark harness, and the plotting utilities—and to draft and copy-edit portions of the manuscript text. The research problem, the modelling choices, the theoretical statements and their proofs, and the experimental design are the authors' own. Every AI-assisted derivation was re-checked by the authors, and every number, table and figure reported in the paper is produced directly by the released code rather than by a language model. Generative AI was not used to generate, modify or fabricate any experimental data or results, and was not used as the basis for our claims about prior work, which were verified against the cited sources. We take responsibility for the final content of this work, including text, claims and artifacts produced with the aid of generative AI.

REFERENCES

K. Aihara, R. Liu, and et al. Dynamical network biomarkers: Theory and applications. *Gene*, 808: 145997, 2022.

Corentin Briat, Ankit Gupta, and Mustafa Khammash. Antithetic integral feedback ensures robust perfect adaptation in noisy biomolecular networks. *Cell Systems*, 2(1):15–26, 2016.

Charlotte Bunne, Geoffrey Schiebinger, Andreas Krause, Aviv Regev, and Marco Cuturi. Optimal transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1):58, 2024.

Marek Capinski and Ekkehard Kopp. *Measure, Integral, and Probability*. Springer-Verlag, London, 2004.

L. Chen, R. Liu, Z. Liu, M. Li, and K. Aihara. Detecting early-warning signals for sudden deterioration of complex diseases by dynamical network biomarkers. *Scientific Reports*, 2:342, 2012.

Shen Chen, Daochuan Li, Dianke Yu, et al. Determination of tipping point in course of $PM_{2.5}$ organic extracts-induced malignant transformation by dynamic network biomarkers. *Journal of Hazardous Materials*, 426:128089, 2022a.

T. Chen, H. L. He, and G. M. Church. Modeling gene expression with differential equations. In *Proceedings of the Pacific Symposium on Biocomputing*, pp. 29–40. World Scientific, 1999.

Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. The most likely evolution of diffusing and vanishing particles: Schrödinger bridges with unbalanced marginals. *SIAM Journal on Control and Optimization*, 60(4):2016–2039, 2022b.

Ricardo Ferreira, Kinuko Ohneda, Masayuki Yamamoto, and Sjaak Philipsen. GATA1 function, a paradigm for transcription factors in hematopoiesis. *Molecular and Cellular Biology*, 25(4): 1215–1227, 2005.

Alexander G Fletcher and James M Osborne. Seven challenges in the multiscale modeling of multicellular tissues. *WIREs mechanisms of disease*, 14(1):e1527, 2022.

Alan D. Friedman. Transcriptional control of granulocyte and monocyte development. *Oncogene*, 26 (47):6816–6828, 2007.

Pilar Frontelo, Deepa Manwani, Mona Galdass, et al. Novel role for EKLF in megakaryocyte lineage commitment. *Blood*, 110(12):3871–3880, 2007.

J. G. D. Gooijer and D. Zerom. On conditional density estimation. *Statistica Neerlandica*, 57: 159–176, 2003.

Peter Hall, Rodney C. L. Wolff, and Qiwei Yao. Methods for estimating a conditional distribution function. *Journal of the American Statistical Association*, 94(445):109–116, 1999.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

Marshall S. Horwitz, Seth J. Corey, H. Leighton Grimes, and Timothy Tidwell. ELANE mutations in cyclic and severe congenital neutropenia: Genetics and pathophysiology. *Hematology/Oncology Clinics of North America*, 27(1):19–41, 2013. doi: 10.1016/j.hoc.2012.10.004.

Brian P Ingalls. *Mathematical modeling in systems biology: an introduction*. MIT press, 2013.

Hiromi Iwasaki and Koichi Akashi. Myeloid lineage commitment from the hematopoietic stem cell. *Immunity*, 26(6):726–740, 2007.

K. Koizumi et al. Suppression of dynamical network biomarker signals at the predisease state (mibyou) before metabolic syndrome in mice by a traditional japanese medicine (kampo formula) bofutsushosan. *Evidence-Based Complementary and Alternative Medicine*, 2020:9129134, 2020.

Soheil Kolouri, Se Rim Park, Matthew Thorpe, Dejan Slepcev, and Gustavo K. Rohde. Optimal mass transport: Signal processing and machine-learning applications. *IEEE Signal Processing Magazine*, 34(4):43–59, 2017.

S. G. Krantz and H. R. Parks. *A Primer of Real Analytic Functions*. Birkhäuser, Boston, 2nd edition, 2002.

J. Lee, I. Krivega, R.K. Dale, and A. Dean. The ldb1 complex co-opts ctcf for erythroid lineage-specific long-range enhancer interactions. *Cell Reports*, 19:2490–2502, 2017. doi: 10.1016/j.celrep.2017.05.072.

Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems - Series A*, 34(4):1533–1574, 2014.

Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos Theodorou. Deep generalized schrödinger bridge. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pp. 9374–9388, 2022.

L. Liu, T. Zhao, M. Ma, and Y. Wang. A new gene regulatory network model based on bp algorithm for interrogating differentially expressed genes of sea urchin. *SpringerPlus*, 5(1):1911, 2016.

R. Liu, P. Chen, K. Aihara, and L. Chen. Identifying early-warning signals of critical transitions with strong noise by dynamical network markers. *Scientific Reports*, 5:17501, 2015.

A. Meister, Y.H Li, B. Choi, and W.H. Wong. Learning a nonlinear dynamical system model of gene regulation: a perturbed steady-state approach. *The Annals of Applied Statistics*, 7(3):1311–1333, 2013.

B. Mityagin. The zero set of a real analytic function. *Mathematical Notes*, 107(3):529–530, 2020.

Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pp. 25858–25889. PMLR, 2023.

M. Oku and K. Aihara. On the covariance matrix of the stationary distribution of a noisy dynamical system. *Nonlinear Theory and Its Applications, IEICE*, 9(2):166–184, 2018.

Adriaan Passemiers, Yves Moreau, and Daniele Raimondi. Fast and accurate inference of gene regulatory networks through robust precision matrix estimation. *Bioinformatics*, 38(10):2802–2809, 2022.

Franziska Paul, Ya'ara Arkin, Amir Giladi, Diego Adhemar Jaitin, Ephraim Kenigsberg, Hadas Keren-Shaul, Deborah Winter, David Lara-Astiaso, Meital Gury, Assaf Weiner, et al. Transcriptional heterogeneity and lineage commitment in myeloid progenitors. *Cell*, 163(7):1663–1677, 2015.

Mehrshad Sadria and Thomas M Bury. Fatenet: an integration of dynamical systems and deep learning for cell fate prediction. *Bioinformatics*, 40(9):btae525, 2024.

Mehrshad Sadria and Anita Layton. scvaeder: integrating deep diffusion models and variational autoencoders for single-cell transcriptomics analysis. *Genome Biology*, 26(1):64, 2025. doi: 10.1186/s13059-025-03519-4.

Mehrshad Sadria and Anita T Layton. Interactions among mtorc, ampk and sirt: a computational model for cell energy balance and metabolism. *Cell Communication and Signaling*, 19(1):57, 2021a.

Mehrshad Sadria and Anita T Layton. Modeling within-host sars-cov-2 infection dynamics and potential treatments. *Viruses*, 13(6):1141, 2021b.

Mehrshad Sadria and Vasu Swaroop. Discovering governing equations of biological systems through representation learning and sparse model discovery. *NAR Genomics and Bioinformatics*, 7(2): lqaf048, 2025.

Mehrshad Sadria, Anita Layton, Sidharta Goyal, and Gary D Bader. Fatecode: Cell fate regulator prediction using classification autoencoder perturbation. *BioRxiv*, pp. 2022–12, 2022.

Mehrshad Sadria, Allen Zhang, and Gary D. Bader. Deep lineage: Single-cell lineage tracing and fate inference using deep learning. *bioRxiv*, 2024. doi: 10.1101/2024.04.25.591126.

Wouter Saelens, Robrecht Cannoodt, Helena Todorov, and Yvan Saeys. A comparison of single-cell trajectory inference methods. *Nature Biotechnology*, 37(5):547–554, 2019.

David W. Sanders, Nancy Kedersha, Daniel S.W. Lee, Amy R. Strom, Victoria Drake, Joshua A. Riback, Dan Bracha, Jorine M. Eeftens, Allana Iwanicki, Alicia Wang, Ming-Tzo Wei, Gena Whitney, Shawn M. Lyons, Paul Anderson, William M. Jacobs, Pavel Ivanov, and Clifford P. Brangwynne. Competing protein-rna interaction networks control multiphase intracellular organization. *Cell*, 181(2):306–324.e28, 2020.

Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.

Yutong Sha, Yuchi Qiu, Peijie Zhou, and Qing Nie. Reconstructing growth and dynamic trajectories from single-cell transcriptomics data. *Nature Machine Intelligence*, 6(1):25–39, 2024.

Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczynski. *Lectures on stochastic programming: modeling and theory*. SIAM, 2014.

X. Shen, H. Sasahara, J. Imura, M. Oku, and K. Aihara. Re-stabilizing large-scale network systems using high-dimension low-sample-size data analysis. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2024a.

X. Shen, N. Shimada, H. Sasahara, and J. Imura. Ultra-early medical treatment-oriented system identification using high-dimension low-sample-size data. *IFAC Journal of Systems and Control*, 27(100245), 2024b.

Jifan Shi, Kazuyuki Aihara, Tiejun Li, and Luonan Chen. Energy landscape decomposition for cell differentiation with proliferation effect. *National Science Review*, 9(8):nwac116, 2022.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2024.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021.

Michael R. Tallack, Thomas Whitington, Wai Shan Yuen, et al. A global role for KLF1 in erythropoiesis revealed by ChIP-seq in primary erythroid cells. *Genome Research*, 20(8):1052–1063, 2010.

Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 9526–9536. PMLR, 2020.

Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 1279–1287. PMLR, 2024.

C. A. Weidmann, A. M. Mustoe, P. B. Jariwala, J. M. Calabrese, and K. M. Weeks. Analysis of rna-protein networks with rnp-map defines functional hubs on rna. *Nature Biotechnology*, 39(3): 347–356, March 2021.

Qian Xu, Guanxun Li, Daniel Osorio, Yan Zhong, Yongjian Yang, Yu-Te Lin, Xiuren Zhang, and James J. Cai. scInTime: A computational method leveraging single-cell trajectory and gene regulatory networks to identify master regulators of cellular differentiation. *Genes*, 13(2):371, 2022. doi: 10.3390/genes13020371.

H. Yasukata, X. Shen, H. Sasahara, J. Imura, M. Oku, and K. Aihara. Design of input assignment and feedback gain for re-stabilizing undirected networks with high-dimension low-sample-size data. *International Journal of Robust and Nonlinear Control*, 33(12):6734–6753, 2023.

Y. Yu, T. Wang, and R. J. Samworth. A useful variant of the davis–kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.

Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through regularized unbalanced optimal transport. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025.

A Detailed background and related work

In biological systems, functions such as mRNA regulation and protein interactions occur through complex networks (Briat et al., 2016). Many diseases stem from abrupt deterioration in these networks, often modeled as bifurcation phenomena (Chen et al., 2012; Sadria & Bury, 2024). Using the snapshot data samples, the *Dynamical Network Biomarker* (DNB) method can predict the stage of the system immediately before the bifurcation occurs, referred to here as the pre-disease stage (Liu et al., 2015). According to DNB theory, certain nodes exhibit amplified fluctuations as they approach a bifurcation point. These dominant directions reflect nodes with large fluctuations, and can be estimated from snapshot data. The DNB theory also enables early disease detection based on the increased fluctuations of specific biomarkers (Aihara et al., 2022). The ability to detect diseases at the pre-disease stage is crucial for early medical intervention. Traditional Japanese medicine has been used to suppress DNB node fluctuations and prevent disease progression (Koizumi et al., 2020). Experimental results in (Chen et al., 2022a) demonstrate that manipulating multiple DNB nodes can significantly alter malignant phenotypes in lung cancer. The success of heuristic approaches has further inspired the development of theoretical frameworks for early medical intervention. Recent studies have explored early treatment via high-dimensional low sample-size (HDLSS) snapshot data. While Yasukata et al. (2023) proposed a single-input method for undirected networks, extensions to directed networks were developed in (Shen et al., 2024a). These works highlight the importance of the system matrix's left eigenvector corresponding to eigenvalues with maximal real parts for optimal input placement, information not directly accessible via principal component analysis (PCA), necessitating system identification by snapshot data. Learning stochastic dynamics from snapshot data has recently gained considerable attention in the machine learning community (Song et al., 2021; Neklyudov et al., 2023; Tong et al., 2024). A dominant approach involves first inferring time-series trajectories from the data, followed by system identification based on the inferred trajectories (Tong et al., 2020). This trajectory inference has been particularly studied in the context of single-cell RNA sequencing (Saelens et al., 2019; Shi et al., 2022; Sha et al., 2024), where uncovering the underlying temporal progression of cellular states is essential. Optimal transport methods have become central tools for such snapshot datasets with temporal resolution (Bunne et al., 2024; Schiebinger et al., 2019), with the Schrödinger bridge (SB) formulation extending these methods to stochastic dynamics by modeling the most likely stochastic paths between two distributions relative to a reference process (Léonard, 2014; Shi et al., 2024; Liu et al., 2022). To improve computational tractability, regularization techniques have also been introduced to these transport-based formulations (Chen et al., 2022b; Zhang et al., 2025), alongside generative models designed to capture the pseudodynamics of the process using different modalities, such as Deep Lineage and scVAEDer (Sadria et al., 2024; Sadria & Layton, 2025). Despite these advances, such methods often suffer from high computational cost and rely on large sample sizes to obtain reliable results, which are rarely satisfied in single-cell data analysis.

A core challenge in modelling gene regulatory systems is their inherently complex multiscale character: slow transcriptional changes interact with much faster protein-level dynamics, and meaningful mechanistic descriptions and models must account for different scales (Fletcher & Osborne, 2022). Classical mechanistic models based on differential equations have been used in systems biology for capturing these interactions across domains from circadian–metabolic coupling and metabolic regulation to mechanistic models of immune and viral dynamics (Sadria & Layton, 2021a;b; Ingalls, 2013). While these models have enabled mechanistic insight across many biological settings, they typically require dense time-series or carefully designed experiments for reliable parameter identification, which are scarce in current high-throughput single-cell genomics. These limitations have motivated powerful trajectory-inference and representation-learning methods that reconstruct temporal progressions from snapshots (for example, transport-based formulations and Schrödinger-bridge approaches). Such tools are useful for recovering likely cellular paths and population-level flows, but their design goals differ from those needed for control: they prioritize path reconstruction (often under specific loss/regularization choices) rather than explicit recovery of the low-level regulatory parameters and multiscale structure required to reason about stability or to design energy-efficient interventions. Complementary recent work demonstrates scalable deep and representation-learning approaches that predict fate changes or extract parsimonious dynamical models from single-cell data and a library-guided sparse discovery framework (Sadria et al., 2022; Sadria & Swaroop, 2025); these advances substantially improve fate prediction and model discovery, but by themselves do not directly yield the identifiable, control-ready parameterizations we target here.

B EARLY DETECTION OF PRE-DISEASE STAGE

Let $\mathbf{C}_{\boldsymbol{\theta}} \in \mathbb{R}^{2n\times 2n}$ denote the covariance matrix of $\mathbf{x}$ when the system is parameterized by $\boldsymbol{\theta}$. As presented in (Oku & Aihara, 2018, Section 3.2), as the parameter $\boldsymbol{\theta}$ approaches the tipping point $\boldsymbol{\theta}_{\mathsf{C}}$, the covariance matrix $\mathbf{C}_{\boldsymbol{\theta}}$ converges to a rank-one matrix if $\lambda_{\mathsf{d},\boldsymbol{\theta}}$ is real, written by

$$\lim_{\lambda_{\mathsf{d},\boldsymbol{\theta}}\to 0^-}(-2\lambda_{\mathsf{d},\boldsymbol{\theta}})\mathbf{C}_{\boldsymbol{\theta}} = \widetilde{D}_{11}\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}^{\top}, \tag{13}$$

where the positive real number $\widetilde{D}_{11}$ is the first diagonal element of $\widetilde{\mathbf{D}} = \mathbf{V}_{\boldsymbol{\theta}}^{-1}\mathbf{D}(\mathbf{V}_{\boldsymbol{\theta}}^{\mathsf{H}})^{-1}$. Here, $\mathbf{V}_{\boldsymbol{\theta}}$ is the matrix whose columns are the right eigenvectors of $\mathbf{A}_{\boldsymbol{\theta}}$, and $\mathbf{V}_{\boldsymbol{\theta}}^{\mathsf{H}}$ denotes its conjugate transpose. On the other hand, the covariance matrix $\mathbf{C}_{\boldsymbol{\theta}}$ converges (after the same rescaling) to a rank-two matrix if $\lambda_{\mathsf{d},\boldsymbol{\theta}}$ is complex. In both cases the dominant eigenvalue of $\mathbf{C}_{\boldsymbol{\theta}}$ diverges as $\lambda_{\mathsf{d},\boldsymbol{\theta}} \to 0^-$.

A group of diagonal elements of $\mathbf{C}_{\boldsymbol{\theta}}$ also increases to infinity, indicating that there exists a group of nodes whose standard deviations become unbounded. In parallel, the covariance between two nodes in this group also grows substantially as $\boldsymbol{\theta}$ approaches the tipping point $\boldsymbol{\theta}_{\mathsf{C}}$. We take the case where $\lambda_{\mathsf{d},\boldsymbol{\theta}}$ is real as an example to illustrate this. The eigenvector $\boldsymbol{\xi}_{\mathsf{d},\boldsymbol{\theta}} = [\boldsymbol{\xi}_{\mathsf{d},\boldsymbol{\theta},\mathsf{r}}^{\top}, \boldsymbol{\xi}_{\mathsf{d},\boldsymbol{\theta},\mathsf{p}}^{\top}]^{\top}$ of $\mathbf{C}_{\boldsymbol{\theta}}$ corresponding to the dominant eigenvalue converges to the right eigenvector of $\mathbf{A}_{\boldsymbol{\theta}}$ associated with the dominant eigenvalue $\lambda_{\mathsf{d},\boldsymbol{\theta}}$, namely,

$$\lim_{\lambda_{\mathsf{d},\boldsymbol{\theta}}\to 0^-}\boldsymbol{\xi}_{\mathsf{d},\boldsymbol{\theta}} = \mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}. \tag{14}$$

By (13), if $\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}(i) \neq 0$ and $\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}(j) \neq 0$, we have

$$\lim_{\lambda_{\mathsf{d},\boldsymbol{\theta}}\to 0^-}|\mathbf{C}_{\boldsymbol{\theta}}(i,j)| = \infty, \tag{15}$$

which is the theoretical explanation for the system's large fluctuations in the pre-disease stage. By (13), the variance of node i, i.e. the diagonal element $\mathbf{C}_{\boldsymbol{\theta}}(i,i)$, grows proportionally to $\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}^2(i)$ near the tipping point. Consequently, the standard deviation of the node, where $\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}(i) \neq 0$, shows a significant increase at the pre-disease stage. The covariance between two nodes ($\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}(i) \neq 0,\ \mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}(j) \neq 0$) also increases significantly in the pre-disease stage.

The above two properties provide a theoretical explanation for the system's large fluctuations in the pre-disease stage, when $\boldsymbol{\theta} \approx \boldsymbol{\theta}_{\mathsf{C}}$. By examining system fluctuations from snapshot data, it becomes possible to detect the pre-disease stage, which is completely model-free. We refer to Chen et al. (2012); Aihara et al. (2022) for more details about indicator selection and computation for pre-disease detection by snapshot data.

C RANDOM VARIABLE IN AN AUGMENTED SPACE

This appendix fixes the probabilistic setting behind the kernel conditional density estimation of Section 4 and Appendix F: the state $\mathbf{X} \in \mathbb{R}^{2n}$ and the sampling time T are treated as one random variable on an augmented space. Consider a random variable $\boldsymbol{\xi} := (\mathbf{X}, T)$ with support $\Xi \subset \mathbb{R}^{2n+1}$, where $\xi := (\mathbf{x}, t)$[2] is a realization of $\boldsymbol{\xi}$. Let $\mathcal{F}$ denote the σ-algebra of subsets of Ξ. Equipped with a probability measure ρ defined on the Borel space $(\Xi, \mathcal{F})$, this forms a probability space $(\Xi, \mathcal{F}, \rho)$. Let $\Xi_{\mathsf{s}} \subseteq \Xi$ be a subset of Ξ. Given a continuous probability density function $p(\xi)$ with support Ξ, the probability that the random variable $\boldsymbol{\xi}$ lies within Ξ_{s} is expressed as $\mathsf{Pr}\{\boldsymbol{\xi} \in \Xi_{\mathsf{s}}\} := \int_{\Xi_{\mathsf{s}}} p(\xi)\,\mathsf{d}\xi$. Suppose that the probability density function $p(\xi)$ is a joint density denoted by $p(\mathbf{x}, t)$. Both $\mathbf{x}$ and t are continuous random variables with marginal probability densities[3] given by $p_{\mathbf{X}}(\mathbf{x}) := \int_{\mathbb{R}^+} p(\mathbf{x}, t)\,\mathsf{d}t,\ p_T(t) := \int_{\mathbb{R}^{2n}} p(\mathbf{x}, t)\,\mathsf{d}\mathbf{x}$. Consider that $p_T(t) > 0$ holds for any $t \in \mathbb{R}^+$ and $p(\mathbf{x}, t) > 0$ holds for any $(\mathbf{x}, t) \in \mathbb{R}^{2n+1}$. A conditional probability density is defined by $p_{\mathbf{X}}^{\mathsf{c}}(\mathbf{x}|T = t) := p(\mathbf{x}, t)/p_T(t)$. We omit "$T =$" in the conditional probability density in the remainder of the paper. Additionally, we also define the conditional probability of having $\mathbf{X} \in \mathcal{X}_{\mathsf{s}} \subseteq \mathbb{R}^{2n}$ when $T = t$ by $\mathsf{Pr}\{\mathbf{X} \in \mathcal{X}_{\mathsf{s}}|t\} := \int_{\mathcal{X}_{\mathsf{s}}} p_{\mathbf{X}}^{\mathsf{c}}(\mathbf{x}|t)\mathsf{d}\mathbf{x}$.

[2] $\xi_k := \left(\mathbf{x}_m^{(k)}, t_k\right)$ can represent a data point from the snapshot dataset.

[3] Although the time point or sample time index may appear as a discrete integer, it is treated in a general sense as a continuous variable. Moreover, a continuous probability density is used to approximate discrete probability over a continuous domain (Capinski & Kopp, 2004).

D BIOLOGICAL INTERPRETATION OF RE-STABILIZATION

Interpretation of re-stabilization. By solving Problem 1 we increase the stability margin of the mRNA-protein regulatory system, making it less sensitive to random fluctuations. The state distribution becomes more concentrated around the healthy equilibrium and less likely to drift toward the disease state under uncertain perturbations. In biological terms, even when the system is close to a critical transition, effective re-stabilization can suppress large fluctuations in regulatory genes and reduce the probability of crossing into the disease stage. Identifying regulatory nodes therefore not only enables early detection but also provides actionable intervention targets.

E PROOF OF PROPOSITION 1

Throughout, $\mathcal{S}$ is the subspace (7) of dimension $p = n^2 + 3n$, $\mathsf{Sym}(2n)$ has dimension $n(2n+1)$, and $\mathbf{D} = \sigma^2\mathbf{I}_{2n}$, $\sigma > 0$, is fixed. Moreover, $\mathcal{L}_{\mathbf{C}} : \mathcal{S} \to \mathsf{Sym}(2n)$ is $\Delta\mathbf{A} \mapsto \Delta\mathbf{A}\mathbf{C} + \mathbf{C}\Delta\mathbf{A}^\top$, with matrix $\mathbf{M}_{\mathcal{S}}(\mathbf{C})$ in any fixed pair of bases. Note $p < n(2n+1)$ exactly when $n > 2$.

Working with $\mathbf{M}_{\mathcal{S}}$ instead of the extended system (6) is a change of coordinates only: the two describe the same solution set, since (5) is precisely the statement $\mathbf{A}_{\boldsymbol{\theta}} \in \mathcal{S}$, but the restricted operator exposes the column rank that governs uniqueness.

E.1 STEP 1: REDUCTION OF GENERICITY TO A SINGLE WITNESS

Let $\mathcal{S}_{\mathsf{stab}} = \{\mathbf{A} \in \mathcal{S} : \mathbf{A} \text{ Hurwitz}\}$, an open subset of $\mathcal{S} \cong \mathbb{R}^p$, and let $\mathcal{O}$ be one of its connected components. For $\mathbf{A} \in \mathcal{S}_{\mathsf{stab}}$, vectorising (3) gives

$$(\mathbf{A} \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{A})\,\mathsf{vec}(\mathbf{C}) = -\,\mathsf{vec}(\mathbf{D})\,, \tag{16}$$

and the eigenvalues of $\mathbf{A} \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{A}$ are the pairwise sums $\lambda_i + \lambda_j$ of eigenvalues of $\mathbf{A}$. Since $\mathbf{A}$ is Hurwitz these are all nonzero, so (16) has the unique solution $\mathbf{C}(\mathbf{A}) = \int_0^\infty e^{\mathbf{A}t}\mathbf{D}e^{\mathbf{A}^\top t}\mathrm{d}t \succ 0$. By Cramer's rule each entry of $\mathbf{C}(\mathbf{A})$ is a ratio of polynomials in the entries of $\mathbf{A}$ whose denominator $\det(\mathbf{A} \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{A})$ does not vanish on $\mathcal{S}_{\mathsf{stab}}$. Hence $\mathbf{A} \mapsto \mathbf{C}(\mathbf{A})$ is real-analytic on $\mathcal{S}_{\mathsf{stab}}$. As $\mathbf{C} \mapsto \mathbf{M}_{\mathcal{S}}(\mathbf{C})$ is linear, the Gram determinant

$$\Phi(\mathbf{A}) := \det\big(\mathbf{M}_{\mathcal{S}}(\mathbf{C}(\mathbf{A}))^\top \mathbf{M}_{\mathcal{S}}(\mathbf{C}(\mathbf{A}))\big) \tag{17}$$

is real-analytic on $\mathcal{O}$, and $\Phi(\mathbf{A}) \neq 0$ if and only if $\mathbf{M}_{\mathcal{S}}(\mathbf{C}(\mathbf{A}))$ has full column rank p, i.e. iff $\mathcal{L}_{\mathbf{C}(\mathbf{A})}$ is injective on $\mathcal{S}$. Therefore $\mathcal{E} = \{\mathbf{A} \in \mathcal{O} : \Phi(\mathbf{A}) = 0\}$ is closed in $\mathcal{O}$.

A real-analytic function on a connected open subset of $\mathbb{R}^p$ that is not identically zero has a zero set of Lebesgue measure zero (Mityagin, 2020; Krantz & Parks, 2002). Consequently, if Φ is nonzero at a single $\mathbf{A}^\star \in \mathcal{O}$, then $\Phi \not\equiv 0$ on $\mathcal{O}$ and $\mathsf{Leb}(\mathcal{E}) = 0$. This reduces the proposition to exhibiting one witness.

E.2 STEP 2: AN EXPLICIT WITNESS

The Lyapunov solution depends linearly on the noise covariance. In particular, if $\mathbf{C}_2(\mathbf{A})$ denotes the solution for $\mathbf{D} = 2\mathbf{I}_{2n}$, then $\mathbf{C}(\mathbf{A}) = (\sigma^2/2)\mathbf{C}_2(\mathbf{A})$ and $\mathbf{M}_{\mathcal{S}}(\mathbf{C}(\mathbf{A})) = (\sigma^2/2)\mathbf{M}_{\mathcal{S}}(\mathbf{C}_2(\mathbf{A}))$. This positive scaling does not change column rank, so it suffices to construct a witness under the normalisation $\mathbf{D} = 2\mathbf{I}_{2n}$. Let $n \geq 3$, $\mathbf{1} = [1, \ldots, 1]^\top \in \mathbb{R}^n$, $\mathbf{J} = \mathbf{1}\mathbf{1}^\top$, $\boldsymbol{\Gamma} = \mathbf{1}\mathbf{e}_1^\top$, and set

$$\mathbf{A}^\star = \begin{bmatrix} -\mathbf{I} & \boldsymbol{\Gamma} \\ 0 & -\mathbf{I} \end{bmatrix}, \qquad \mathbf{C}^\star = \begin{bmatrix} \mathbf{I} + \frac{1}{2}\mathbf{J} & \frac{1}{2}\boldsymbol{\Gamma} \\ \frac{1}{2}\boldsymbol{\Gamma}^\top & \mathbf{I} \end{bmatrix}. \tag{18}$$

Then $\mathbf{A}^\star \in \mathcal{S}$ (with $\Delta_{\mathsf{r}} = \Delta_{\mathsf{p}} = \mathbf{I}$, $\Gamma_{\mathsf{r}} = 0$, $\Gamma_{\mathsf{p}} = \boldsymbol{\Gamma}$) and $\mathbf{A}^\star$ is Hurwitz, its spectrum being $\{-1\}$; hence $\mathbf{A}^\star \in \mathcal{S}_{\mathsf{stab}}$.

(a) $\mathbf{C}^\star = \mathbf{C}(\mathbf{A}^\star)$. Using $\boldsymbol{\Gamma}\boldsymbol{\Gamma}^\top = \mathbf{1}\mathbf{e}_1^\top\mathbf{e}_1\mathbf{1}^\top = \mathbf{J}$,

$$\mathbf{A}^\star\mathbf{C}^\star = \begin{bmatrix} -\big(\mathbf{I} + \frac{1}{2}\mathbf{J}\big) + \frac{1}{2}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^\top & -\frac{1}{2}\boldsymbol{\Gamma} + \boldsymbol{\Gamma} \\ -\frac{1}{2}\boldsymbol{\Gamma}^\top & -\mathbf{I} \end{bmatrix} = \begin{bmatrix} -\mathbf{I} & \frac{1}{2}\boldsymbol{\Gamma} \\ -\frac{1}{2}\boldsymbol{\Gamma}^\top & -\mathbf{I} \end{bmatrix},$$

whose symmetric part is $-\mathbf{I}_{2n}$. Hence $\mathbf{A}^\star\mathbf{C}^\star + \mathbf{C}^\star\mathbf{A}^{\star\top} + 2\mathbf{I} = 0$ as required. Moreover $\mathbf{C}^\star \succ 0$: its Schur complement with respect to the $(2,2)$ block is $\mathbf{I} + \frac{1}{2}\mathbf{J} - \frac{1}{4}\boldsymbol{\Gamma}\boldsymbol{\Gamma}^\top = \mathbf{I} + \frac{1}{4}\mathbf{J} \succ 0$.

(b) $\mathcal{L}_{\mathbf{C}^\star}$ is injective on $\mathcal{S}$. Write $\mathbf{P} = \frac{1}{2}\mathbf{\Gamma}$, $\mathbf{Q} = \mathbf{I} + \frac{1}{2}\mathbf{J}$, and take $\Delta\mathbf{A} = \left[\begin{smallmatrix} \mathbf{D}_\mathsf{r} & \mathbf{B} \\ \mathbf{G} & \mathbf{D}_\mathsf{p} \end{smallmatrix}\right] \in \mathcal{S}$ with $\mathbf{D}_\mathsf{r} = \mathsf{diag}(r_i)$, $\mathbf{D}_\mathsf{p} = \mathsf{diag}(s_i)$, $\mathbf{G} = \mathsf{diag}(g_i)$. Since $\mathbf{C}^\star$ is symmetric, the condition $\mathcal{L}_{\mathbf{C}^\star}(\Delta\mathbf{A}) = 0$ says that

$$\Delta\mathbf{A}\,\mathbf{C}^\star = \begin{bmatrix} \mathbf{D}_\mathsf{r}\mathbf{Q} + \mathbf{B}\mathbf{P}^\top & \mathbf{D}_\mathsf{r}\mathbf{P} + \mathbf{B} \\ \mathbf{G}\mathbf{Q} + \mathbf{D}_\mathsf{p}\mathbf{P}^\top & \mathbf{G}\mathbf{P} + \mathbf{D}_\mathsf{p} \end{bmatrix} \quad \text{is skew-symmetric.}$$

Lower-right block. $\mathbf{G}\mathbf{P} + \mathbf{P}^\top\mathbf{G} + 2\mathbf{D}_\mathsf{p} = 0$. Here $(\mathbf{G}\mathbf{P})_{ij} = \frac{1}{2}g_i\delta_{j1}$ and $(\mathbf{P}^\top\mathbf{G})_{ij} = \frac{1}{2}g_j\delta_{i1}$. Reading the entries with $i \neq j$ and $j = 1$ gives $g_i = 0$ for all $i \geq 2$; the diagonal entries then give $s_i = 0$ for $i \geq 2$ and $g_1 + 2s_1 = 0$. Thus $\mathbf{G} = g_1\mathbf{e}_1\mathbf{e}_1^\top$ and $\mathbf{D}_\mathsf{p} = -\frac{g_1}{2}\mathbf{e}_1\mathbf{e}_1^\top$.

Off-diagonal blocks. $\mathbf{D}_\mathsf{r}\mathbf{P} + \mathbf{B} + \left(\mathbf{G}\mathbf{Q} + \mathbf{D}_\mathsf{p}\mathbf{P}^\top\right)^\top = 0$, i.e. $\mathbf{B} = -\mathbf{D}_\mathsf{r}\mathbf{P} - \mathbf{Q}\mathbf{G} - \mathbf{P}\mathbf{D}_\mathsf{p}$. Every term has all columns but the first equal to zero, so $\mathbf{B} = \mathbf{b}\mathbf{e}_1^\top$ with, writing $\mathbf{r} = [r_1, \ldots, r_n]^\top$ and using $\mathbf{Q}\mathbf{e}_1 = \mathbf{e}_1 + \frac{1}{2}\mathbf{1}$,

$$\mathbf{b} = -\tfrac{1}{2}\mathbf{r} - g_1\mathbf{e}_1 - \tfrac{g_1}{4}\mathbf{1}. \tag{19}$$

Upper-left block. $\mathbf{D}_\mathsf{r}\mathbf{Q} + \mathbf{Q}\mathbf{D}_\mathsf{r} + \mathbf{B}\mathbf{P}^\top + \mathbf{P}\mathbf{B}^\top = 0$. Since $\mathbf{D}_\mathsf{r}\mathbf{Q} + \mathbf{Q}\mathbf{D}_\mathsf{r} = 2\mathbf{D}_\mathsf{r} + \frac{1}{2}(\mathbf{D}_\mathsf{r}\mathbf{J} + \mathbf{J}\mathbf{D}_\mathsf{r})$, $\mathbf{B}\mathbf{P}^\top = \frac{1}{2}\mathbf{b}\mathbf{1}^\top$ and $\mathbf{P}\mathbf{B}^\top = \frac{1}{2}\mathbf{1}\mathbf{b}^\top$, the (i,j) entry reads $2r_i\delta_{ij} + \frac{1}{2}(u_i + u_j) = 0$ with $u_i := r_i + b_i$. For $i \neq j$ this gives $u_i + u_j = 0$ for every pair, which for $n \geq 3$ forces $u \equiv 0$; the diagonal then gives $2r_i + u_i = 0$, hence $\mathbf{D}_\mathsf{r} = 0$ and $\mathbf{b} = -\mathbf{r} = 0$, so $\mathbf{B} = 0$. Substituting $\mathbf{b} = 0$, $\mathbf{r} = 0$ into (19) leaves $g_1\mathbf{e}_1 + \frac{g_1}{4}\mathbf{1} = 0$, whose $i \geq 2$ component gives $g_1 = 0$. Hence $\mathbf{G} = 0$, $\mathbf{D}_\mathsf{p} = 0$ and $\Delta\mathbf{A} = 0$, proving injectivity. Combining with Step 1 completes the proof. □

Scope of genericity and link to finite-sample recovery. We emphasize the parameter space over which the genericity statement is made. For a data-generating structured stable Jacobian $\mathbf{A}_0 \in \mathcal{O}$, let $\mathbf{C}_0 := \mathbf{C}(\mathbf{A}_0)$ denote its population covariance determined by the Lyapunov equation. The estimator does not observe $\mathbf{C}_0$ directly, but instead constructs a data-dependent estimate $\widehat{\mathbf{C}}$. The model-generated covariance family $\mathbf{C}(\mathcal{O})$ is the image of the p-dimensional parameter space $\mathcal{O} \subseteq \mathcal{S} \cong \mathbb{R}^p$ and has Hausdorff dimension at most p. Since the positive-definite cone in $\mathsf{Sym}(2n)$ has dimension $n(2n+1) > p$, genericity over all $\mathbf{C} \succ 0$ would not characterize the covariances generated by the model. Proposition 1 is therefore formulated with respect to Lebesgue measure on the structured parameter space $\mathcal{O}$, which is also the parameter space used in $\widehat{\mathcal{P}}_{\boldsymbol{\theta}}(\mathcal{D}_{\mathsf{snap}})$. The proposition establishes population-level injectivity at $\mathbf{C}_0$; finite-sample recovery additionally requires control of $\|\widehat{\mathbf{C}} - \mathbf{C}_0\|_\mathsf{F}$. This covariance-estimation error is bounded in Theorem 5 of Appendix H and is then propagated to $\|\widehat{\mathbf{A}} - \mathbf{A}_0\|_\mathsf{F}$ in the proof of Theorem 2.

F KERNEL CONDITIONAL DENSITY ESTIMATION: EXPLICIT FORMULAE

This appendix collects the details of the Nadaraya–Watson (NW) construction summarised in Section 4. The underlying probability space of the augmented random variable $(\mathbf{X}, T)$, together with the joint density $p(\mathbf{x}, t)$ and the conditional density $p^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t)$ that the estimator below targets, is formalised in Appendix C. The first step estimates the joint density $p(\mathbf{x}, t)$ from $\mathcal{D}_{\mathsf{snap}}$ by kernel density estimation,

$$\widehat{p}\,(\mathbf{x}, t) = \frac{1}{NHh^{2n+1}} \sum_{k,m} K_{\mathbf{X}}\left(\frac{\mathbf{x} - \mathbf{x}_m^{(k)}}{h}\right) K_T\left(\frac{t - t_k}{h}\right), \tag{20}$$

where $K_{\mathbf{X}}(\cdot)$ and $K_T(\cdot)$ are kernels for $\mathbf{X}$ and t and h is the bandwidth, required to satisfy the standard consistency conditions $h \to 0$ and $NHh^{2n} \to \infty$ (Gooijer & Zerom, 2003). Many kernels are admissible in practice, including uniform, triangular, biweight, triweight, Epanechnikov and Gaussian. The NW conditional density estimator is then

$$\widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t) = \left\{ \sum_{k,m} K_{\mathbf{X}}\left(\frac{\mathbf{x} - \mathbf{x}_m^{(k)}}{h}\right) K_T\left(\frac{t - t_k}{h}\right) \right\} \Big/ \left\{ \sum_{k,m} K_T\left(\frac{t - t_k}{h}\right) \right\}. \tag{21}$$

Each sample $\mathbf{x}_m^{(k)}$ is associated with its episode time t_k, so there are effectively N samples of $\mathbf{x}$ attached to each t_k even though t_k itself occurs once per episode. This is precisely why the time kernel K_T lets neighbouring episodes contribute to the estimate at t_k: the estimator is driven by the pooled count $M = NH$, which is what makes the HDLSS regime $N < 2n$ tractable.

G Proof of Theorem 1

As preparation we first establish consistency of the covariance estimate.

Theorem 4. *Let $\widehat{\mathbf{C}}_{\boldsymbol{\theta}}$ be the sample covariance of M_{s} i.i.d. draws from the Nadaraya–Watson estimator $\widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\cdot \mid t_k)$ built from $\mathcal{D}_{\mathsf{snap}}$, and suppose the bandwidth satisfies $h \to 0$ and $NHh^{2n} \to \infty$. Then $\widehat{\mathbf{C}}_{\boldsymbol{\theta}} \xrightarrow{p} \mathbf{C}_{\boldsymbol{\theta}}$ as $M_{\mathsf{s}} \to \infty$ and $M = NH \to \infty$.*

Proof. The proof consists of two steps.

Step 1: Convergence of $\widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k)$ to $p^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k)$, and convergence of corresponding covariance.

By the results in (Gooijer & Zerom, 2003; Hall et al., 1999), under the stated bandwidth conditions $h \to 0$ and $NHh^{2n} \to \infty$, the NW conditional density estimator satisfies

$$\sup_{\mathbf{x} \in \mathbb{R}^{2n}} |\widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k) - p^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k)| \xrightarrow{p} 0 \quad \text{as} \quad N, H \to \infty.$$

Since the true density $p^{\mathsf{c}}_{\mathbf{X}}$ is Gaussian $\mathcal{N}(\mathbf{0}^{2n}, \mathbf{C}_{\boldsymbol{\theta}})$ with finite second moments, and the uniform convergence holds, we can invoke standard results on the convergence of integrals of bounded functions with respect to densities (e.g., Theorem 2.1 in (Gooijer & Zerom, 2003)).

Let $g(\mathbf{x}) = \mathbf{x}\mathbf{x}^\top$, which is polynomially bounded. Then, as $N, H \to \infty$,

$$\int g(\mathbf{x}) \widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k) d\mathbf{x} \xrightarrow{p} \int g(\mathbf{x}) p^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k) d\mathbf{x} = \mathbf{C}_{\boldsymbol{\theta}}. \tag{22}$$

Thus, the covariance matrix of $\widehat{p}^{\mathsf{c}}_{\mathbf{X}}$ converges in probability to $\mathbf{C}_{\boldsymbol{\theta}}$.

Step 2: Convergence of empirical covariance $\widehat{\mathbf{C}}_{\boldsymbol{\theta}}$ to the covariance of $\widehat{p}^{\mathsf{c}}_{\mathbf{X}}$.

Now consider $\{\widehat{\mathbf{x}}_m^{(k)}\}_{m=1}^{M_{\mathsf{s}}}$ as i.i.d. samples from $\widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k)$. By the Weak Law of Large Numbers, for any $\epsilon_1 > 0$ and $\delta_1 > 0$, there exists M_0 such that for all $M_{\mathsf{s}} \geq M_0$, we have

$$\mathsf{Pr}\left\{ \left\| \widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \int \mathbf{x}\mathbf{x}^\top \widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k) d\mathbf{x} \right\|_{\mathsf{F}} < \epsilon_1 \right\} \geq 1 - \delta_1. \tag{23}$$

Combination of the two steps: ϵ-δ argument.

Finally, we combine the two steps carefully. Let $\epsilon > 0$ and $\delta > 0$ be arbitrary.

- From Step 1, there exists (N_0, H_0) such that for all $N \geq N_0, H \geq H_0$,
$$\mathsf{Pr}\left\{ \left\| \int \mathbf{x}\mathbf{x}^\top \widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k) d\mathbf{x} - \mathbf{C}_{\boldsymbol{\theta}} \right\|_{\mathsf{F}} < \epsilon/2 \right\} \geq 1 - \delta/2.$$
- From Step 2, for these fixed (N, H) and for sufficiently large $M_{\mathsf{s}} \geq M_0$, we have
$$\mathsf{Pr}\left\{ \left\| \widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \int \mathbf{x}\mathbf{x}^\top \widehat{p}^{\mathsf{c}}_{\mathbf{X}}(\mathbf{x} \mid t_k) d\mathbf{x} \right\|_{\mathsf{F}} < \epsilon/2 \right\} \geq 1 - \delta/2.$$

Now applying the union bound, for $N \geq N_0$, $H \geq H_0$ and $M_{\mathsf{s}} \geq M_0$, we obtain:

$$\mathsf{Pr}\left\{ \left\| \widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \mathbf{C}_{\boldsymbol{\theta}} \right\|_{\mathsf{F}} < \epsilon \right\} \geq 1 - \delta.$$

Since ϵ and δ are arbitrary, we conclude that:

$$\widehat{\mathbf{C}}_{\boldsymbol{\theta}} \xrightarrow{p} \mathbf{C}_{\boldsymbol{\theta}} \quad \text{as} \quad M_{\mathsf{s}}, N, H \to \infty.$$

□

By Theorem 4, for any given $\mathbf{A}_{\boldsymbol{\theta}}$, it holds that

$$J\left(a, \widehat{\mathbf{C}}_{\boldsymbol{\theta}}\right) \xrightarrow{p} J\left(a, \mathbf{C}_{\boldsymbol{\theta}}\right), \quad \text{as } N, H \to \infty. \tag{24}$$

That is, for any fixed a, the objective function evaluated with the estimated covariance converges in probability to its true counterpart. By invoking Shapiro et al. (2014, Proposition 5.1), this pointwise convergence on a compact set implies that

$$J\left(a, \widehat{\mathbf{C}}_{\boldsymbol{\theta}}\right) \to J\left(a, \mathbf{C}_{\boldsymbol{\theta}}\right) \tag{25}$$

holds uniformly on any compact subset of $\mathbb{R}^p$, $p = n^2 + 3n$, the space of free coordinates of $\mathcal{S}$, i.e. of the vectors collecting the structurally nonzero entries of $\mathbf{A}$. Then, by Shapiro et al. (2014, Proposition 5.2), the convergence of the objective functions guarantees that the minimizer of $(\widehat{\mathcal{P}}_{\boldsymbol{\theta}}(\mathcal{D}_{\mathsf{snap}}))$ converges in probability to the minimizer of $(\mathcal{P}_{\boldsymbol{\theta}})$ as $N, H \to \infty$.

The minimiser of $(\mathcal{P}_{\boldsymbol{\theta}})$ over $\mathcal{S}$ is unique and equal to $\mathbf{A}_{\boldsymbol{\theta}}$ precisely under the identifiability condition (9), which is the hypothesis $\sigma_\star > 0$ of Theorem 1 and holds for almost every parameter of the model class by Proposition 1. Under that hypothesis it follows that

$$\widehat{\mathbf{A}}_{\boldsymbol{\theta}} \xrightarrow{p} \mathbf{A}_{\boldsymbol{\theta}} \tag{26}$$

as the pooled sample size $M = NH \to \infty$ with $h \to 0$ and $NHh^{2n} \to \infty$, which completes the proof. The rate is governed by the pooled count $M = NH$ rather than by the N cells of a single episode, which is what the time kernel buys.

H Proof of Theorem 2

As preparation of the proof, we first give the following results regarding the confidence bound on the covariance matrix estimation.

Theorem 5. *Suppose Assumption 2 holds for some L_β and $\beta \in (0,1)$, and let $\epsilon > 0$ and $\delta \in (0,1)$ be given. Write $d := 2n$ for the state dimension. Assume in addition that the bandwidth is fine enough that the smoothing bias obeys*

$$\mathsf{Pr}\left\{\left\|\mathbb{E}_{\widehat{p}}\left[\mathbf{x}\mathbf{x}^\top\right] - \mathbf{C}_{\boldsymbol{\theta}}\right\|_{\mathsf{F}} \le \epsilon/2\right\} \;\ge\; 1 - \delta/2, \tag{27}$$

a condition that Theorem 4 guarantees for h sufficiently small and NHh^{2n} sufficiently large, and which separates the smoothing bias from the Monte-Carlo error controlled below.

Then, if the number of samples M_{s} entering $\widehat{\mathbf{C}}_{\boldsymbol{\theta}}$ satisfies:

$$M_{\mathsf{s}} \ge \frac{8d^2 L_\beta^4}{\epsilon^2} \log\left(\frac{4d^2}{\delta}\right) = \frac{32n^2 L_\beta^4}{\epsilon^2} \log\left(\frac{16n^2}{\delta}\right),$$

then with probability at least $1 - \beta - \delta$, we have:

$$\left\|\widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \mathbf{C}_{\boldsymbol{\theta}}\right\|_{\mathsf{F}} \le \epsilon.$$

Proof. Let $\epsilon > 0$ and $\delta \in (0,1)$ be given.

We split the error $\left\|\widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \mathbf{C}_{\boldsymbol{\theta}}\right\|_{\mathsf{F}}$ into two terms:

$$\left\|\widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \mathbf{C}_{\boldsymbol{\theta}}\right\|_{\mathsf{F}} \le \underbrace{\left\|\widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \mathbb{E}_{\widehat{p}}[\mathbf{x}\mathbf{x}^\top]\right\|_{\mathsf{F}}}_{X_1} + \underbrace{\left\|\mathbb{E}_{\widehat{p}}[\mathbf{x}\mathbf{x}^\top] - \mathbf{C}_{\boldsymbol{\theta}}\right\|_{\mathsf{F}}}_{X_2}.$$

Step 1: Control of X_2, the smoothing bias.

X_2 is exactly the quantity appearing in the hypothesis (27), so

$$\mathsf{Pr}\left\{X_2 \le \epsilon/2\right\} \ge 1 - \delta/2.$$

This term is a property of the bandwidth, not of the number of samples drawn: Step 1 of Theorem 4 shows it tends to zero in probability as $h \to 0$ with $NHh^{2n} \to \infty$, which is why it is imposed as a hypothesis here rather than bounded by M_{s}.

Step 2: Control of X_1 using Hoeffding inequality.

Under Assumption 2, for each entry (i, j) of the matrix $\mathbf{x}\mathbf{x}^\top$, we have:

$$|x^{(i)}x^{(j)}| \leq L_\beta^2 \quad \text{with probability at least } 1-\beta.$$

Conditioning on the event where this bound holds, we apply Hoeffding inequality (Hoeffding, 1963) to each entry of the matrix. For each (i, j):

$$\mathsf{Pr}\left\{\left|\frac{\sum_{m=1}^{M_{\mathsf{s}}} \mathbf{x}_m^{(i)}\mathbf{x}_m^{(j)} - \mathbb{E}_{\widehat{p}}^{i,j}}{M_{\mathsf{s}}}\right| \geq \epsilon'\right\} \leq 2\exp\left(-\frac{2M_{\mathsf{s}}\epsilon'^2}{4L_\beta^4}\right), \tag{28}$$

where

$$\mathbb{E}_{\widehat{p}}^{i,j} := \mathbb{E}_{\widehat{p}}[\mathbf{x}^{(i)}\mathbf{x}^{(j)}]. \tag{29}$$

Since $\mathbf{x}\mathbf{x}^\top$ is a $d \times d$ matrix with $d = 2n$, the Frobenius norm is controlled entrywise by $\|\cdot\|_{\mathsf{F}} \leq d\max_{ij}|\cdot|$, so achieving $X_1 \leq \epsilon/2$ requires the entrywise accuracy

$$\epsilon' := \frac{\epsilon}{2d} = \frac{\epsilon}{4n}.$$

Applying the union bound over the d^2 matrix entries then gives

$$\mathsf{Pr}\{X_1 \leq \epsilon/2\} \geq 1 - 2d^2\exp\left(-\frac{M_{\mathsf{s}}\epsilon^2}{8d^2L_\beta^4}\right).$$

To ensure this probability is at least $1-\delta/2$, it suffices to choose M_{s} such that:

$$2d^2\exp\left(-\frac{M_{\mathsf{s}}\epsilon^2}{8d^2L_\beta^4}\right) \leq \delta/2.$$

Solving this inequality yields the stated bound:

$$M_{\mathsf{s}} \geq \frac{8d^2L_\beta^4}{\epsilon^2}\log\left(\frac{4d^2}{\delta}\right) = \frac{32n^2L_\beta^4}{\epsilon^2}\log\left(\frac{16n^2}{\delta}\right).$$

Final step: Union bound.

Finally, applying union bound over the two steps, we obtain:

$$\mathsf{Pr}\left\{\left\|\widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \mathbf{C}_{\boldsymbol{\theta}}\right\|_{\mathsf{F}} \leq \epsilon\right\} \geq 1 - \beta - \delta.$$

□

Proof. **Step 1: Sample covariance approximation.** By Theorem 5, for M_{s} large enough, we have with probability at least $1-\beta-\delta$,

$$\left\|\widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \mathbf{C}_{\boldsymbol{\theta}}\right\|_{\mathsf{F}} \leq \epsilon.$$

Step 2: Perturbation of optimal solution $\widehat{\mathbf{A}}_{\boldsymbol{\theta}}$.

We give the structured least-squares perturbation argument, which is where the identifiability margin $\sigma_\star$ enters.

Write $a \in \mathbb{R}^p$ for the free coordinates of $\mathbf{A}_{\boldsymbol{\theta}}$ in $\mathcal{S}$, so that $\|a\|_2 = \|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}$, and set $\mathbf{M} := \mathbf{M}_{\mathcal{S}}(\mathbf{C}_{\boldsymbol{\theta}})$, $\widehat{\mathbf{M}} := \mathbf{M}_{\mathcal{S}}(\widehat{\mathbf{C}}_{\boldsymbol{\theta}})$ and $d := -\mathsf{vech}(\mathbf{D})$. Then $(\mathcal{P}_{\boldsymbol{\theta}})$ and $(\widehat{\mathcal{P}}_{\boldsymbol{\theta}}\,(\mathcal{D}_{\mathsf{snap}}))$ read $\min_a \|\mathbf{M}a - d\|_2^2$ and $\min_a \|\widehat{\mathbf{M}}a - d\|_2^2$, and $\mathbf{M}a = d$ holds exactly.

(i) Lipschitz dependence on the covariance. Since $\mathbf{C} \mapsto \mathbf{M}_{\mathcal{S}}(\mathbf{C})$ is linear, $\widehat{\mathbf{M}} - \mathbf{M} = \mathbf{M}_{\mathcal{S}}(\Delta\mathbf{C})$ with $\Delta\mathbf{C} = \widehat{\mathbf{C}}_{\boldsymbol{\theta}} - \mathbf{C}_{\boldsymbol{\theta}}$, and for any $\Delta\mathbf{A} \in \mathcal{S}$,

$$\left\|\Delta\mathbf{A}\,\Delta\mathbf{C} + \Delta\mathbf{C}\,\Delta\mathbf{A}^{\top}\right\|_{\mathsf{F}} \le 2\,\|\Delta\mathbf{C}\|_2\,\|\Delta\mathbf{A}\|_{\mathsf{F}} \le 2\,\|\Delta\mathbf{C}\|_{\mathsf{F}}\,\|\Delta\mathbf{A}\|_{\mathsf{F}}\,,$$

so $\|\widehat{\mathbf{M}} - \mathbf{M}\|_2 \le 2\|\Delta\mathbf{C}\|_{\mathsf{F}} \le 2\epsilon$ on the event of Step 1.

(ii) Residual control. As $\widehat{a}$ minimises $\|\widehat{\mathbf{M}}\,\cdot - d\|_2$, $\|\widehat{\mathbf{M}}\widehat{a} - d\|_2 \le \|\widehat{\mathbf{M}}a - d\|_2 = \|(\widehat{\mathbf{M}} - \mathbf{M})a\|_2$. Hence

$$\left\|\widehat{\mathbf{M}}(\widehat{a} - a)\right\|_2 \le \left\|\widehat{\mathbf{M}}\widehat{a} - d\right\|_2 + \left\|d - \widehat{\mathbf{M}}a\right\|_2 \le 2\left\|\widehat{\mathbf{M}} - \mathbf{M}\right\|_2 \|a\|_2 \le 4\epsilon\,\|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}\,.$$

(iii) Stability of the smallest singular value. By Weyl's inequality and (i), $\sigma_{\min}(\widehat{\mathbf{M}}) \ge \sigma_\star - 2\epsilon \ge \sigma_\star/2$ whenever $\epsilon \le \sigma_\star/4$. In particular $\widehat{\mathbf{M}}$ still has full column rank, so $\widehat{a}$ is unique.

Combining (ii) and (iii),

$$\left\|\widehat{\mathbf{A}}_{\boldsymbol{\theta}} - \mathbf{A}_{\boldsymbol{\theta}}\right\|_{\mathsf{F}} = \|\widehat{a} - a\|_2 \le \frac{\left\|\widehat{\mathbf{M}}(\widehat{a} - a)\right\|_2}{\sigma_{\min}(\widehat{\mathbf{M}})} \le \frac{4\epsilon\,\|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}}{\sigma_\star/2} = \frac{8\,\|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}}{\sigma_\star}\,\epsilon,$$

which is the claim. Equivalently, in terms of the modulus $\mu = 2\sigma_\star^2$ of the restricted Hessian, the bound reads $8\|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}\epsilon\sqrt{2/\mu}$.

□

I Proof of Theorem 3

We bound the perturbation of the dominant spectral projector, which is the object whose diagonal gives the participation factors (12). The argument is carried out with the resolvent rather than with a Hermitian eigenvector bound, because $\mathbf{A}_{\boldsymbol{\theta}}$ is not symmetric; the classical $1/\delta_\lambda$ estimate is recovered at the end, and this is also what defines the constant $c_{\boldsymbol{\theta}}$ used in Theorem 3. Writing the proof for a group of m eigenvalues costs nothing and covers, at no extra effort, the symmetric networks of Appendix J whose dominant eigenvalue is repeated; the reader interested only in the simple case may set $m = 1$ throughout, in which case Λ_{d} is the single dominant eigenvalue, $\mathbf{P}_{\mathsf{d}}$ is the rank-one projector $\mathbf{v}_{\mathsf{d}}\mathbf{w}_{\mathsf{d}}^{\top}/\mathbf{w}_{\mathsf{d}}^{\top}\mathbf{v}_{\mathsf{d}}$, and δ_λ is the eigengap (11).

Setup. Let $\Lambda_{\mathsf{d}} \subset \mathsf{spec}(\mathbf{A}_{\boldsymbol{\theta}})$ be the dominant group of m eigenvalues of largest real part and $\delta_\lambda = \mathsf{dist}(\Lambda_{\mathsf{d}}, \mathsf{spec}(\mathbf{A}_{\boldsymbol{\theta}}) \setminus \Lambda_{\mathsf{d}}) > 0$, which for $m = 1$ is the eigengap (11). Let Γ_{d} be the boundary of the $\delta_\lambda/2$-neighbourhood of Λ_{d}, a union of at most m circles of radius $\delta_\lambda/2$, so its total length obeys $\ell(\Gamma_{\mathsf{d}}) \le \pi m \delta_\lambda$. By construction Γ_{d} encloses Λ_{d} and no other eigenvalue, and the Riesz projector is

$$\mathbf{P}_{\mathsf{d}} = \frac{1}{2\pi\mathrm{i}} \oint_{\Gamma_{\mathsf{d}}} \mathbf{R}(z)\,\mathrm{d}z, \qquad \mathbf{R}(z) := (z\mathbf{I} - \mathbf{A}_{\boldsymbol{\theta}})^{-1}, \qquad r_{\mathsf{d}} := \sup_{z \in \Gamma_{\mathsf{d}}} \|\mathbf{R}(z)\|_2 < \infty.$$

Step 1: resolvent perturbation. Put $\Delta\mathbf{A} := \widehat{\mathbf{A}}_{\boldsymbol{\theta}} - \mathbf{A}_{\boldsymbol{\theta}}$ and suppose $r_{\mathsf{d}}\|\Delta\mathbf{A}\|_2 \le \frac{1}{2}$. For $z \in \Gamma_{\mathsf{d}}$, $z\mathbf{I} - \widehat{\mathbf{A}}_{\boldsymbol{\theta}} = (\mathbf{I} - \mathbf{R}(z)\Delta\mathbf{A})\,(z\mathbf{I} - \mathbf{A}_{\boldsymbol{\theta}})$ with $\|\mathbf{R}(z)\Delta\mathbf{A}\|_2 \le \frac{1}{2}$, so the Neumann series converges, $z\mathbf{I} - \widehat{\mathbf{A}}_{\boldsymbol{\theta}}$ is invertible on Γ_{d}, and

$$\left\|\left(z\mathbf{I} - \widehat{\mathbf{A}}_{\boldsymbol{\theta}}\right)^{-1} - \mathbf{R}(z)\right\|_2 \le \|\mathbf{R}(z)\|_2 \sum_{k \ge 1} \|\mathbf{R}(z)\Delta\mathbf{A}\|_2^k \le \frac{r_{\mathsf{d}}^2\,\|\Delta\mathbf{A}\|_2}{1 - \frac{1}{2}} = 2r_{\mathsf{d}}^2\,\|\Delta\mathbf{A}\|_2\,.$$

In particular Γ_{d} avoids $\mathsf{spec}(\widehat{\mathbf{A}}_{\boldsymbol{\theta}})$, so the perturbed projector $\widehat{\mathbf{P}}_{\mathsf{d}}$ defined by the same contour is well posed and, by continuity of the spectrum, encloses exactly m eigenvalues of $\widehat{\mathbf{A}}_{\boldsymbol{\theta}}$.

Step 2: projector bound. Integrating the previous display over Γ_{d},

$$\left\|\widehat{\mathbf{P}}_{\mathsf{d}} - \mathbf{P}_{\mathsf{d}}\right\|_2 \leq \frac{\ell(\Gamma_{\mathsf{d}})}{2\pi}\, 2r_{\mathsf{d}}^2\, \|\Delta\mathbf{A}\|_2 \leq m\, \delta_\lambda\, r_{\mathsf{d}}^2\, \|\Delta\mathbf{A}\|_2 \,. \tag{30}$$

Step 3: from the projector to the participation factors. Since $\kappa_{\boldsymbol{\theta}}(j) = [\mathbf{P}_{\mathsf{d}}]_{jj} = \mathbf{e}_j^\top \mathbf{P}_{\mathsf{d}} \mathbf{e}_j$,

$$\max_j |\widehat{\kappa}_{\boldsymbol{\theta}}(j) - \kappa_{\boldsymbol{\theta}}(j)| \leq \left\|\widehat{\mathbf{P}}_{\mathsf{d}} - \mathbf{P}_{\mathsf{d}}\right\|_2 .$$

Step 4: combination with Theorem 2. On the event of Theorem 2 we have $\|\Delta\mathbf{A}\|_2 \leq \|\Delta\mathbf{A}\|_{\mathsf{F}} \leq 8\|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}\epsilon/\sigma_\star$, and the hypothesis $\epsilon \leq \sigma_\star/(16 r_{\mathsf{d}}\|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}})$ is exactly $r_{\mathsf{d}}\|\Delta\mathbf{A}\|_2 \leq \frac{1}{2}$, so Step 1 applies. Substituting into (30),

$$\max_j |\widehat{\kappa}_{\boldsymbol{\theta}}(j) - \kappa_{\boldsymbol{\theta}}(j)| \leq \left\|\widehat{\mathbf{P}}_{\mathsf{d}} - \mathbf{P}_{\mathsf{d}}\right\|_2 \leq \frac{8\, m\, \delta_\lambda\, r_{\mathsf{d}}^2\, \|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}}{\sigma_\star}\, \epsilon = \frac{c_{\boldsymbol{\theta}}\, \|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}}}{\sigma_\star\, \delta_\lambda}\, \epsilon, \qquad c_{\boldsymbol{\theta}} := 8\, m\, \delta_\lambda^2\, r_{\mathsf{d}}^2,$$

which is the statement of Theorem 3, with $\epsilon_0 := \sigma_\star/(16 r_{\mathsf{d}}\|\mathbf{A}_{\boldsymbol{\theta}}\|_{\mathsf{F}})$ the threshold quoted there. □

The constant, and reduction to the classical bound. If $\mathbf{A}_{\boldsymbol{\theta}}$ is normal and λ_{d} is simple, then $m = 1$ and $r_{\mathsf{d}} = \sup_{z\in\Gamma_{\mathsf{d}}} 1/\mathsf{dist}(z, \mathsf{spec}(\mathbf{A}_{\boldsymbol{\theta}})) = 2/\delta_\lambda$, so $c_{\boldsymbol{\theta}} = 8\,\delta_\lambda^2\,(2/\delta_\lambda)^2 = 32$ as quoted in Theorem 3, and (30) becomes $\|\widehat{\mathbf{P}}_{\mathsf{d}} - \mathbf{P}_{\mathsf{d}}\|_2 \leq 4\|\Delta\mathbf{A}\|_2/\delta_\lambda$, i.e. the $1/\delta_\lambda$ scaling of Davis–Kahan-type eigenvector estimates (Yu et al., 2015). In general $c_{\boldsymbol{\theta}}$ exceeds 32 only through the non-normality of $\mathbf{A}_{\boldsymbol{\theta}}$, since $\delta_\lambda r_{\mathsf{d}}/2$ is the factor by which the resolvent is amplified relative to a normal matrix with the same spectrum. Stating the result through $\mathbf{P}_{\mathsf{d}}$ is what allows the same guarantee to cover networks whose modules reach criticality together, where $m > 1$ and a single eigenvector is not determined by the model; the separation δ_λ is then measured between the dominant group and the remaining spectrum, and is positive in both synthetic models of Section 5 and Appendix J ($\delta_\lambda = 0.516$ with $m = 3$ for the synchronous modular network, $\delta_\lambda = 0.703$ with $m = 1$ for the localized one).

Consequences for target selection. The projector bound is what makes the diagonal design of Lemma 1 usable from finite data: the first-order eigenvalue shift is proportional to the participation factor $\kappa_{\boldsymbol{\theta}}(s_i)$ at the actuated node, so bounding $\max_j |\widehat{\kappa}_{\boldsymbol{\theta}}(j) - \kappa_{\boldsymbol{\theta}}(j)|$ directly bounds the error in the quantity used to rank candidate sites, and the ranking is preserved whenever the gap between the leading and trailing candidates exceeds twice that error. Two consequences are worth noting. Working with the spectral projector covers networks whose modules reach criticality simultaneously, where the dominant eigenvalue is repeated and no single eigenvector is determined by the model, while reproducing the classical eigenvector product when the critical mode is simple. And both quantities entering the bound, $\sigma_\star$ and δ_λ, are computable from the identified model, so the reliability of a proposed intervention can be assessed before it is applied.

J Simulation Model Details

Here, we provide a complete description of the gene-protein regulatory network model used to generate the synthetic data for our experiments. The model is parameterised by the number of modules R and the number of genes per module n_{mod}, giving $n = R\, n_{\mathrm{mod}}$ genes and a state dimension $2n$. Throughout the paper we use $R = 3$ and $n_{\mathrm{mod}} = 15$, hence $n = 45$ genes and $2n = 90$ state variables. Two instantiations of the same network and the same governing equations are evaluated, differing only in how the bifurcation thresholds of the modules are coupled (Appendix J.2): the *synchronous modular* network of the main text, in which one shared threshold drives all modules to criticality together, and the *localized modular* network of Appendix J.5, in which module-specific thresholds place a single module at criticality.

J.1 Network Structure and Topology

The network consists of $2n$ state variables: the concentrations of the n mRNAs and of the n proteins. As illustrated in Figure 4, the $n = 45$ genes are partitioned into $R = 3$ disjoint modules of $n_{\mathrm{mod}} = 15$ genes. Within each module the protein of the module's last gene acts as the master regulator: it

activates the transcription of all n_{mod} genes of its own module — and of no other module — via Hill-type regulation, while every mRNA is translated into its own protein. The master protein thereby receives a positive feedback signal through the translation of its own gene, creating one self-reinforcing loop per module, which gives rise to bistability and bifurcation phenomena.

In the global indexing used throughout the paper, the master regulator of module r is $P_{\mu(r)}$ with $\mu(r) = r\, n_{\text{mod}}$, i.e. P_{15}, P_{30} and P_{45}. Since the modules do not interact, $\mathbf{A}_{\boldsymbol{\theta}}$ is block diagonal with one $2n_{\text{mod}} \times 2n_{\text{mod}}$ block per module, and the concentration of $P_{\mu(r)}$ acts as a proxy for the state of module r. Because the participation factor is supported on both the mRNA and the protein of a master gene, the master-regulator node set of the $2n$-dimensional state is $\{\text{mRNA}_{15}, \text{mRNA}_{30}, \text{mRNA}_{45}, P_{15}, P_{30}, P_{45}\}$.

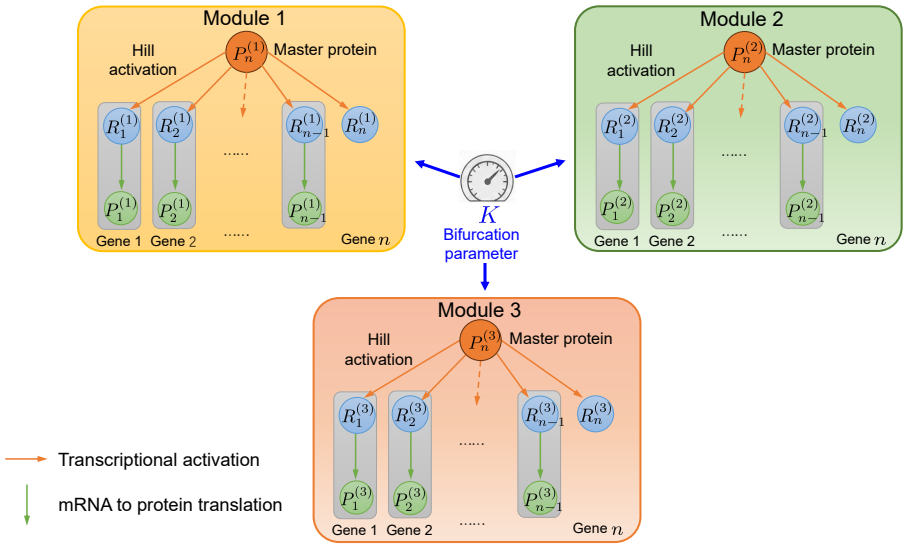

Figure 4: Modular master-regulator topology used by both synthetic models ($R = 3$ modules; the figure writes n for the module size $n_{\text{mod}} = 15$). $R_i^{(r)}$ and $P_i^{(r)}$ are the mRNA and the protein of gene i of module r, i.e. m_j and P_j with global index $j = (r-1)\, n_{\text{mod}} + i$ in the SDEs of Appendix J.2; the master protein $P_n^{(r)}$ is $P_{\mu(r)}$. Orange arrows: the master protein activates the transcription of every gene of its own module, including its own, through a Hill term with half-activation threshold K_r. Green arrows: each mRNA is translated into its protein; the translation of the master gene closes the module's self-reinforcing loop. The two models of Appendix J.2 differ only in the coupling of the thresholds: the dial K marks the single shared threshold $K_1 = \cdots = K_R = K_{\text{master}}$ of the synchronous modular network, whereas the localized modular network gives each module its own K_r and sweeps only K_1.

J.2 System Dynamics and Governing Equations

Both synthetic models used in this paper are generated by *one* set of governing equations and differ only in how the bifurcation thresholds of the R modules are coupled. Let $r(i) \in \{1, \ldots, R\}$ denote the index of the module containing gene i, and recall that the master regulator of module r is the protein $P_{\mu(r)}$ with $\mu(r) = r\, n_{\text{mod}}$. Writing K_r for the activation threshold of module r, the temporal evolution of the mRNA (m_i) and protein (P_i) concentrations of each gene $i \in \{1, \ldots, n\}$ is described by the system of stochastic differential equations (SDEs)

$$\frac{dm_i}{dt} = \underbrace{\alpha_{\text{basal},i}}_{\text{Basal Rate}} + \underbrace{\alpha_{\text{activated}} \frac{P_{\mu(r(i))}^{n_{\text{master}}}}{K_{r(i)}^{n_{\text{master}}} + P_{\mu(r(i))}^{n_{\text{master}}}}}_{\text{Activated Transcription (Hill function, own module's master)}} - \underbrace{\gamma_m m_i}_{\text{Degradation}} + \xi_{m,i}(t),$$

$$\frac{dP_i}{dt} = \underbrace{\beta_i m_i}_{\text{Translation}} - \underbrace{\gamma_{p,i} P_i}_{\text{Degradation}} + \xi_{P,i}(t),$$

where:

- The activated transcription term is a standard Hill function representing the cooperative binding of the master regulator $P_{\mu(r(i))}$ of gene i's own module to the promoter region of gene i. Genes are therefore driven only by the master of their own module, which makes $\mathbf{A}_{\boldsymbol{\theta}}$ block diagonal.
- K_r is the Michaelis-Menten constant of module r, i.e. the concentration of $P_{\mu(r)}$ required for half-maximal activation. These are the parameters that are varied to induce the bifurcation, and their coupling defines the two models below.
- n_{master} is the Hill coefficient, representing the cooperativity of binding.
- $\xi_{m,i}(t)$ and $\xi_{P,i}(t)$ are stochastic terms representing noise (Appendix J.3).

Model 1: synchronous modular network (main text). All modules share a single threshold,

$$K_1 = K_2 = \cdots = K_R \;=\; K_{\text{master}},$$

and the per-gene parameters $(\alpha_{\text{basal},i}, \beta_i, \gamma_{p,i})$ are tiled identically across modules, so the R modules are isomorphic. Consequently all modules approach their saddle-node bifurcation at the same value of K_{master}, the dominant eigenvalue of $\mathbf{A}_{\boldsymbol{\theta}}$ has multiplicity R, and the participation factor is distributed equally over the $2R = 6$ master nodes. Snapshots are taken at $H = 10$ equispaced values of K_{master} in $[K_{\text{bif}} - 3,\ K_{\text{bif}} - 0.4]$, where $K_{\text{bif}} = 6.2421$ is the common saddle-node threshold. Since the dominant eigenspace is R-dimensional, κ is evaluated as the diagonal of the spectral projector onto that eigenspace, which coincides with $\mathbf{w}_{\mathsf{d}}(i)\mathbf{v}_{\mathsf{d}}(i)$ whenever the dominant eigenvalue is simple.

Model 2: localized modular network (Appendix J.5). Each module keeps its own threshold. Module 1 is driven towards criticality while the remaining modules are pinned far below theirs,

$$K_1 \in [K_1^{\text{bif}} - 3,\ K_1^{\text{bif}} - 0.05], \qquad K_2 = 0.35\, K_{\text{bif}}, \qquad K_3 = 0.45\, K_{\text{bif}},$$

with $K_1^{\text{bif}} = 6.2385$. Modules 2 and 3 therefore remain on their high-expression branch far from any bifurcation, the dominant eigenvalue is simple, and the dominant mode is confined to module 1. The participation factor is then supported on the two master nodes $\{P_{15}, \text{mRNA}_{15}\}$ and κ reduces to $\mathbf{w}_{\mathsf{d}}(i)\mathbf{v}_{\mathsf{d}}(i)$. Snapshots are again taken at $H = 10$ equispaced values of K_1.

J.3 Noise Model

The stochastic terms $\xi(t)$ represent additive Gaussian white noise, which models random fluctuations in the biochemical reactions. The SDEs are of the form $dX_t = f(X_t)dt + \sigma dW_t$, where $f(X_t)$ is the deterministic drift part of the equations above, dW_t is a standard Wiener process, and σ is the noise strength. The noise covariance matrix is $\mathbf{D} = \sigma^2 \mathbf{I}_{2n}$.

For the numerical simulations, we use the Euler-Maruyama method with a time step Δt. The noise term for each state variable at each time step is implemented as:

$$\text{Noise} = \sigma\sqrt{\Delta t} \cdot \mathcal{N}(0, 1)$$

where $\mathcal{N}(0, 1)$ is a random variable drawn from a standard normal distribution. After each Euler–Maruyama step the state is projected onto the non-negative orthant (concentrations are clipped at zero), matching the released implementation. The noise strength σ was set to a constant value of 0.05 for both models (i.e., $\mathbf{D} = 0.0025 \cdot \mathbf{I}_{90}$).

Snapshots are generated on the high-expression branch as the system approaches the tipping point. For each selected threshold value, we simulate 2000 independent realizations (representing individual cells) using Euler-Maruyama integration over a simulation horizon of $T = 800$ time units with step size $\Delta t = 0.1$. The final state of each realization is taken as a single-cell measurement, yielding one snapshot of 2000 cells. This process is repeated over the $H = 10$ threshold values of the model in question, producing snapshot datasets with gradual shifts in covariance structure as criticality is approached. Ground-truth Jacobians are computed at the deterministic steady state on the high-expression branch for each threshold value; the Jacobian of the last episode provides the reference support pattern used by the identification methods.

J.4 Parameter Settings

Both models share the same parameter specification, given in the table below. The per-gene parameters are drawn once for a single module of $n_{\mathrm{mod}} = 15$ genes and then tiled across the $R = 3$ modules, so that the modules are structurally and parametrically identical; the two models differ only in the coupling of the module thresholds K_r described in Appendix J.2.

Parameter	Value / Distribution	Description
R	3	Number of modules.
n_{mod}	15	Genes per module ($n = 45$, state dimension $2n = 90$).
$\alpha_{\mathrm{basal},i}$	$\mathcal{U}(0.03, 0.06)$	Basal transcription rate for gene i.
$\alpha_{\mathrm{activated}}$	6.0	Maximum activated transcription rate.
K_r	Varied	Bifurcation parameter: activation threshold of module r.
n_{master}	4	Hill coefficient for master-regulator activation.
β_i	$\mathcal{U}(1.6, 2.4)$	Translation rate for mRNA i.
γ_m	1.0	mRNA degradation rate (uniform for all genes).
$\gamma_{p,i}$	$\mathcal{U}(0.8, 1.2)$	Protein degradation rate for protein i.
σ	0.05	Additive Gaussian noise strength ($\mathbf{D} = 0.0025 \cdot \mathbf{I}_{90}$).
H	10	Number of snapshot episodes (threshold values).
Cells per episode	2000	Independent Euler-Maruyama realizations per threshold.

Table 2: Parameter specification for the two modular bifurcation models. Per-gene values are sampled i.i.d. from the indicated uniform distributions for one module and tiled across the $R = 3$ modules, making the modules identical; the models differ only in the coupling of the thresholds K_r (Appendix J.2).

J.5 Detailed Simulation Results: Localized Modular Network

The main text drives all three modules to criticality simultaneously. Here we report the complementary *localized modular* instantiation of the same 90-dimensional network (Model 2 of Appendix J.2), in which each module carries its own threshold and only module 1 is driven towards its saddle-node while modules 2 and 3 are held far below theirs. The dominant eigenvalue is then simple and its mode is confined to module 1, so the participation factor is supported on the two-node master set $\mathcal{M} = \{P_{15}, \mathrm{mRNA}_{15}\}$, with $|\kappa_i| = 0.746$ and 0.665 respectively, and κ reduces exactly to $\mathbf{w}_{\mathsf{d}}(i)\mathbf{v}_{\mathsf{d}}(i)$. Retaining the strict criterion that *all* k top-ranked nodes must be master regulators, the chance levels are 2.2% for $k = 1$ and $2.5 \times 10^{-2}\%$ for $k = 2$. All other settings match the main text: 100 Monte Carlo trials at each $N \in \{10, 20, 30, 50, 70, 100, 150\}$, with all three methods receiving identical cell subsamples.

Figure 5 shows that the proposed method again localises its signal exactly: κ is supported on P_{15} and mRNA_{15} and is numerically zero on the remaining 88 nodes, so the two regulators of the critical module are isolated from the rest of the network even though the two dormant modules are structurally identical to the critical one. The scInTime statistic places its largest value inside the critical module (0.61 at P_{15}), but the surrounding proteins reach 0.60, so the regulator is not separated from the background.

Figure 6 shows that this localised setting sharpens the comparison in the sample-starved regime. With $N = 10$ cells per episode — nine times fewer than the state dimension — the proposed method already identifies the regulator in 40% of trials, against 5% for SOTA and 0% for scInTime, with paired McNemar significances of $p = 2.8 \times 10^{-9}$ and $p = 1.8 \times 10^{-12}$ respectively. The proposed method reaches 100% at $N = 50$ and holds it, whereas SOTA requires $N = 150$ to match it and scInTime stays at 0% at every sample size ($p < 10^{-11}$ throughout). Moreover the $k = 1$ and $k = 2$ curves of the proposed method coincide: whenever it ranks a master regulator first it recovers the *entire* two-node master set, which is exactly the information required to place a diagonal re-stabilising input.

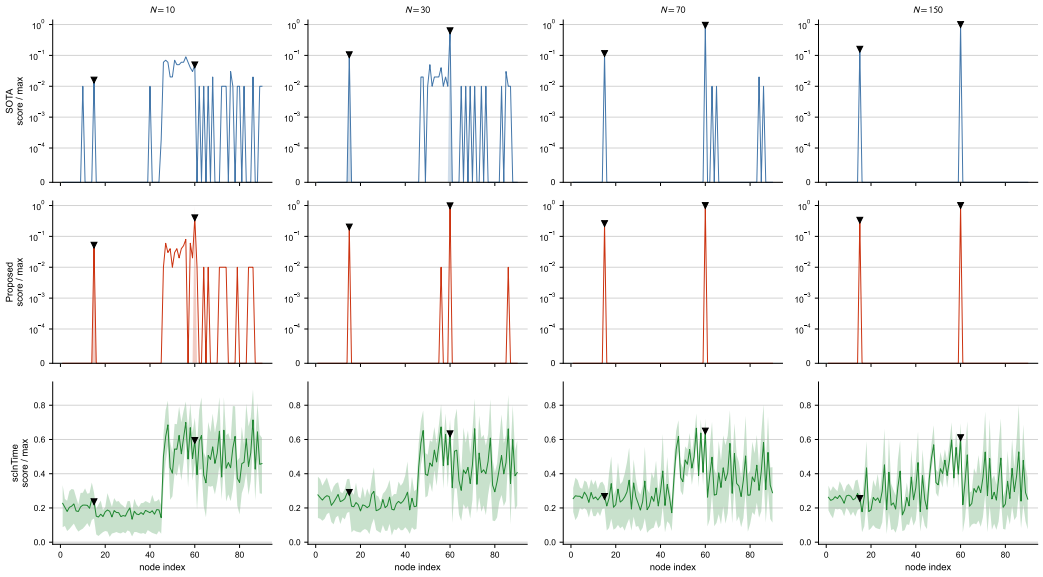

Figure 5: Localized modular network: per-node mean score (line) and interquartile range (band) over 100 trials, normalised by the per-trial maximum; black triangles mark the two master regulators mRNA_{15} and P_{15}. Rows: SOTA, Proposed, scInTime; columns: increasing number of cells per episode N. The proposed method concentrates κ on the two regulators of the critical module and is numerically zero on every other node.

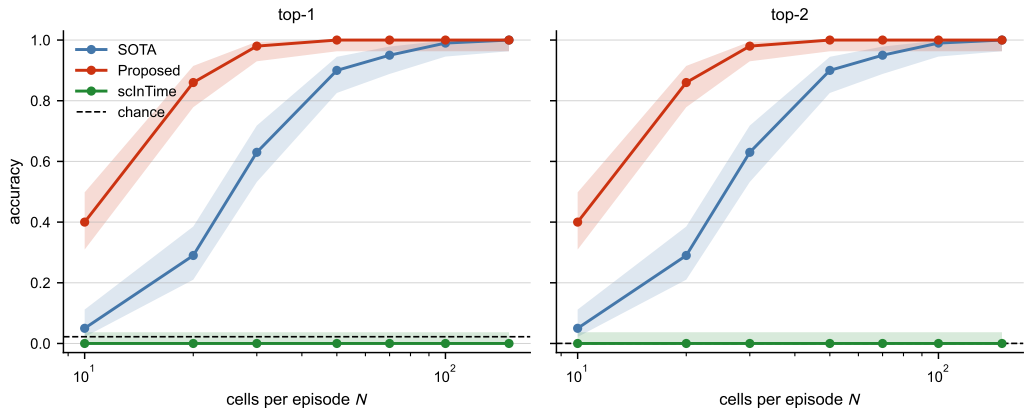

Figure 6: Localized modular network: regulatory-node identification accuracy versus cells per episode N, where a trial counts as correct only if *all* k top-ranked nodes are master regulators ($k = 1, 2$). Bands are Wilson 95% confidence intervals over 100 trials and the dashed line is the chance level. The proposed method leads at every N, and its $k = 1$ and $k = 2$ curves coincide, i.e. it recovers the complete master set whenever it ranks a master regulator first.

K EXPERIMENTAL DATASET, PRE-PROCESSING, AND BIOLOGICAL VALIDATION

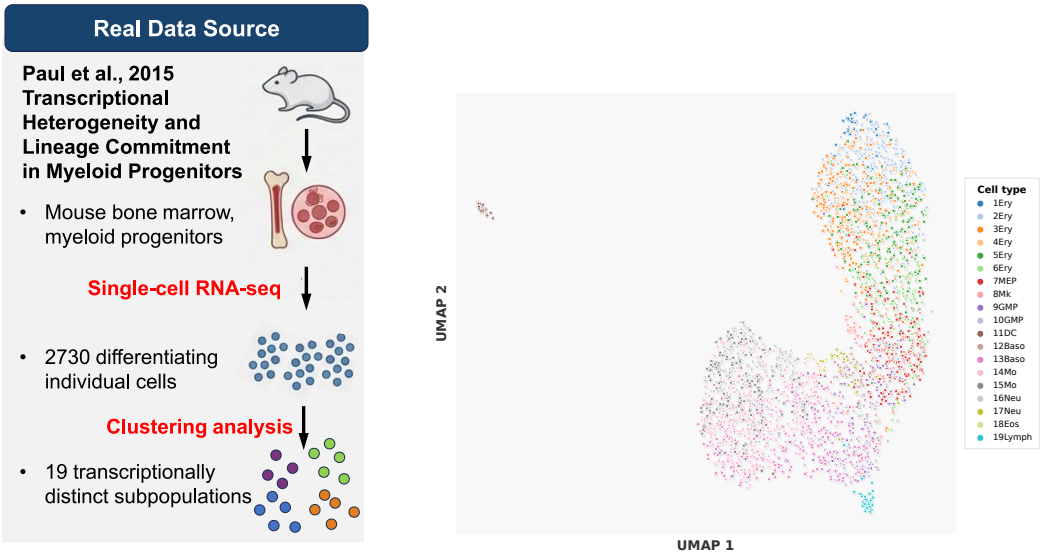

(a) Data source overview

(b) Cell-type clusters (UMAP)

Figure 7: Real data source and cell-type landscape from the Paul et al. (2015) scRNA-seq dataset. (a) Paul et al. (2015) profiled 2,730 mouse bone marrow myeloid progenitors via MARS-seq; clustering identified 19 transcriptionally distinct subpopulations primed toward seven hematopoietic fates (GEO: GSE72857). (b) UMAP visualization of the 2,730 cells coloured by the 19 lineage clusters.

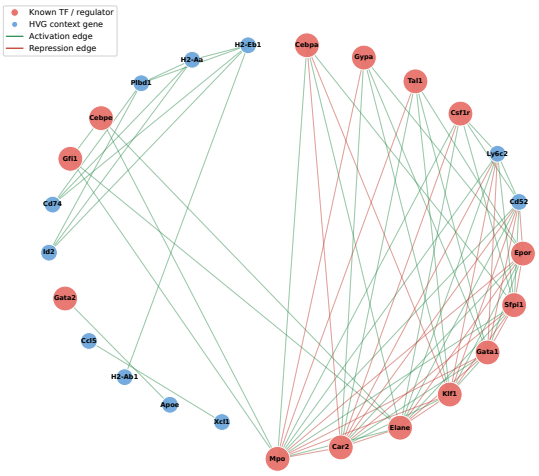

Figure 8: Inferred signed structural GRN prior $\widehat{\mathbf{A}}_{\mathrm{ref}}$ constructed from a 50-gene panel, showing key regulatory genes and their connections.

K.1 MRNA-LEVEL REDUCTION

When $\gamma_{p,i} \gg \gamma_{m,i}$, the protein concentration quasi-statically satisfies $P_i \approx (\beta_i/\gamma_{p,i})m_i$. Substituting into (1) yields a reduced n-dimensional SDE with system matrix

$$\mathbf{A}_{\boldsymbol{\theta}}^{\mathrm{red}} = -\Delta_{\mathrm{r},\boldsymbol{\theta}} + \Gamma_{\mathrm{p},\boldsymbol{\theta}}\,\mathrm{diag}\Big(\frac{\beta_i}{\gamma_{p,i}}\Big)\,, \tag{31}$$

which preserves the sparsity structure of $\mathbf{A}_{\boldsymbol{\theta}}$ and satisfies (5). This model is used in Section 6.

K.2 Dataset and pre-processing

The Paul et al. (2015) dataset was downloaded from NCBI GEO (accession GSE72857). Raw UMI counts were normalized to 10^4 counts per cell, followed by a $\log(x+1)$ transformation. To define our $n = 50$ gene panel, we employed a hybrid selection strategy: (i) 26 curated hematopoietic transcription factors and regulator genes to serve as a control-relevant prior, and (ii) 24 additional highly variable genes (HVGs). These HVGs were identified and grouped into 20 mean-expression bins; within-bin dispersion is z-score normalized, excluding any genes already present in the curated set. To construct the reference structural network prior $\widehat{\mathbf{A}}_{\text{ref}}$, pairwise Spearman correlation coefficients (r_s) were computed across all cells for the selected 50 genes. Edges demonstrating statistically significant correlations ($|r_s| \geq 0.3$ and $p < 0.05$) were retained and assigned a sign to denote the regulatory relationship (activation for $r_s > 0$; repression for $r_s < 0$). To finalize the system matrix, diagonal entries were set to -1, and the retained off-diagonal entries were assigned values of ± 0.3.

K.3 Episode selection, covariance estimation, and system matrix reconstruction

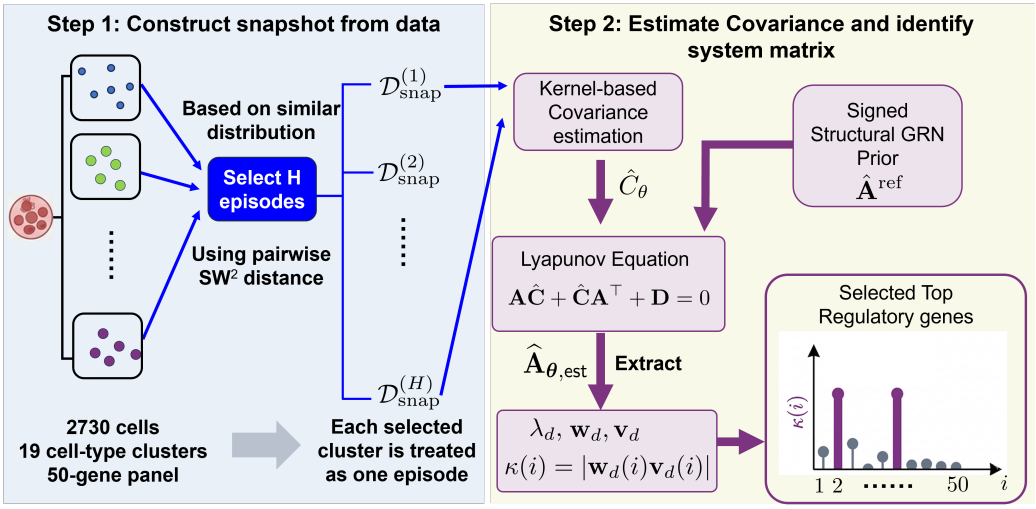

Figure 9: Experimental validation framework. Step 1: $H = 15$ episodes are selected from the 19 cell-type clusters via pairwise SW_2 distances to form the snapshot dataset. Step 2: a kernel-based covariance estimate $\widehat{\mathbf{C}}_{\boldsymbol{\theta}}$ is computed and, together with the signed structural GRN prior $\widehat{\mathbf{A}}_{\text{ref}}$, used to solve $(\widehat{\mathcal{P}}_{\boldsymbol{\theta}}(\mathcal{D}_{\text{snap}}))$, yielding $\widehat{\mathbf{A}}_{\boldsymbol{\theta},\text{est}}$ and the per-gene regulatory indicator $\kappa(i)$.

The overall experimental pipeline is illustrated in Figure 9.

Episode selection. Pairwise SW_2 distances between all 19 cell-type distributions were computed in the 50-dimensional HVG space (1,000 projections, seed 42). From all $\binom{19}{15} = 3{,}876$ subsets of 15 cell types, those with maximum pairwise $\text{SW}_2 \leq 1.334$ were retained. The subset with the smallest maximum distance, spanning both erythroid and myeloid states, was selected. SW_2 is preferred over Euclidean distance for its robustness to cell-count imbalance (Kolouri et al., 2017).

The 15 selected episodes and their cell counts are listed in Table 3.

Covariance estimation and system matrix reconstruction. As shown in Table 3, cell counts per episode range from 31 (19Lymph) to 373 (14Mo). From each episode k ($k = 1, \ldots, H$; $H = 15$), we draw $N = 25$ cells uniformly at random without replacement and centre them by the per-episode mean: $\tilde{\mathbf{x}}_m^{(k)} = \mathbf{x}_m^{(k)} - \bar{\mathbf{x}}^{(k)}$, where $\bar{\mathbf{x}}^{(k)} = N^{-1}\sum_{m=1}^{N}\mathbf{x}_m^{(k)}$. Pooling these centred vectors across all episodes yields $M = H \cdot N$ samples in the n-dimensional state space, which are rescaled to $[-1, 1]$

Table 3: Cell counts of the 15 selected episodes (cell types) from the Paul et al. (2015) dataset. Episodes are ordered by their index in the differentiation trajectory.

k	Cell type	Cells	k	Cell type	Cells
1	1Ery	43	9	7MEP	167
2	2Ery	329	10	8Mk	68
3	3Ery	246	11	9GMP	63
4	4Ery	124	12	10GMP	153
5	5Ery	180	13	12Baso	69
6	6Ery	173	14	13Baso	300
7	14Mo	373	15	19Lymph	31
8	15Mo	186			

via a global min–max transform. Let $\bar{\mathbf{x}}_j$ denote the normalised centred data point after the min–max transform.

Following the kernel conditional density estimation framework of Section 4, we use a Gaussian kernel directly in the n-dimensional state space with bandwidth $h = 1.0$; the query point for sample i is taken to be $\bar{\mathbf{x}}_i$ itself, so that the kernel weight measures the local sample density around each observation. The kernel weight for data point i is then

$$w_i = \frac{\sum_{j=1}^{M} K\left(\frac{\bar{\mathbf{x}}_i - \bar{\mathbf{x}}_j}{h}\right)}{\sum_{l=1}^{M}\sum_{j=1}^{M} K\left(\frac{\bar{\mathbf{x}}_l - \bar{\mathbf{x}}_j}{h}\right)}, \qquad M = H \cdot N, \tag{32}$$

where $K(\cdot)$ is the Gaussian kernel. The covariance estimate is computed as $\widehat{\mathbf{C}}_{\boldsymbol{\theta}} = c^{-1}\sum_{i=1}^{M} w_i\, \tilde{\mathbf{x}}_i\, \tilde{\mathbf{x}}_i^{\top}$ with normalisation constant $c = 5$, using the original (un-normalised) centred vectors.

Given $\widehat{\mathbf{C}}_{\boldsymbol{\theta}}$ and the diffusion matrix $\mathbf{D} = \sigma^2\mathbf{I}_n$ ($\sigma = 0.1$), we solve $(\widehat{\mathcal{P}}_{\boldsymbol{\theta}}(\mathcal{D}_{\mathsf{snap}}))$ subject to the structural constraint (5) imposed by $\widehat{\mathbf{A}}_{\mathrm{ref}}$, obtaining the optimal estimate $\widehat{\mathbf{A}}_{\boldsymbol{\theta},\mathsf{est}}$. The dominant right and left eigenvectors $\mathbf{v}_{\mathsf{d},\boldsymbol{\theta}}$ and $\mathbf{w}_{\mathsf{d},\boldsymbol{\theta}}$ of $\widehat{\mathbf{A}}_{\boldsymbol{\theta},\mathsf{est}}$ are then extracted and the participation factor κ_i of (12) is computed as in Section 6.

Run-time identifiability check. The identifiability condition (9) is not assumed here but tested. In this experiment the free subspace $\mathcal{S}$ is the sparsity pattern of $\widehat{\mathbf{A}}_{\mathrm{ref}}$ rather than the mRNA–protein block pattern of (7), so Proposition 1 does not apply verbatim: its Step 1 reduction holds for any fixed pattern, but the witness of Step 2 is specific to (7). We therefore assemble $\mathbf{M}_{\mathcal{S}}(\widehat{\mathbf{C}}_{\boldsymbol{\theta}})$ for the inferred pattern and verify numerically that it has full column rank. Because $\mathbf{C} \mapsto \mathbf{M}_{\mathcal{S}}(\mathbf{C})$ is linear, $\sigma_\star$ carries the scale of $\mathbf{D}$, so we report the scale-free relative margin $\sigma_\star/\|\widehat{\mathbf{C}}_{\boldsymbol{\theta}}\|_2$ alongside $\widehat{\mathbf{A}}_{\boldsymbol{\theta},\mathsf{est}}$. A vanishing or near-vanishing margin signals that the recovered κ ranking is not trustworthy for the given prior, which makes the diagnostic a practical safeguard rather than a theoretical caveat.

K.4 Biological descriptions of top-5 genes

Mpo ($\kappa = 0.508$) is the canonical granulocyte marker and among the most differentially expressed genes separating myeloid from erythroid progenitors in Paul et al. (2015). Its dominance is consistent with the granulocyte–monocyte bifurcation being the primary transition in the episode sequence (Iwasaki & Akashi, 2007). *Klf1* ($\kappa = 0.415$) is the master regulator of erythropoiesis (Tallack et al., 2010), activating erythroid genes while repressing myeloid fate via PU.1 downregulation (Frontelo et al., 2007).

Gata1 ($\kappa = 0.384$) is indispensable for erythroid and megakaryocyte development (Ferreira et al., 2005). Its co-identification with Klf1 at ranks 2 and 3 reflects their known cooperation on shared erythroid target promoters (Tallack et al., 2010). *Car2* ($\kappa = 0.345$) is a downstream erythroid effector of the Gata1–Klf1 circuit (Lee et al., 2017), confirming that regulatory influence propagates correctly

through the estimated network. *Elane* ($\kappa = 0.304$) is a neutrophil-specific serine protease regulated by C/EBPα (Friedman, 2007). Loss-of-function mutations cause congenital neutropenia (Horwitz et al., 2013), establishing it as a critical node of granulocyte commitment.

L LIMITATIONS

While our framework provides a principled way to identify mRNA-protein regulatory dynamics from snapshot data and design re-stabilization strategies, several limitations remain. First, although the kernel conditional density estimator mitigates the high-dimension low-sample-size (HDLSS) challenge, it introduces additional computational cost and may still be sensitive to bandwidth selection.

M BROADER IMPACT

This work contributes to the emerging intersection of machine learning, biology, and medicine by providing a data-driven framework for identifying regulatory dynamics and designing early interventions in disease progression. From a biological perspective, the ability to detect pre-disease stages and suggest re-stabilization strategies has the potential to inform ultra-early treatment, shifting medical practice from reactive treatment to preventive intervention. This aligns with ongoing efforts in precision medicine, where computational tools guide targeted therapies at the molecular level. From a machine learning perspective, our study highlights how structural priors and dynamical systems theory can enhance learning in the high-dimension low-sample-size (HDLSS) regime, which frequently arises in single-cell analysis and other scientific domains. These insights may inspire future work on combining domain knowledge with statistical learning for better data efficiency and interpretability. At the same time, caution is required when interpreting computationally identified regulatory nodes as clinical intervention targets. Translational applications will require rigorous experimental validation and ethical considerations to ensure safety and effectiveness.

N EXPERIMENTS COMPUTE RESOURCES

All experiments were conducted on a MacBook equipped with an Apple M4 chip and 32GB of unified memory.