DATA-DRIVEN EDGE-RESOLVED PERTURBATION PROPAGATION IN NONLINEAR DYNAMICAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Analysis of perturbation propagation in nonlinear networks has traditionally relied on explicit governing equations. When the governing equations are unknown, resolving source-specific perturbation responses into contributions from individual interactions and organizing them into source-connected propagation structures from observed trajectories remains challenging. To address this challenge, we introduce a data-driven method for finite-time perturbation analysis using observed trajectories and known network structure. A graph-constrained neural vector field is learned from the trajectories, and its time-varying Jacobian is evaluated along an observed unperturbed reference trajectory. The resulting Jacobian drives nonautonomous variational dynamics to infer source-specific first-order perturbation responses. We then apply variation of constants to decompose the response of each affected non-source node into time-integrated incoming-edge contributions and evaluate these contributions at the node's inferred arrival time. These arrival-conditioned contributions are combined with arrival ordering to construct predecessor graphs and extract source-connected propagation paths. Across three nonlinear network systems, our method recovers state-dependent Jacobian structure, affected-node sets, propagation timing, and arrival-conditioned edge contributions. The resulting predecessor graphs closely agree with those constructed from the ground-truth Jacobian and retain most of the corresponding contribution mass.

1 INTRODUCTION

Perturbation analysis is a well-established approach for studying response and propagation in nonlinear dynamical networks. Perturbation propagation is shaped jointly by network topology, local node dynamics, and interactions (Strogatz, 2001; Barzel & Barabási, 2013; Hens et al., 2019; Ji et al., 2023). This interplay determines not only which nodes respond and when, but also which interactions and pathways mediate propagation from a source to downstream nodes (Harush & Barzel, 2017). In nonlinear systems, the effective interactions governing perturbation propagation can vary with the evolving system state even when the underlying network structure remains fixed (Kirst et al., 2016). When governing equations are available, these aspects of perturbation propagation can be analyzed directly from the dynamics. In many empirical settings, however, the underlying dynamics are not explicitly known, leaving observed trajectories as the primary source of dynamical information. This creates a central challenge: how can perturbation propagation from a specific source be inferred from observed trajectories and resolved at the level of individual interactions when the governing dynamics are unknown?

Existing work on perturbation propagation has characterized how disturbances spread through networks at different levels of resolution. At the node level, perturbative analyses have quantified source–target responses following localized disturbances (Barzel & Barabási, 2013). Spatiotemporal analyses have further characterized propagation times and revealed how network topology and local dynamics jointly shape the temporal spread of perturbations (Hens et al., 2019). Green's-function formulations have described transient sensitivities between perturbed and responding nodes (Perumal et al., 2009). Extending this perspective to nonlinear and time-dependent networks, Randi & Leifer (2021) formulate direct-edge Green's-function kernels together with connected response

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.

functions that capture both direct and indirect propagation. At a finer structural level, Harush & Barzel (2017) quantify the contribution of individual nodes, edges, and paths by selectively freezing network components and measuring the resulting change in perturbation-induced information flow. Path-based formulations further decompose steady-state source–target responses into contributions associated with multiple source–target paths (Bontorin & De Domenico, 2023). These studies provide complementary descriptions of local interactions and network-mediated propagation, but they organize edge-level information as response kernels, intervention effects, or path-based steady-state quantities. What remains less directly resolved is how the transient response of a particular target node at a finite stage of propagation can be decomposed into the contributions accumulated through its individual incoming interactions. Node-level responses characterize the timing and magnitude of a target's perturbation response, whereas the contributions of upstream interactions to that response remain unresolved. Resolving these contributions makes it possible to identify the interactions that transmit the response into each affected node and to organize them into source-connected propagation structure. An edge-resolved transient decomposition is therefore needed to link target-node responses to time-accumulated contributions from individual incoming interactions.

Resolving transient perturbation propagation at the level of individual interactions requires a data-driven representation of the state-dependent local dynamics. Observed trajectories can now support dynamical reconstruction at several levels of representation. Explicit governing relations can be sought through symbolic regression (Schmidt & Lipson, 2009; Udrescu & Tegmark, 2020) or through sparse identification of parsimonious dynamical models (Brunton et al., 2016; Champion et al., 2019; Kaheman et al., 2020). When an explicit closed-form model is not required, operator-based approaches provide reduced representations of system evolution (Williams et al., 2015; Lusch et al., 2018), while neural continuous-time models learn flexible nonlinear dynamics directly from trajectories (Chen et al., 2018; Zang & Wang, 2020). For interacting systems, graph-based formulations further incorporate relational structure into the learned dynamics (Sanchez-Gonzalez et al., 2018; 2020; Murphy et al., 2021). Beyond reproducing state evolution, a growing body of work seeks to recover the interactions that generate it, including interpretable interaction laws (Cranmer et al., 2020; Gao & Yan, 2022; Hu et al., 2025; Yu et al., 2026) and relational structure inferred directly from multivariate observations (Casadiego et al., 2017; Kipf et al., 2018). State-dependent local interactions can also be characterized more directly through local linearization, perturbation-based interaction inference, or Jacobian estimation (Deyle et al., 2016; Sarmah et al., 2022; Eisen et al., 2025). These developments make increasingly detailed local dynamical information accessible from observational data. A key challenge is therefore to connect data-driven dynamical reconstruction with source-specific perturbation propagation analysis, so that recovered local dynamics can be used to resolve how transient responses develop through network interactions over time.

To bridge the gap between recovering local dynamics from data and resolving source-specific perturbation propagation at the level of individual interactions, we develop a data-driven method that operates on observed trajectories without requiring explicit governing equations. We first learn a graph-constrained neural vector field from the observed trajectories and evaluate its time-varying Jacobian along an unperturbed reference trajectory. The resulting Jacobian defines a non-autonomous variational system for inferring source-specific first-order perturbation responses. We then apply variation of constants to decompose the response of each affected non-source node into time-integrated contributions from its individual incoming interactions, evaluated at the node's inferred arrival time. Finally, these arrival-conditioned edge contributions are combined with arrival ordering to construct predecessor graphs and extract source-connected propagation paths.

Our main contributions are as follows:

- **Data-driven source-specific perturbation analysis.** We combine learned trajectory-dependent local dynamics with variational analysis to infer source-specific first-order responses, affected nodes, and propagation timing without explicit governing equations.
- **Edge-resolved transient response decomposition.** We decompose each affected node's transient response into time-integrated contributions from its individual incoming interactions, evaluated at the node's inferred arrival time.
- **Propagation structure recovery.** We organize arrival-conditioned edge contributions by response ordering into predecessor graphs and source-connected paths.

2 METHOD

2.1 PROBLEM FORMULATION

We consider a network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \boldsymbol{A})$ with node set $\mathcal{V} = \{1, \dots, N\}$, edge set $\mathcal{E}$, and known adjacency matrix $\boldsymbol{A} \in \mathbb{R}^{N \times N}$. The scalar node states form $\boldsymbol{x}(t) = [x_1(t), \dots, x_N(t)]^\top \in \mathbb{R}^N$, governed by an unknown continuous-time nonlinear vector field:

$$\frac{\mathrm{d}\boldsymbol{x}(t)}{\mathrm{d}t} = F(\boldsymbol{x}(t); \boldsymbol{A}). \tag{1}$$

We observe $\mathcal{D} = \{\boldsymbol{x}^{(m)}(t_k)\}_{m=1,\dots,M;\, k=0,\dots,K}$, comprising M trajectories sampled at $t_0 < \cdots < t_K$. The analytical form of F is unavailable to the inference procedure.

Given an unperturbed reference trajectory $\bar{\boldsymbol{x}}(t)$ observed over the analysis horizon $[t_0, T]$, we consider an initial-condition perturbation applied to source node s at time t_0:

$$\boldsymbol{x}^{(\epsilon,s)}(t_0) = \bar{\boldsymbol{x}}(t_0) + \epsilon \boldsymbol{e}_s, \tag{2}$$

where $\boldsymbol{e}_s$ is the s-th standard basis vector and ϵ is the perturbation amplitude. The perturbed trajectory $\boldsymbol{x}^{(\epsilon,s)}(t)$ evolves under the same nonlinear dynamics as the reference trajectory over $[t_0, T]$. The corresponding nonlinear perturbation response is

$$\Delta \boldsymbol{x}^{\mathrm{NL}}_{\epsilon,s}(t) = \boldsymbol{x}^{(\epsilon,s)}(t) - \bar{\boldsymbol{x}}(t), \qquad t \in [t_0, T]. \tag{3}$$

Our task is to infer source-specific, edge-resolved perturbation propagation from observed trajectories and known network structure when the governing dynamics are unknown. The nonlinear response in Equation 3 is used only as a reference for experimental evaluation and is not provided to the inference procedure. The central problem is to transform observational data into a finite-time description of how perturbations are transmitted through individual network interactions and how these interactions form source-connected propagation paths.

2.2 FRAMEWORK OVERVIEW

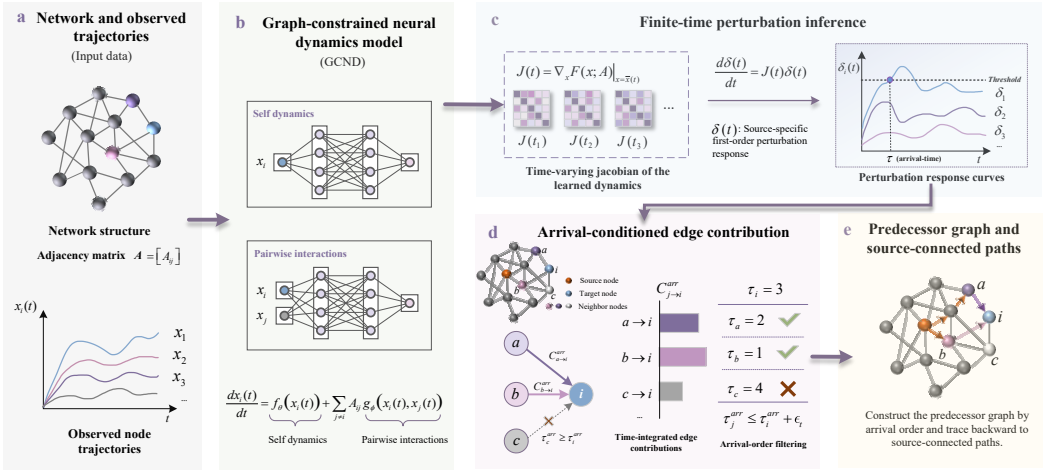

Figure 1: **Framework for data-driven edge-resolved perturbation propagation.** (a) Network and observed trajectories provide the network structure and observed node trajectories as inputs. (b) The graph-constrained neural dynamics model (GCND) learns intrinsic node dynamics and pairwise interactions. (c) Finite-time perturbation inference evaluates the learned time-varying Jacobian along the unperturbed reference trajectory and integrates the variational dynamics to obtain source-specific perturbation responses and arrival times. (d) Arrival-conditioned edge contribution uses variation of constants to obtain time-integrated incoming-edge contributions, while arrival ordering identifies temporally admissible predecessors. (e) The predecessor graph and source-connected paths organize these predecessors into a propagation structure from which source-connected paths are extracted.

Figure 1 summarizes the proposed framework, which maps observed network trajectories and known network structure to source-specific edge-resolved perturbation propagation. Given these inputs, we first learn a Graph-Constrained Neural Dynamics (GCND) model that provides a differentiable graph-constrained vector field with local sensitivities aligned to structural edges. Its time-varying Jacobian is evaluated along the unperturbed reference trajectory and used to infer source-specific first-order perturbation responses and arrival times. Variation of constants then resolves each affected non-source node's response into time-integrated contributions from its individual incoming interactions, evaluated at the node's arrival time. These arrival-conditioned edge contributions are combined with arrival ordering to construct predecessor graphs and extract source-connected propagation paths.

2.3 Graph-Constrained Neural Dynamics

GCND uses shared pairwise functions and aggregation over structural neighbors, following established principles of interaction networks and graph neural vector fields (Battaglia et al., 2016; Gilmer et al., 2017; Vasiliauskaite & Antulov-Fantulin, 2024). GCND parameterizes the vector field at node i as

$$F_{\theta,\phi,i}(\boldsymbol{x};\boldsymbol{A}) = f_\theta(x_i) + \sum_{j\neq i} A_{ij}\, g_\phi(x_i, x_j)\,, \tag{4}$$

where, f_θ represents the intrinsic dynamics shared across nodes, and g_ϕ represents the pairwise interaction function shared across structural edges. We use the convention that A_{ij} encodes a permitted or weighted interaction from node j to node i. For a symmetric structural connection, the two dynamical directions are evaluated separately because their local derivatives can differ with the states of the interacting nodes. This pairwise parameterization aligns off-diagonal local derivatives with structural interactions. For $i \neq j$, the derivative of $F_{\theta,\phi,i}$ with respect to x_j can be nonzero only if $A_{ij} \neq 0$. This edge-aligned differentiable structure makes GCND suitable for the subsequent perturbation analysis, which requires trajectory-dependent Jacobians whose off-diagonal entries correspond to individual network interactions.

Let $\mathcal{I}_{\text{train}}$ be the set of training trajectory and time index pairs. We minimize the mean-squared error against derivative targets estimated numerically from observed trajectories:

$$\mathcal{L}_{\text{dyn}} = \frac{1}{|\mathcal{I}_{\text{train}}|N} \sum_{(m,k)\in\mathcal{I}_{\text{train}}} \left\| F_{\theta,\phi}\Big(\boldsymbol{x}^{(m)}(t_k);\boldsymbol{A}\Big) - \dot{\boldsymbol{x}}^{(m)}_{\text{tar}}(t_k) \right\|_2^2. \tag{5}$$

2.4 Jacobian Along the Reference Trajectory

Differentiating GCND gives $\boldsymbol{J}(\boldsymbol{x}) = \nabla_{\boldsymbol{x}} F_{\theta,\phi}(\boldsymbol{x};\boldsymbol{A})$, with $J_{ij}(\boldsymbol{x}) = \partial F_{\theta,\phi,i}(\boldsymbol{x};\boldsymbol{A})/\partial x_j$. For $i \neq j$, this entry measures the instantaneous sensitivity of node i's state derivative to x_j and vanishes when $A_{ij} = 0$. The diagonal entry measures sensitivity to x_i, including the intrinsic term and receiver-state dependence of incoming interactions.

Perturbation inference is conditioned on the observed unperturbed reference trajectory $\bar{\boldsymbol{x}}(t)$, along which we evaluate

$$\boldsymbol{J}(t) = \nabla_{\boldsymbol{x}} F_{\theta,\phi}(\boldsymbol{x};\boldsymbol{A})\big|_{\boldsymbol{x}=\bar{\boldsymbol{x}}(t)}\,, \qquad t \in [t_0, T]. \tag{6}$$

This Jacobian tracks local first-order interactions as the reference state evolves. Numerical evaluation and interpolation are described in Appendix B. More generally, the downstream analysis applies to any differentiable dynamics model that provides reliable trajectory-dependent Jacobians with off-diagonal entries aligned to the network interactions.

2.5 Finite-Time Perturbation Inference

For a perturbation of amplitude $\epsilon \neq 0$ applied at source node s, we infer the first-order variational response $\boldsymbol{\delta}(t)$ about the unperturbed reference trajectory $\bar{\boldsymbol{x}}(t)$ by integrating the variational system with the learned Jacobian $\boldsymbol{J}(t)$ from Equation 6:

$$\frac{\mathrm{d}\boldsymbol{\delta}(t)}{\mathrm{d}t} = \boldsymbol{J}(t)\boldsymbol{\delta}(t), \qquad \boldsymbol{\delta}(t_0) = \epsilon \boldsymbol{e}_s. \tag{7}$$

Equivalently, the solution is $\boldsymbol{\delta}(t) = \boldsymbol{\Phi}(t, t_0)\epsilon \boldsymbol{e}_s$, where $\boldsymbol{\Phi}(t, t_0)$ is the state-transition matrix generated by $\boldsymbol{J}(t)$, satisfying $\partial_t \boldsymbol{\Phi}(t, t_0) = \boldsymbol{J}(t)\boldsymbol{\Phi}(t, t_0)$ and $\boldsymbol{\Phi}(t_0, t_0) = \boldsymbol{I}$ (Baake & Schlägel, 2011).

We detect node responses using $\eta_\epsilon = \alpha|\epsilon|$, where $\alpha > 0$ is a dimensionless detection ratio. Since the variational response scales linearly with ϵ, α specifies a common relative response level across perturbation amplitudes. For each non-source node i, we define the arrival time and peak response magnitude over $[t_0, T]$ as

$$\tau_i^{\mathrm{arr}} = \inf\{t \in [t_0, T] : |\delta_i(t)| \geq \eta_\epsilon\}, \qquad p_i^{\mathrm{peak}} = \max_{t\in[t_0,T]} |\delta_i(t)|. \tag{8}$$

We adopt $\inf \varnothing = +\infty$, so nodes that remain below the threshold are classified as unaffected within the analysis horizon. The inferred affected-node set is $\mathcal{A}_s = \{i \neq s : \tau_i^{\mathrm{arr}} < \infty\}$. We set the source arrival time to $\tau_s^{\mathrm{arr}} = t_0$ and exclude the source from affected-node evaluation.

2.6 Arrival-Conditioned Edge Contributions

To identify how the inferred response of a node is assembled from its incoming interactions, we write the variational dynamics of node i as

$$\frac{\mathrm{d}\delta_i(t)}{\mathrm{d}t} = J_{ii}(t)\delta_i(t) + \sum_{j\neq i} J_{ij}(t)\delta_j(t). \tag{9}$$

For scalar node states, we define the diagonal propagator

$$G_i(t, u) = \exp\left(\int_u^t J_{ii}(v)\,\mathrm{d}v\right). \tag{10}$$

Applying variation of constants gives

$$\delta_i(t) = G_i(t, t_0)\delta_i(t_0) + \sum_{j\neq i} C_{j\to i}(t), \tag{11}$$

where the time-integrated contribution entering node i through interaction $j \to i$ is

$$C_{j\to i}(t) = \int_{t_0}^t G_i(t, u)J_{ij}(u)\delta_j(u)\,\mathrm{d}u. \tag{12}$$

For every non-source node $i \neq s$, $\delta_i(t_0) = 0$, and Equation 11 therefore reduces to the exact additive decomposition $\delta_i(t) = \sum_{j\neq i} C_{j\to i}(t)$. The product $J_{ij}(u)\delta_j(u)$ is the instantaneous first-order input transmitted through interaction $j \to i$, while $G_i(t, u)$ accounts for its subsequent amplification or attenuation under the target node's diagonal local dynamics. The signed integral preserves reinforcement and cancellation over time.

Each $C_{j\to i}(t)$ quantifies the contribution to the response of node i through its final incoming interaction $j \to i$. Because $\delta_j(u)$ already contains the upstream propagation history reaching node j, this contribution aggregates those histories without explicitly enumerating all time-ordered walks.

For each affected non-source node $i \in \mathcal{A}_s$, we evaluate every incoming-interaction contribution at the node's inferred arrival time,

$$C_{j\to i}^{\mathrm{arr}} = C_{j\to i}(\tau_i^{\mathrm{arr}}). \tag{13}$$

The resulting $C_{j\to i}^{\mathrm{arr}}$ gives the additive contribution of interaction $j \to i$ when node i first crosses the detection threshold. We compute it for every structural incoming interaction before applying any predecessor selection; these arrival-conditioned contributions provide the edge-level quantities used to construct the propagation structure in the next section.

2.7 Arrival-Conditioned Predecessor Graphs and Source-Connected Paths

The Peano–Baker expansion represents multi-step perturbation propagation through time-ordered products of Jacobian matrices (Baake & Schlägel, 2011), expressing finite-time responses through temporally ordered interaction chains. To obtain an interpretable propagation structure from these

edge-level contributions, we organize them using inferred node arrival times and construct a predecessor graph. For each affected non-source node, a neighboring node with a structural edge into the target is considered a temporally admissible predecessor if its inferred response arrives no later than that of the target, up to a numerical tolerance. This criterion organizes the accumulated edge contributions according to the arrival order of detectable node responses and enables the extraction of source-connected propagation paths.

For $i \in \mathcal{A}_s$, the admissible predecessor set is

$$\mathcal{P}_i = \left\{ j \neq i : A_{ij} \neq 0,\ \tau_j^{\mathrm{arr}} \leq \tau_i^{\mathrm{arr}} + \varepsilon_t \right\}, \tag{14}$$

where $\varepsilon_t \geq 0$ is a numerical tolerance for arrival-time comparisons. Unaffected nodes have $\tau_j^{\mathrm{arr}} = +\infty$ and are excluded.

The arrival-conditioned predecessor graph is $\mathcal{G}_s^{\mathrm{arr}} = (\mathcal{V}_s^{\mathrm{arr}}, \mathcal{E}_s^{\mathrm{arr}})$, with $\mathcal{V}_s^{\mathrm{arr}} = \{s\} \cup \mathcal{A}_s$ and

$$\mathcal{E}_s^{\mathrm{arr}} = \left\{ j \to i : i \in \mathcal{A}_s,\ j \in \mathcal{P}_i,\ \left| C_{j\to i}^{\mathrm{arr}} \right| > 0 \right\}. \tag{15}$$

Each edge retains its signed arrival-conditioned weight $C_{j\to i}^{\mathrm{arr}}$.

We traverse the predecessor graph backward from affected non-source nodes to extract simple paths connecting them to the source node s. Traversal rules and search limits are detailed in Appendix C.

3 Experiments

3.1 Experimental Setup

Dynamical Systems and Complex Networks. We use Mutualistic–SW, Neuronal–Brain, and SIS–Gamma as the three primary benchmark systems in the main text (Hens et al., 2019; Zang & Wang, 2020). These pairings span different dynamical families, topology classes, network sizes, and edge-weighting schemes. The Small-world network is synthetic and unweighted, the Brain network is an empirical weighted connectome, and Gamma300 is an unweighted scale-free network with 300 nodes and 589 undirected edges. Table 1 summarizes the dynamics and network configurations of the three primary benchmark systems. To further evaluate the framework across different dynamics–topology combinations, we report several additional pairings in Appendix F.

Table 1: **Network dynamics and complex networks.**

System	Node dynamics	Complex Networks	Nodes	Edges	Weighted
Mutualistic–SW	$\frac{\mathrm{d}x_i}{\mathrm{d}t} = b + x_i\left(1 - \frac{x_i}{k}\right)\left(\frac{x_i}{c} - 1\right) + \sum_{j=1}^{N} A_{ij} \frac{x_i x_j}{d + e x_i + h x_j}$	Small-world	400	1,186	No
Neuronal–Brain	$\frac{\mathrm{d}x_i}{\mathrm{d}t} = -B x_i + D \tanh(x_i) + \sum_{j=1}^{N} A_{ij} \tanh(x_j)$	Brain	998	17,865	Yes
SIS–Gamma	$\frac{\mathrm{d}x_i}{\mathrm{d}t} = -x_i + \sum_{j=1}^{N} A_{ij}(1 - x_i) x_j$	Gamma300	300	589	No

Model Training and Local-Dynamics Evaluation. GCND was trained on trajectory-derived derivative targets with validation-based model selection. Governing equations supplied synthetic observations and post-training evaluation references. Because the subsequent variational analysis is governed directly by the local Jacobian, trajectory accuracy alone does not establish whether the learned dynamics recover the perturbation sensitivities required for propagation inference (Park et al., 2024). Rollout evaluation used all 10 test trajectories; Jacobian evaluation used 20 states per trajectory, covering diagonal entries and off-diagonal structural-edge entries, with J_{ij} and J_{ji} evaluated separately. Additional implementation and training details are provided in Appendix A.

Perturbation Response and Edge Contribution Evaluation. We used one fixed test trajectory per system and 20 sources covering a broad range of source connectivity values. Perturbations had amplitude $\epsilon = r\sigma_x$, with $r \in \{0.1\%, 0.3\%, 1\%, 3\%, 5\%, 10\%\}$, where σ_x is the across-node standard deviation at perturbation onset. The full sweep assessed amplitude dependence; detailed edge contribution and propagation-structure analyses used $r = 1\%$. Nonlinear responses came from paired unperturbed and perturbed rollouts. Analytical and learned time-varying Jacobians generated variational responses $\boldsymbol{\delta}^{\star}(t)$ and $\widehat{\boldsymbol{\delta}}(t)$, respectively, along the same unperturbed reference

trajectory. Arrival detection used $\eta_\epsilon = \alpha|\epsilon|$ with $\alpha = 0.01$. Ground-truth edge contributions used the analytical Jacobian and its variational response in the same variation-of-constants procedure as the learned model. Each contribution was evaluated at its respective variational arrival time. Direct comparisons included all nonself structural incoming edges of non-source targets affected by both responses, before predecessor selection. Appendix D gives source-selection rules, metrics, valid-source counts, and statistical procedures.

3.2 Fidelity of Learned Dynamics for Perturbation Analysis

Before evaluating perturbation propagation, we assess held-out trajectory fidelity and trajectory-dependent local sensitivities. Median time-averaged relative L_2 rollout errors were 3.24%, 7.42%, and 0.272% for Mutualistic–SW, Neuronal–Brain, and SIS–Gamma, respectively (Figure 2). Learned and analytical Jacobians also agree closely, with median Pearson correlations above 0.9980 on structural edges and above 0.9996 on the diagonal, while relative Frobenius errors remain below 3.97% and 3.22%, respectively. These results indicate that the learned vector fields reproduce held-out trajectories and preserve the edge-aligned local sensitivities required for perturbation analysis.

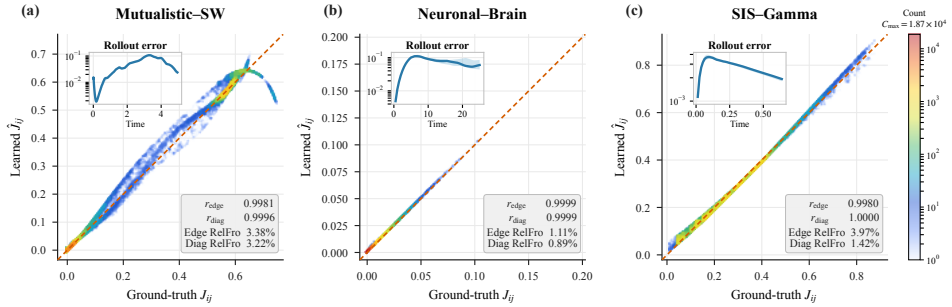

Figure 2: **Learned dynamics and Jacobian fidelity across three networked systems.** Scatter plots compare learned and analytical off-diagonal Jacobian entries over 200 states per system; dashed lines denote identity. Insets show median relative L_2 rollout errors, and boxes summarize Jacobian correlations and relative Frobenius errors. Each panel displays 50,000 edge–state pairs.

3.3 Finite-Time Perturbation Inference

Accurate local Jacobians must remain informative after temporal integration to support finite-time perturbation inference. We therefore evaluated the responses generated by the learned time-varying Jacobians along the reference trajectory. Figure 3 compares three response quantities: the nonlinear perturbation response, the variational response generated by the analytical Jacobian, and the variational response generated by the learned Jacobian. Their pairwise discrepancies characterize complementary aspects of the inference process. The difference between the analytical-Jacobian variational response and the nonlinear perturbation response reflects finite-amplitude departure from the first-order approximation, while the difference between the learned-Jacobian variational response and the analytical-Jacobian variational response reflects the discrepancy associated with Jacobian learning. The remaining comparison, between the learned-Jacobian variational response and the nonlinear perturbation response, measures the end-to-end inference discrepancy.

These discrepancies exhibit distinct amplitude-dependent patterns across the three systems (Figure 3(a–c)). In Mutualistic–SW, the analytical-Jacobian variational response closely approximates the nonlinear response at the smallest perturbation amplitude, while their discrepancy increases markedly as r grows. By comparison, the learned-Jacobian variational response maintains a relatively stable discrepancy from the analytical-Jacobian variational response across the amplitude sweep, and the growth of the end-to-end discrepancy tracks the increasing departure of the nonlinear response from the first-order variational approximation. Neuronal–Brain remains comparatively stable over smaller perturbations before showing a moderate increase at larger r, whereas SIS–Gamma maintains low and relatively stable response discrepancies throughout the evaluated range. This behavior is consistent with the local first-order approximation underlying variational dynamics. Sufficiently small perturbations are governed to leading order by the Jacobian along the

reference trajectory (Khalil, 2002; Yamane & Yoden, 2001). As the perturbation grows, higher-order nonlinear effects become increasingly relevant (Khalil, 2002). The amplitude sweep therefore motivates $r = 1\%$ as a common small-amplitude operating point for the detailed edge-contribution and propagation-structure analyses that follow.

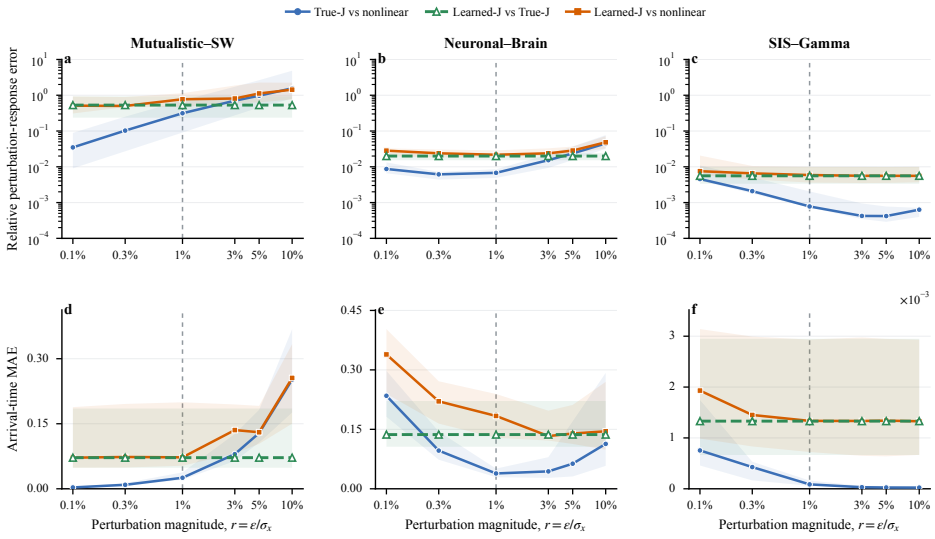

Figure 3: **Finite-time perturbation inference across perturbation amplitudes.** Columns show Mutualistic–SW, Neuronal–Brain, and SIS–Gamma; rows show time-weighted relative response error and arrival time MAE. Lines and shading denote the median and IQR across 20 sources.

Arrival-time behavior provides a complementary view of node-level propagation fidelity (Figure 3(d–f)). The arrival-time discrepancy between the learned- and analytical-Jacobian variational responses remains nearly invariant with r across all three systems, consistent with the linear scaling of the variational response and the amplitude-scaled detection threshold. At $r = 1\%$, the learned-Jacobian responses closely reproduce the nonlinear affected-node sets and propagation order across all three systems (Table 2). Peak-response rankings are also strongly preserved for Mutualistic–SW and Neuronal–Brain, while SIS–Gamma shows lower agreement.

Table 2: **Finite-time perturbation agreement at $r = 1\%$.** Values are source-level medians comparing learned-Jacobian variational responses with nonlinear responses. Confidence intervals and valid-source counts are reported in Appendix D.

System	Affected-node $F_1 \uparrow$	Arrival-time MAE $\downarrow$	Arrival-order $\rho \uparrow$	Peak-response $\rho \uparrow$
Mutualistic–SW	0.949	0.0725	0.989	0.997
Neuronal–Brain	0.997	0.1838	1.000	0.997
SIS–Gamma	1.000	0.0013	1.000	0.792

To assess whether retaining the trajectory dependence of $J(t)$ is necessary, we compare the learned time-varying Jacobian with counterparts frozen at perturbation onset and near steady state. Both fixed-Jacobian variants substantially degrade response recovery, particularly in Mutualistic–SW and Neuronal–Brain, and analytical fixed Jacobians show the same pattern. Appendix E reports the full amplitude sweeps and corresponding propagation-structure comparisons.

3.4 Edge-Resolved Perturbation Propagation

To evaluate edge-resolved perturbation propagation, we assess arrival-conditioned edge contributions, their organization into source-connected propagation structures, and their functional relevance.

Before predecessor selection, learned arrival-conditioned edge contributions closely match their analytical-Jacobian references, with median signed Pearson correlations of 0.9961, 0.9999, and 1.0000 and relative L_2 errors of 8.00%, 1.16%, and 0.123% for Mutualistic–SW, Neuronal–Brain, and SIS–Gamma, respectively.

The resulting predecessor graphs also agree closely with the analytical references. Median predecessor-edge Jaccard similarities are 0.865, 0.991, and 1.000; median node-level exact predecessor-set agreement rates are 0.810, 0.956, and 1.000; and median contribution-mass recalls are 0.9926, 0.9946, and 1.0000. Figure 4 illustrates this agreement for source 20 in Mutualistic–SW, where the learned predecessor graph reproduces the main propagation branches of the analytical reference from intermediate to late propagation. Source-connected paths likewise show high exact path match rates and path-edge Jaccard similarities among targets for which both analyses resolve a path. Detailed path diagnostics and fixed-Jacobian comparisons are reported in Appendix E.3.

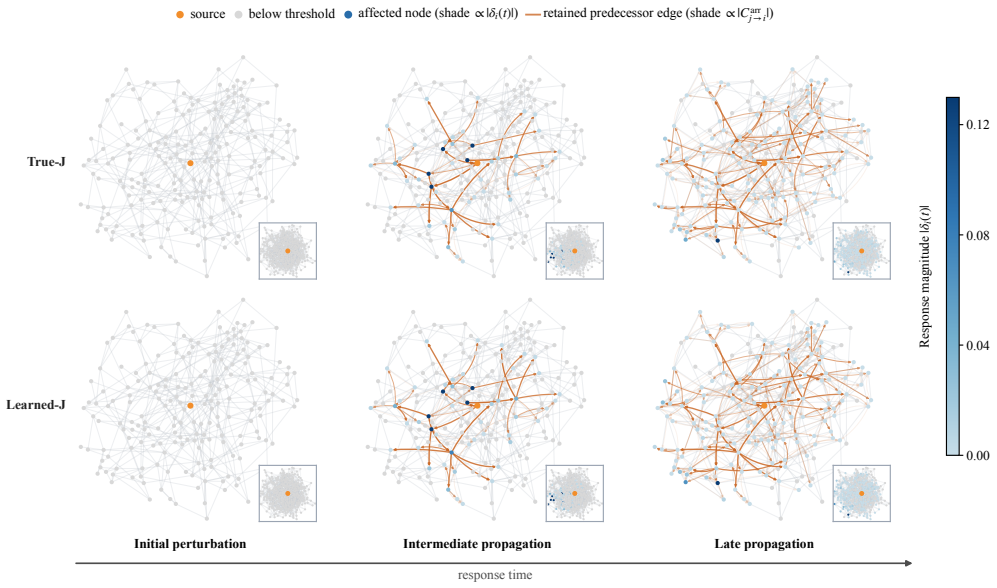

Figure 4: **Edge-resolved propagation pathways.** Predecessor graphs for source 20 in Mutualistic–SW at $r = 1\%$ are shown for the analytical and learned time-varying Jacobians at initial, intermediate, and late propagation stages. Node and predecessor-edge shades encode $|\delta_i(t)|$ and $|C^{\mathrm{arr}}_{j\to i}|$, respectively; insets show the complete network.

We further test functional relevance by removing the learned predecessor edges in each target's source-connected ancestry from the ground-truth nonlinear system and comparing the response with three random removals matched in edge count and approximate receiver indegree. Across the three systems, median targeted-minus-random differences in peak-response suppression were 0.986, 1.396, and 1.000; targeted removal prevented arrival in 83.3%, 95.0%, and 100% of conditions, compared with 0%, 1.7%, and 0% under random removal. Holm-adjusted tests show significant targeted-minus-random differences in peak suppression and arrival delay across all three systems, with full statistical results reported in Appendix D.

4 Conclusion

We presented a data-driven method for inferring source-specific edge-resolved perturbation propagation from observed network trajectories when the governing dynamics are unknown. Across three nonlinear network systems, the learned time-varying Jacobians recover perturbation responses, propagation timing, and arrival-conditioned incoming-edge contributions with high fidelity. These contributions yield predecessor graphs and source-connected paths that closely agree with analytical-Jacobian references, providing an edge-resolved description of perturbation propagation from observed trajectories.

REFERENCES

Michael Baake and Ulrike Schlägel. The Peano–Baker series. *Proceedings of the Steklov Institute of Mathematics*, 275(1):155–159, 2011. doi: 10.1134/S0081543811080098.

Baruch Barzel and Albert-László Barabási. Universality in network dynamics. *Nature Physics*, 9 (10):673–681, 2013. doi: 10.1038/nphys2741.

Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, and Koray Kavukcuoglu. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems*, volume 29, 2016.

Sebastiano Bontorin and Manlio De Domenico. Multi pathways temporal distance unravels the hidden geometry of network-driven processes. *Communications Physics*, 6:129, 2023. doi: 10.1038/s42005-023-01204-1.

Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.

Jose Casadiego, Mor Nitzan, Sarah Hallerberg, and Marc Timme. Model-free inference of direct network interactions from nonlinear collective dynamics. *Nature Communications*, 8(1):2192, 2017. doi: 10.1038/s41467-017-02288-4.

Kathleen Champion, Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116 (45):22445–22451, 2019. doi: 10.1073/pnas.1906995116.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Miles Cranmer, Alvaro Sanchez Gonzalez, Peter Battaglia, Rui Xu, Kyle Cranmer, David Spergel, and Shirley Ho. Discovering symbolic models from deep learning with inductive biases. In *Advances in Neural Information Processing Systems*, volume 33, pp. 17429–17442, 2020.

Ethan R. Deyle, Robert M. May, Stephan B. Munch, and George Sugihara. Tracking and forecasting ecosystem interactions in real time. *Proceedings of the Royal Society B: Biological Sciences*, 283 (1822):20152258, 2016. doi: 10.1098/rspb.2015.2258.

Adam J. Eisen, Mitchell Ostrow, Sarthak Chandra, Leo Kozachkov, Earl K. Miller, and Ila R. Fiete. Characterizing control between interacting subsystems with deep Jacobian estimation. In *Advances in Neural Information Processing Systems*, volume 38, pp. 51361–51406, 2025. doi: 10.52202/085713-1711.

Ting-Ting Gao and Gang Yan. Autonomous inference of complex network dynamics from incomplete and noisy data. *Nature Computational Science*, 2:160–168, 2022. doi: 10.1038/s43588-022-00217-0.

Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1263–1272. PMLR, 2017.

Uzi Harush and Baruch Barzel. Dynamic patterns of information flow in complex networks. *Nature Communications*, 8(1):2181, 2017. doi: 10.1038/s41467-017-01916-3.

Chittaranjan Hens, Uzi Harush, Simi Haber, Reuven Cohen, and Baruch Barzel. Spatiotemporal signal propagation in complex networks. *Nature Physics*, 15(4):403–412, 2019. doi: 10.1038/s41567-018-0409-0.

Jiao Hu, Jiaxu Cui, and Bo Yang. Learning interpretable network dynamics via universal neural symbolic regression. *Nature Communications*, 16(1):6226, 2025. doi: 10.1038/s41467-025-61575-7.

Peng Ji, Jiachen Ye, Yu Mu, Wei Lin, Yang Tian, Chittaranjan Hens, Matjaž Perc, Yang Tang, Jie Sun, and Jürgen Kurths. Signal propagation in complex networks. *Physics Reports*, 1017:1–96, 2023. doi: 10.1016/j.physrep.2023.03.005.

Kadierdan Kaheman, J. Nathan Kutz, and Steven L. Brunton. SINDy-PI: A robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 476(2242):20200279, 2020. doi: 10.1098/rspa.2020.0279.

Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River, NJ, third edition, 2002. ISBN 0-13-067389-7.

Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2688–2697. PMLR, 2018.

Christoph Kirst, Marc Timme, and Demian Battaglia. Dynamic information routing in complex networks. *Nature Communications*, 7(1):11061, 2016. doi: 10.1038/ncomms11061.

Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1):4950, 2018. doi: 10.1038/s41467-018-07210-0.

Charles Murphy, Edward Laurence, and Antoine Allard. Deep learning of contagion dynamics on complex networks. *Nature Communications*, 12(1):4720, 2021. doi: 10.1038/s41467-021-24732-2.

Jeongjin Park, Nicole Tianjiao Yang, and Nisha Chandramoorthy. When are dynamical systems learned from time series data statistically accurate? In *Advances in Neural Information Processing Systems*, volume 37, pp. 43975–44008, 2024. doi: 10.52202/079017-1396.

Thanneer M. Perumal, Yan Wu, and Rudiyanto Gunawan. Dynamical analysis of cellular networks based on the Green's function matrix. *Journal of Theoretical Biology*, 261(2):248–259, 2009. doi: 10.1016/j.jtbi.2009.07.037.

Francesco Randi and Andrew M. Leifer. Nonequilibrium Green's functions for functional connectivity in the brain. *Physical Review Letters*, 126(11):118102, 2021. doi: 10.1103/PhysRevLett.126.118102.

Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin Riedmiller, Raia Hadsell, and Peter Battaglia. Graph networks as learnable physics engines for inference and control. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4470–4479. PMLR, 2018.

Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 8459–8468. PMLR, 2020.

Deepraj Sarmah, Gregory R. Smith, Mehdi Bouhaddou, Alan D. Stern, James Erskine, and Marc R. Birtwistle. Network inference from perturbation time course data. *npj Systems Biology and Applications*, 8(1):42, 2022. doi: 10.1038/s41540-022-00253-6.

Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85, 2009. doi: 10.1126/science.1165893.

Steven H. Strogatz. Exploring complex networks. *Nature*, 410(6825):268–276, 2001. doi: 10.1038/35065725.

Silviu-Marian Udrescu and Max Tegmark. AI Feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16):eaay2631, 2020. doi: 10.1126/sciadv.aay2631.

Vaiva Vasiliauskaite and Nino Antulov-Fantulin. Generalization of neural network models for complex network dynamics. *Communications Physics*, 7(1):348, 2024. doi: 10.1038/s42005-024-01837-w.

Matthew O. Williams, Ioannis G. Kevrekidis, and Clarence W. Rowley. A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015. doi: 10.1007/s00332-015-9258-5.

Shozo Yamane and Shigeo Yoden. Finite-time evolution of small perturbations superposed on a chaotic solution: Experiment with an idealized barotropic model. *Journal of the Atmospheric Sciences*, 58(9):1066–1078, 2001. doi: 10.1175/1520-0469(2001)058⟨1066:FTEOSP⟩2.0.CO;2.

Zihan Yu, Jingtao Ding, and Yong Li. Discovering network dynamics with neural symbolic regression. *Nature Computational Science*, 6:156–168, 2026. doi: 10.1038/s43588-025-00893-8.

Chengxi Zang and Fei Wang. Neural dynamics on complex networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 892–902, 2020. doi: 10.1145/3394486.3403132.

A Model Implementation and Training Protocol

For all three primary benchmark systems, GCND follows the graph-constrained decomposition in Equation 4. The intrinsic function f_θ and pairwise interaction function g_ϕ are implemented as multilayer perceptrons shared across nodes and structural edges, respectively. Both networks operate on scalar node states and use SiLU activations. Each pairwise message is weighted by the corresponding entry of $\boldsymbol{A}$, and incoming messages are aggregated by summation. The prescribed network structure remains fixed throughout training. Table 3 summarizes the model architecture and training configuration.

Table 3: **GCND architecture and training settings for the three primary benchmark systems.** The trajectory split gives the numbers of training, validation, and test trajectories.

Setting	Mutualistic–SW	Neuronal–Brain	SIS–Gamma
Intrinsic network f_θ	$1 \to 64 \to 64 \to 1$	$1 \to 64 \to 1$	$1 \to 64 \to 1$
Interaction network g_ϕ	$2 \to 256 \to 256 \to 1$	$2 \to 128 \to 1$	$2 \to 128 \to 1$
Hidden activation	SiLU	SiLU	SiLU
Optimizer	Adam	Adam	Adam
Learning rate	1×10^{-3}	3×10^{-3}	1×10^{-2}
Batch size	128	64	64
Maximum training epochs	300	100	100
Gradient clipping norm	10	10	10
Training objective	Derivative MSE	Derivative MSE	Derivative MSE
Trajectory split (train/validation/test)	200/10/10	200/10/10	200/10/10

Trajectory identities were assigned to the training, validation, and test sets before constructing state–derivative pairs. The allocation was performed at the trajectory level, with all sampled states and time points from a trajectory retained in the same split. Each system used 200 trajectories for parameter optimization, 10 trajectories for checkpoint selection, and 10 trajectories for held-out evaluation. The reference trajectory used in the subsequent perturbation analyses was designated as a test trajectory before model training. The same split was retained throughout the rollout, Jacobian, perturbation-response, edge-attribution, and propagation experiments.

Derivative targets were estimated from the sampled trajectories by finite differences. Adam optimization minimized the derivative reconstruction loss in Equation 5 with the fixed learning rates listed in Table 3. Gradient clipping used a maximum norm of 10. For each system, the checkpoint with the lowest validation derivative mean-squared error was selected. Rollout and direct Jacobian evaluations used all 10 test trajectories, while the perturbation-response, edge-attribution, and propagation analyses used the designated reference trajectory from the test split. Fixed random seeds governed trajectory generation, split assignment, and model initialization.

B Jacobian Computation and Variational Integration

B.1 Jacobian Computation

The pairwise structure of GCND allows the Jacobian to be assembled from derivatives of the intrinsic and interaction functions. For the vector field defined in Section 2.3, its entries are

$$J_{ij}(\boldsymbol{x}) = A_{ij}\, \partial_2 g_\phi(x_i, x_j), \qquad i \neq j,$$
$$J_{ii}(\boldsymbol{x}) = f'_\theta(x_i) + \sum_{j \neq i} A_{ij}\, \partial_1 g_\phi(x_i, x_j).$$

Here, ∂_1 and ∂_2 denote derivatives with respect to the receiver and sender states, respectively. We compute these derivatives using automatic differentiation, with the trained parameters and network structure held fixed. The diagonal entry combines the intrinsic derivative with the receiver-state dependence of all incoming interactions. For a symmetric structural connection, the two directional entries are evaluated separately at the corresponding node states.

Each model uses a single fixed scalar scale shared across nodes. Input states are divided by this scale, and predicted state derivatives are multiplied by the same scale. The two factors cancel under the chain rule, so the computed derivatives give the Jacobian with respect to the original state variables.

B.2 Evaluation Along the Reference Trajectory

Perturbation inference is conditioned on the observed unperturbed reference trajectory over the analysis horizon. We evaluate the learned Jacobian at each sampled reference state, obtaining $\boldsymbol{J}_k = \nabla_{\boldsymbol{x}} F_{\theta,\phi}(\bar{\boldsymbol{x}}(t_k); \boldsymbol{A})$. At intermediate times, we use piecewise-linear interpolation:

$$\boldsymbol{J}_{\text{interp}}(t) = (1-\lambda)\boldsymbol{J}_k + \lambda \boldsymbol{J}_{k+1}, \qquad \lambda = \frac{t - t_k}{t_{k+1} - t_k}, \qquad t \in [t_k, t_{k+1}].$$

Interpolation is applied directly to the Jacobian entries. The interval-specific weight accommodates nonuniform observation times. It also preserves structural zeros and reproduces the evaluated Jacobian at each observation time. This interpolant provides the numerical approximation to the trajectory-dependent Jacobian used in the variational dynamics.

B.3 Variational Integration and Numerical Implementation

We store the diagonal entries and the off-diagonal entries associated with directed structural edges separately. At each integration stage, the Jacobian-vector product is computed by combining the diagonal contribution with the incoming-edge sum for each node. Storage at each sampled time therefore scales with the number of nodes and directed structural edges.

The variational dynamics are integrated using fourth-order Runge–Kutta, with the interpolated diagonal and off-diagonal entries evaluated at each integration stage. For a fixed reference trajectory, the variational system is homogeneous and linear in the perturbation response. We therefore integrate the response with initial condition $\boldsymbol{e}_s$ once for each source and multiply the resulting response by ϵ for each perturbation amplitude. The stored Jacobian sequence is shared across sources and amplitudes along the same reference trajectory. For experimental comparison, analytical Jacobians are evaluated along that reference trajectory and used with the same interpolation and variational integration procedure.

C Predecessor Graph Construction and Path Traversal

C.1 Nonzero Predecessors and Numerical Tolerance

For an affected node i other than the source, let

$$\mathcal{P}_i^+ = \left\{ j \in \mathcal{P}_i : \left| C_{j\to i}^{\text{arr}} \right| > 0 \right\}, \qquad M_i = \sum_{j \in \mathcal{P}_i^+} \left| C_{j\to i}^{\text{arr}} \right|. \tag{16}$$

We use $\varepsilon_C = 10^{-15}$ as the numerical tolerance for M_i. The retained predecessor set is

$$\mathcal{R}_i = \begin{cases} \mathcal{P}_i^+, & M_i > \varepsilon_C, \\ \varnothing, & M_i \le \varepsilon_C. \end{cases} \tag{17}$$

Nodes with no retained predecessors are marked unresolved. For source s, the retained edge set is

$$\mathcal{E}_{s,\text{ret}}^{\text{arr}} = \{ (j \to i) : i \in \mathcal{A}_s,\ j \in \mathcal{R}_i \}. \tag{18}$$

The absolute values prevent cancellation between positive and negative contributions when assessing whether M_i exceeds the numerical tolerance. Each retained edge carries its signed arrival-conditioned contribution.

C.2 Arrival-Order Tolerance and Path Traversal

For deterministic traversal, predecessors are ordered by decreasing $|C_{j\to i}^{\text{arr}}|$, with equal magnitudes ordered by the fixed input edge index.

We use $\varepsilon_t = 10^{-12}$ in Equation 14. Consequently, a predecessor may have an estimated arrival time up to ε_t later than its target, and arrival-ordered refers to ordering up to this numerical tolerance. Equal or numerically indistinguishable arrival times can create cycles in the retained predecessor graph. During backward traversal, each branch therefore maintains a set of visited nodes. Rejecting predecessor choices that revisit a node ensures that every reported path is simple. A branch terminates when it reaches the source, has no unvisited admissible predecessor, or reaches the depth limit.

The multi-predecessor path diagnostics in the main analysis report at most 20 paths per target with a beam width of 50 and a maximum traversal depth of 32 for all three systems. These bounds apply to the path search; $C^{\mathrm{arr}}_{j\to i}$, the retained predecessor sets, $\mathcal{E}^{\mathrm{arr}}_{s,\mathrm{ret}}$, and the predecessor-graph metrics use the complete retained graph.

D Statistical Protocol and Uncertainty Estimation

D.1 Evaluation Metrics

For a held-out trajectory m, the relative rollout error over the full state is

$$e^{(m)}_{\mathrm{roll}}(t_k) = \frac{\left\|\widehat{\boldsymbol{x}}^{(m)}(t_k) - \boldsymbol{x}^{(m)}(t_k)\right\|_2}{\left\|\boldsymbol{x}^{(m)}(t_k)\right\|_2 + \varepsilon_{\mathrm{roll}}}, \qquad \bar{e}^{(m)}_{\mathrm{roll}} = \frac{1}{K}\sum_{k=1}^{K} e^{(m)}_{\mathrm{roll}}(t_k), \tag{19}$$

where the shared initial state t_0 is omitted from the average. For direct Jacobian evaluation, let $\boldsymbol{j}_{\mathcal{E}}(t_k)$ and $\widehat{\boldsymbol{j}}_{\mathcal{E}}(t_k)$ collect the ground-truth and learned entries $J^{\star}_{ij}(t_k)$ and $\widehat{J}_{ij}(t_k)$, respectively, satisfying $i \neq j$ and $A_{ij} \neq 0$. Directional entries are retained separately. The metrics at each state are

$$r_{\mathrm{edge}}(t_k) = \mathrm{corr_P}\left(\boldsymbol{j}_{\mathcal{E}}(t_k), \widehat{\boldsymbol{j}}_{\mathcal{E}}(t_k)\right), \qquad E_{J,\mathcal{E}}(t_k) = \frac{\left\|\widehat{\boldsymbol{j}}_{\mathcal{E}}(t_k) - \boldsymbol{j}_{\mathcal{E}}(t_k)\right\|_2}{\left\|\boldsymbol{j}_{\mathcal{E}}(t_k)\right\|_2 + \varepsilon_J}, \tag{20}$$

where $\mathrm{corr_P}$ denotes Pearson correlation and the ε terms are numerical stabilizers.

The response metric in Figure 3 is the time-weighted relative error of a candidate response $\boldsymbol{a}$ with respect to a reference response $\boldsymbol{b}$:

$$E_{\mathrm{resp}}(\boldsymbol{a} \mid \boldsymbol{b}) = \frac{\left[\sum_{k=0}^{K} w_k \left\|\boldsymbol{a}(t_k) - \boldsymbol{b}(t_k)\right\|_2^2\right]^{1/2}}{\left[\sum_{k=0}^{K} w_k \left\|\boldsymbol{b}(t_k)\right\|_2^2\right]^{1/2} + \varepsilon_{\mathrm{num}}}, \tag{21}$$

with trapezoidal weights

$$w_0 = \frac{t_1 - t_0}{2}, \qquad w_K = \frac{t_K - t_{K-1}}{2}, \qquad w_k = \frac{t_{k+1} - t_{k-1}}{2} \quad (1 \le k < K), \tag{22}$$

and $\varepsilon_{\mathrm{num}} = 10^{-12}$. Each of Figure 3(a–c) shows three comparisons: $E_{\mathrm{resp}}(\boldsymbol{\delta}^{\star} \mid \Delta\boldsymbol{x}^{\mathrm{NL}}_{\epsilon,s})$, $E_{\mathrm{resp}}(\widehat{\boldsymbol{\delta}} \mid \boldsymbol{\delta}^{\star})$, and $E_{\mathrm{resp}}(\widehat{\boldsymbol{\delta}} \mid \Delta\boldsymbol{x}^{\mathrm{NL}}_{\epsilon,s})$. The same metric is used in Appendix E.

For a response $\boldsymbol{a}$, define $\mathcal{A}(\boldsymbol{a}) = \{i \neq s : \tau^{\mathrm{arr}}_i(\boldsymbol{a}) < \infty\}$ and $p_i(\boldsymbol{a}) = \max_{0\le k\le K}|a_i(t_k)|$. Precision, recall, and F_1 for affected nodes are

$$P = \frac{|\mathcal{A}(\boldsymbol{a}) \cap \mathcal{A}(\boldsymbol{b})|}{\max(|\mathcal{A}(\boldsymbol{a})|, 1)}, \qquad R = \frac{|\mathcal{A}(\boldsymbol{a}) \cap \mathcal{A}(\boldsymbol{b})|}{\max(|\mathcal{A}(\boldsymbol{b})|, 1)}, \qquad F_1 = \frac{2PR}{\max(P + R, 10^{-30})}, \tag{23}$$

so that an empty predicted or reference set yields $F_1 = 0$. Arrival time MAE is

$$E_{\mathrm{arr}}(\boldsymbol{a}, \boldsymbol{b}) = \frac{1}{|\mathcal{S}(\boldsymbol{a}, \boldsymbol{b})|} \sum_{i \in \mathcal{S}(\boldsymbol{a},\boldsymbol{b})} \left|\tau^{\mathrm{arr}}_i(\boldsymbol{a}) - \tau^{\mathrm{arr}}_i(\boldsymbol{b})\right|, \qquad \mathcal{S}(\boldsymbol{a}, \boldsymbol{b}) = \mathcal{A}(\boldsymbol{a}) \cap \mathcal{A}(\boldsymbol{b}), \tag{24}$$

and Spearman correlation of arrival order uses the same nodes. Spearman correlation of peak responses compares $\{p_i(\boldsymbol{a})\}_{i\neq s}$ and $\{p_i(\boldsymbol{b})\}_{i\neq s}$ over all nodes other than the source.

Direct edge-contribution evaluation uses $r = 1\%$ and $\eta_\epsilon = \alpha|\epsilon|$ with $\alpha = 0.01$ for both variational responses. For each source, let $\widehat{\boldsymbol{c}}_s$ and $\boldsymbol{c}^{\star}_s$ collect the learned and analytical-Jacobian reference

contributions over all nonself structural incoming edges of targets affected by both responses. The entries are $\widehat{C}_{j\to i}(\widehat{\tau}_i^{\mathrm{arr}})$ and $C^{\star}_{j\to i}(\tau_i^{\star,\mathrm{arr}})$, respectively: each is evaluated at the arrival time of its own variational response. The edge set is shared between the two vectors and includes zero contributions before predecessor graph construction. We compute

$$r_{C,s} = \mathrm{corr}_{\mathrm{P}}(\widehat{\boldsymbol{c}}_s, \boldsymbol{c}_s^{\star}), \qquad E_{C,s} = \frac{\|\widehat{\boldsymbol{c}}_s - \boldsymbol{c}_s^{\star}\|_2}{\|\boldsymbol{c}_s^{\star}\|_2}. \tag{25}$$

Pearson correlation uses signed values and requires at least three finite pairs with nonconstant vectors. Vectors satisfying $\|\boldsymbol{v} - \bar{v}\mathbf{1}\|_2 < 10^{-13}|\bar{v}|$ are treated as numerically constant. The relative error requires $\|\boldsymbol{c}_s^{\star}\|_2 > 10^{-12}$ and uses no denominator stabilizer. Both metrics are computed once per source and summarized over 20 valid sources per system in Table 4.

For predecessor-graph evaluation, let $\mathcal{E}_s^{\star,\mathrm{arr}}$ and $\widehat{\mathcal{E}}_s^{\mathrm{arr}}$ denote the ground-truth and learned retained edge sets for source s, and let $\mathcal{R}_i^{\star}$ and $\widehat{\mathcal{R}}_i$ denote their predecessor sets at node i. Define $\mathcal{T}_s = \{i : \mathcal{R}_i^{\star} \cup \widehat{\mathcal{R}}_i \neq \varnothing\}$. The three metrics for set agreement are

$$J_{\mathrm{edge},s} = \frac{|\mathcal{E}_s^{\star,\mathrm{arr}} \cap \widehat{\mathcal{E}}_s^{\mathrm{arr}}|}{|\mathcal{E}_s^{\star,\mathrm{arr}} \cup \widehat{\mathcal{E}}_s^{\mathrm{arr}}|}, \tag{26}$$

$$A_{\mathrm{exact},s} = \frac{1}{|\mathcal{T}_s|} \sum_{i\in\mathcal{T}_s} \mathbf{1}\left[\mathcal{R}_i^{\star} = \widehat{\mathcal{R}}_i\right], \tag{27}$$

$$J_{\mathrm{node},s} = \frac{1}{|\mathcal{T}_s|} \sum_{i\in\mathcal{T}_s} \frac{|\mathcal{R}_i^{\star} \cap \widehat{\mathcal{R}}_i|}{|\mathcal{R}_i^{\star} \cup \widehat{\mathcal{R}}_i|}. \tag{28}$$

Recall of ground-truth contribution mass is

$$R_s^{\mathrm{mass}} = \frac{\sum_i \sum_{j\in\widehat{\mathcal{R}}_i\cap\mathcal{P}_i^{\star}} \left|C^{\star,\mathrm{arr}}_{j\to i}\right|}{\sum_i \sum_{j\in\mathcal{P}_i^{\star}} \left|C^{\star,\mathrm{arr}}_{j\to i}\right|}, \tag{29}$$

where $\mathcal{P}_i^{\star}$ is the ground-truth temporally admissible incoming set. Computed for each source, this ratio assigns zero recovered mass to an affected target for which the learned analysis retains no predecessor.

D.2 Evaluation Units, Source Selection, and Aggregation

All statistics were computed separately for each system and, where applicable, each perturbation amplitude. For held-out rollout evaluation, the analysis unit was a test trajectory ($n = 10$ per system). The relative L_2 error over the full state was averaged after the initial time, omitting the exactly zero error at the shared initial state. Medians across trajectories are reported in Section 3.2, whereas Figure 2 shows the median across trajectories at each time point.

For direct Jacobian evaluation, Pearson correlation and relative Frobenius error were computed separately for off-diagonal entries on structural edges and diagonal self-dynamics at each of 20 selected states from each test trajectory. This yielded 200 state measurements per system, nested within 10 trajectories. Figure 2(a–c) visualizes the pooled edge entries across these measurements.

For perturbation and predecessor-graph analyses, the source node was the analysis unit. Each perturbation amplitude comprised 20 source conditions per system. Node and target quantities were reduced within each source before being summarized across sources. Main tables report medians; IQRs denote the empirical 25th–75th percentiles across the relevant analysis units.

Source selection used the same eligibility rule in all three systems. Candidate sources were drawn from high-, intermediate-, and low-connectivity groups, together with randomly sampled nodes. Candidates were screened at $r = 1\%$, using the detection threshold $\eta_\epsilon = 0.01|\epsilon|$ over the system-specific observation window $[t_0, T]$. Sources with no detectable response at any non-source node were excluded, and selection continued until 20 eligible sources were obtained per system. Each selected source set was then held fixed across the perturbation-amplitude sweep, direct edge contribution evaluation, predecessor construction, nonlinear interventions, and fixed-Jacobian comparisons. For Neuronal–Brain, source connectivity is measured by the sum of absolute outgoing network weights, $\sum_i |A_{is}|$.

Table 4: **Direct edge-contribution agreement at $r = 1\%$, evaluated at the respective variational arrival times before predecessor selection.** Values are source-level median [IQR] and 95% bootstrap confidence intervals. Relative L_2 errors are percentages.

Metric	Median [IQR]	95% CI	n_{valid}
Mutualistic–SW			
Signed Pearson r_C	0.9961 [0.9793, 0.9984]	[0.9852, 0.9984]	20
Relative L_2 (%)	8.00 [5.08, 18.02]	[5.15, 15.24]	20
Neuronal–Brain			
Signed Pearson r_C	0.9999 [0.9988, 1.0000]	[0.9996, 1.0000]	20
Relative L_2 (%)	1.16 [0.77, 4.31]	[0.90, 2.49]	20
SIS–Gamma			
Signed Pearson r_C	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Relative L_2 (%)	0.123 [0.074, 0.149]	[0.077, 0.145]	20

The final source sets are listed below in ascending order, using zero-based node indices in the corresponding adjacency matrices. This presentation order does not specify the candidate screening order.

Mutualistic–SW: 14, 20, 33, 93, 160, 165, 166, 167, 177, 200, 297, 301, 328, 333, 351, 356, 360, 366, 367, 388.

Neuronal–Brain: 176, 321, 323, 351, 367, 409, 414, 430, 437, 553, 599, 696, 748, 817, 829, 840, 873, 915, 916, 967.

SIS–Gamma: 0, 4, 5, 6, 7, 8, 28, 41, 42, 45, 50, 198, 199, 226, 230, 236, 256, 260, 283, 289.

D.3 Uncertainty Estimation and Validity Criteria

We estimated 95% percentile bootstrap confidence intervals for medians using 5,000 resamples with replacement. Bootstrap units were source nodes for perturbation, predecessor-graph, and intervention metrics and test trajectories for rollout metrics. For Jacobian metrics at each state, trajectories formed clusters, retaining all 20 states from each sampled trajectory. IQRs describe the dispersion of evaluated units, whereas bootstrap intervals quantify the resampling uncertainty of the median.

A Spearman statistic required at least three finite pairs and two distinct values in each vector. These metric-specific requirements were applied after source selection; a selected source could therefore contribute to some metrics but not others. Undefined values were omitted by metric, and the number of valid sources n_{valid} is reported below.

D.4 Matched Removal Intervention Tests

For each source, targets were chosen near the 25th, 50th, and 75th percentiles of nonlinear arrival times among affected nodes connected to the source in both retained predecessor graphs. Duplicate selections were removed, giving up to three targets per source. For each selected target, the intervention removed the source-connected ancestry subgraph within the retained predecessor graph, including its upstream edges. The same edges were removed from both the unperturbed and perturbed nonlinear rollouts used to compute the intervention response. The analysis used 20 sources per system.

Let $\Delta x_i^{\text{orig}}(t)$ and $\Delta x_i^{\text{rem}}(t)$ denote the ground-truth nonlinear perturbation responses before and after removal of the edge set for each target. Let p_i^{orig} and p_i^{rem} denote their peak response magnitudes.

Table 5: **Rollout and Jacobian uncertainty:** median [IQR] and 95% bootstrap confidence intervals.

Metric	Median [IQR]	95% CI	n
Mutualistic–SW			
Rollout error (%)	3.24 [3.21, 3.36]	[3.20, 3.38]	10 traj.
r_{edge}	0.9981 [0.9960, 0.9996]	[0.9980, 0.9982]	200 states
Edge RelFro (%)	3.38 [1.94, 4.61]	[3.30, 3.41]	200 states
r_{diag}	0.9996 [0.9991, 0.9999]	[0.9996, 0.9996]	200 states
Diag RelFro (%)	3.22 [0.90, 4.60]	[3.20, 3.29]	200 states
Neuronal–Brain			
Rollout error (%)	7.42 [6.68, 8.68]	[6.57, 8.97]	10 traj.
r_{edge}	0.9999 [0.9999, 0.9999]	[0.9999, 0.9999]	200 states
Edge RelFro (%)	1.11 [1.06, 1.17]	[1.06, 1.17]	200 states
r_{diag}	0.9999 [0.9999, 1.0000]	[0.9999, 1.0000]	200 states
Diag RelFro (%)	0.89 [0.85, 0.99]	[0.85, 1.11]	200 states
SIS–Gamma			
Rollout error (%)	0.272 [0.263, 0.286]	[0.262, 0.288]	10 traj.
r_{edge}	0.9980 [0.9980, 0.9982]	[0.9980, 0.9980]	200 states
Edge RelFro (%)	3.97 [3.63, 4.02]	[3.96, 3.97]	200 states
r_{diag}	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	200 states
Diag RelFro (%)	1.42 [1.29, 1.43]	[1.41, 1.42]	200 states

Table 6: **Response uncertainty across sources at $r = 1\%$:** median [IQR] and 95% bootstrap confidence intervals.

Metric	Median [IQR]	95% CI	n_{valid}
Mutualistic–SW			
Affected F_1	0.949 [0.918, 0.978]	[0.929, 0.978]	20
Arrival MAE	0.0725 [0.0531, 0.1996]	[0.0568, 0.1895]	20
Arrival ρ	0.9887 [0.9777, 0.9965]	[0.9791, 0.9959]	20
Peak ρ	0.9965 [0.9927, 0.9970]	[0.9931, 0.9969]	20
Neuronal–Brain			
Affected F_1	0.9968 [0.9940, 0.9980]	[0.9941, 0.9978]	20
Arrival MAE	0.1838 [0.1285, 0.2392]	[0.1318, 0.2334]	20
Arrival ρ	0.9999 [0.9997, 1.0000]	[0.9998, 1.0000]	20
Peak ρ	0.9969 [0.9676, 0.9999]	[0.9748, 0.9999]	20
SIS–Gamma			
Affected F_1	1.000 [1.000, 1.000]	[1.000, 1.000]	20
Arrival MAE	0.0013 [0.0007, 0.0029]	[0.0008, 0.0024]	20
Arrival ρ	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	11
Peak ρ	0.7921 [0.7311, 0.9241]	[0.7505, 0.8753]	20

The three principal intervention endpoints are

$$S_{\mathrm{peak},i} = \frac{p_i^{\mathrm{orig}} - p_i^{\mathrm{rem}}}{p_i^{\mathrm{orig}} + \varepsilon_{\mathrm{num}}}, \tag{30}$$

$$D_{\mathrm{arr},i} = \tau_i^{\mathrm{arr,rem}} - \tau_i^{\mathrm{arr,orig}}, \tag{31}$$

$$D_{\mathrm{traj},i} = \frac{\left[\sum_k \left|\Delta x_i^{\mathrm{rem}}(t_k) - \Delta x_i^{\mathrm{orig}}(t_k)\right|^2\right]^{1/2}}{\left[\sum_k \left|\Delta x_i^{\mathrm{orig}}(t_k)\right|^2\right]^{1/2} + \varepsilon_{\mathrm{num}}}. \tag{32}$$

Table 7: **Predecessor-graph agreement across sources at $r = 1\%$:** median [IQR] and 95% bootstrap confidence intervals.

Metric	Median [IQR]	95% CI	n_{valid}
Mutualistic–SW			
Predecessor-edge Jaccard	0.865 [0.766, 0.902]	[0.787, 0.898]	20
Exact predecessor agreement	0.810 [0.637, 0.885]	[0.646, 0.875]	20
Node-level predecessor set Jaccard	0.879 [0.777, 0.922]	[0.796, 0.911]	20
Contribution mass recall	0.9926 [0.9082, 0.9972]	[0.9307, 0.9965]	20
Neuronal–Brain			
Predecessor-edge Jaccard	0.991 [0.985, 0.995]	[0.985, 0.994]	20
Exact predecessor agreement	0.956 [0.913, 0.976]	[0.918, 0.971]	20
Node-level predecessor set Jaccard	0.990 [0.980, 0.994]	[0.981, 0.993]	20
Contribution mass recall	0.9946 [0.9889, 0.9958]	[0.9908, 0.9954]	20
SIS–Gamma			
Predecessor-edge Jaccard	1.000 [1.000, 1.000]	[1.000, 1.000]	20
Exact predecessor agreement	1.000 [1.000, 1.000]	[1.000, 1.000]	20
Node-level predecessor set Jaccard	1.000 [1.000, 1.000]	[1.000, 1.000]	20
Contribution mass recall	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20

If removal prevents threshold crossing within $[t_0, T]$, we set $\tau_i^{\text{arr,rem}} = T$ when computing the delay capped at the analysis window. We also evaluated the analogous change in the full response trajectory and the Jaccard similarity between affected sets. These five endpoints define the family used for multiplicity adjustment below.

For each source and target, targeted predecessor-edge removal was paired with three random edge sets having the same edge count and approximately matched receiver indegree. We subtracted the median random control effect from the targeted effect at each target, then took the median difference across targets within each source. Two-sided Wilcoxon signed-rank tests compared these differences with zero across sources. Holm adjustment was applied across the 30 tested combinations of three systems, two predecessor-graph constructions, and five intervention endpoints; adjusted $p < 0.05$ was used as the significance criterion. Table 8 reports the three principal contrasts for learned predecessor graphs; positive values indicate a larger effect of targeted than matched random removal. Arrival elimination occurred in 50/60, 57/60, and 49/49 interventions based on learned predecessor graphs in Mutualistic–SW, Neuronal–Brain, and SIS–Gamma, respectively, compared with 0/180, 3/180, and 0/147 matched random interventions. These rates summarize target conditions, while the statistical contrasts use source-level aggregation.

E Fixed-Jacobian Comparisons

E.1 Baseline Construction and Freeze State Selection

The fixed-Jacobian experiments used the same reference trajectory, perturbation amplitudes, integration horizon, sampling grid, and arrival threshold as the time-varying analysis, with the same 20-source panel for each system. All response comparisons in the main text and this appendix use the same source-level evaluation records. We compared

$$\begin{aligned} \frac{\mathrm{d}\widehat{\boldsymbol{\delta}}_{\mathrm{TV}}}{\mathrm{d}t} &= \widehat{\boldsymbol{J}}(t)\widehat{\boldsymbol{\delta}}_{\mathrm{TV}}, \qquad \frac{\mathrm{d}\widehat{\boldsymbol{\delta}}_0}{\mathrm{d}t} = \widehat{\boldsymbol{J}}_0\widehat{\boldsymbol{\delta}}_0, \\ \frac{\mathrm{d}\widehat{\boldsymbol{\delta}}_{\mathrm{ns}}}{\mathrm{d}t} &= \widehat{\boldsymbol{J}}_{\mathrm{ns}}\widehat{\boldsymbol{\delta}}_{\mathrm{ns}}, \qquad \widehat{\boldsymbol{\delta}}_{\mathrm{TV}}(t_0) = \widehat{\boldsymbol{\delta}}_0(t_0) = \widehat{\boldsymbol{\delta}}_{\mathrm{ns}}(t_0) = \epsilon \boldsymbol{e}_s, \end{aligned} \tag{33}$$

where

$$\widehat{\boldsymbol{J}}_0 = \widehat{\boldsymbol{J}}(t_0), \qquad \widehat{\boldsymbol{J}}_{\mathrm{ns}} = \left.\frac{\partial F_{\theta,\phi}(\boldsymbol{x}; \boldsymbol{A})}{\partial \boldsymbol{x}}\right|_{\boldsymbol{x}=\boldsymbol{x}_{\mathrm{ns}}}, \tag{34}$$

Table 8: **Differences between targeted and random intervention effects across sources.** Values are median [IQR] with 95% bootstrap confidence intervals. Arrival delay is reported in each system's time units.

Effect	Median [IQR]	95% CI	p_{adj}	n
Mutualistic–SW				
Peak suppression	0.986 [0.941, 1.000]	[0.944, 0.999]	5.72×10^{-5}	20
Arrival delay	1.523 [1.282, 1.655]	[1.290, 1.617]	5.72×10^{-5}	20
Target trajectory change	0.962 [0.947, 0.975]	[0.951, 0.975]	5.72×10^{-5}	20
Neuronal–Brain				
Peak suppression	1.396 [0.979, 1.862]	[0.987, 1.764]	5.72×10^{-5}	20
Arrival delay	409.848 [405.335, 410.668]	[405.653, 410.458]	5.72×10^{-5}	20
Target trajectory change	0.251 [-0.500, 0.877]	[-0.438, 0.828]	0.869	20
SIS–Gamma				
Peak suppression	1.000 [1.000, 1.000]	[1.000, 1.000]	5.72×10^{-5}	20
Arrival delay	2.511 [2.506, 2.517]	[2.508, 2.516]	5.72×10^{-5}	20
Target trajectory change	0.994 [0.976, 0.998]	[0.977, 0.998]	5.72×10^{-5}	20

Table 9: **Near-steady freeze states and their spectral abscissae, reported as ground truth / learned.**

System	Freeze time	Relative speed	Terminal drift	Spectral abscissa (ground truth / learned)
Mutualistic–SW	25.000	1.38×10^{-8}	1.18×10^{-9}	$\boldsymbol{J}(t_0)$: $13.122/6.190$ $\boldsymbol{J}(\boldsymbol{x}_{\mathrm{ns}})$: $-11.465/-11.412$
Neuronal–Brain	417.500	6.36×10^{-6}	3.20×10^{-4}	$\boldsymbol{J}(t_0)$: $0.580/0.577$ $\boldsymbol{J}(\boldsymbol{x}_{\mathrm{ns}})$ ground truth: -1.17×10^{-4} $\boldsymbol{J}(\boldsymbol{x}_{\mathrm{ns}})$ learned: -1.24×10^{-3}
SIS–Gamma	2.538	1.63×10^{-3}	5.66×10^{-4}	$\boldsymbol{J}(t_0)$: $-0.786/-0.804$ $\boldsymbol{J}(\boldsymbol{x}_{\mathrm{ns}})$: $-1.576/-1.617$

and $\boldsymbol{x}_{\mathrm{ns}}$ is a near-steady state selected by the criteria below. All baselines solve the transient perturbation task beginning at $\bar{\boldsymbol{x}}(t_0)$; the fixed models hold the Jacobian at onset or at $\boldsymbol{x}_{\mathrm{ns}}$. Corresponding oracle baselines replace $\widehat{\boldsymbol{J}}$ with Jacobians computed from the governing equations.

For Neuronal–Brain and SIS–Gamma, the final saved reference states satisfied relative speed and terminal drift tolerances of 10^{-2}. For Mutualistic–SW, the near-steady state was obtained at $t = 25$, where both diagnostics were below 10^{-6}. Table 9 reports the selected states and the spectral abscissa $\alpha_{\mathrm{sp}}(\boldsymbol{J}) = \max_\ell \mathrm{Re}\,\lambda_\ell(\boldsymbol{J})$ of the ground-truth and learned fixed matrices. The positive onset values in Mutualistic–SW and Neuronal–Brain identify unstable modes that can accumulate when a single onset Jacobian is propagated over the full horizon.

E.2 Response Recovery Across Perturbation Amplitudes

Figure 5 compares the learned time-varying Jacobian with Jacobians frozen at perturbation onset and near steady state. In Mutualistic–SW and Neuronal–Brain, onset freezing produces unstable amplification over the long horizon. Near-steady freezing suppresses most of the original transient response. At $r = 1\%$, the median response error of the onset/near-steady fixed models is $2.88 \times 10^3/1.000$ in Mutualistic–SW and $7.52 \times 10^{100}/0.999$ in Neuronal–Brain, with F_1 scores for affected nodes of $0.405/0.082$ and $0.378/0.009$, respectively. In SIS–Gamma, the two freeze points emphasize different summaries: near-steady freezing yields a response error of 0.104 versus 0.212 at onset, while onset freezing yields an F_1 score for affected nodes of 0.880 versus 0.681 near steady. The time-varying analysis yields a median affected-node F_1 of 1.000, with an IQR of [1.000, 1.000].

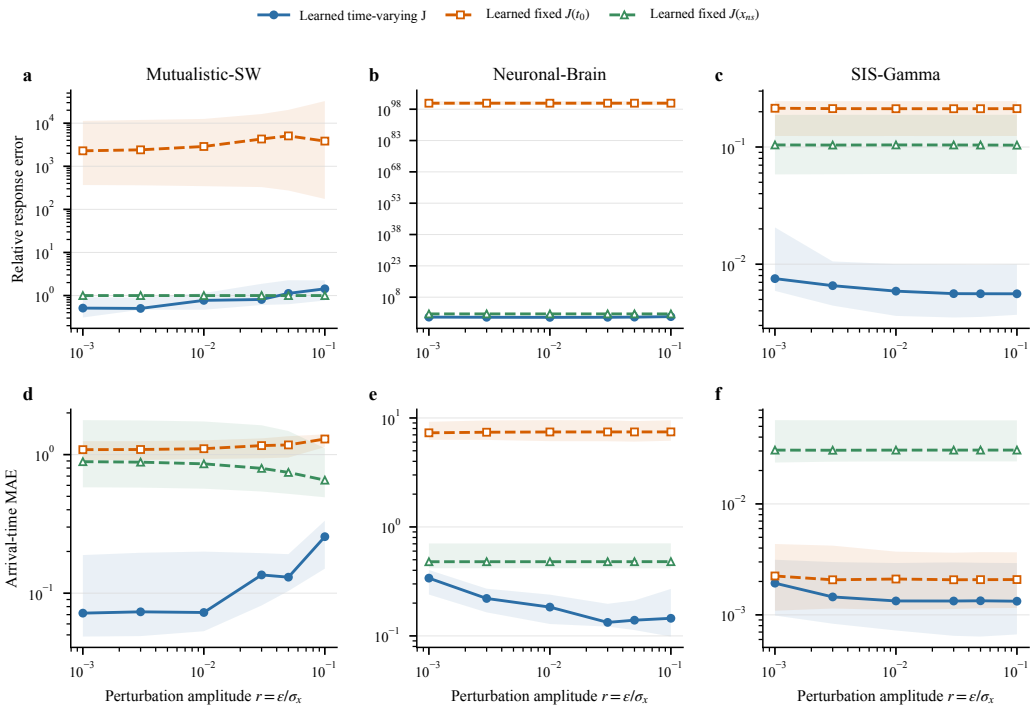

Figure 5: **Fixed-Jacobian comparison across perturbation amplitudes.** Columns show Mutualistic–SW, Neuronal–Brain, and SIS–Gamma. The upper row reports E_{resp} against the ground-truth nonlinear response, and the lower row reports arrival time MAE over shared affected nodes other than the source. Lines and shading denote medians and IQRs over valid sources within each 20-source panel. Arrival time MAE is defined only when the two responses share an affected non-source node; at $r = 1\%$, the near-steady model has 20, 11, and 18 valid sources in the three systems, respectively.

Oracle fixed Jacobians show the same pattern. At $r = 1\%$, the oracle response errors for onset freezing are 94.2, 2.42×10^{101}, and 0.173 for Mutualistic–SW, Neuronal–Brain, and SIS–Gamma, respectively. The oracle errors for near-steady freezing are 1.000, 0.999, and 0.104, with corresponding F_1 scores for affected nodes of 0.082, 0.009, and 0.667. Their similarity to the learned fixed-Jacobian results shows that substantial response errors persist under single-point freezing even when the Jacobian is computed from the governing equations.

E.3 Path Diagnostics

Figure 6 reports path diagnostics at $r = 1\%$ for both fixed and time-varying models. Exact predecessor and path metrics are conditional on targets for which both candidate and ground-truth time-varying analyses resolve a source-connected path. Jaccard similarity between path edge sets is averaged over those targets. Jaccard similarity between predecessor-edge sets compares the complete unions of recovered edges and quantifies propagation support. The predecessor-edge Jaccard after onset freezing is 0.110, 0.165, and 0.750 for Mutualistic–SW, Neuronal–Brain, and SIS–Gamma, respectively; the near-steady values are 0.025, 0.00034, and 0.500. Conditional path agreement describes the resolved targets, while predecessor-edge Jaccard measures agreement over the full retained edge sets.

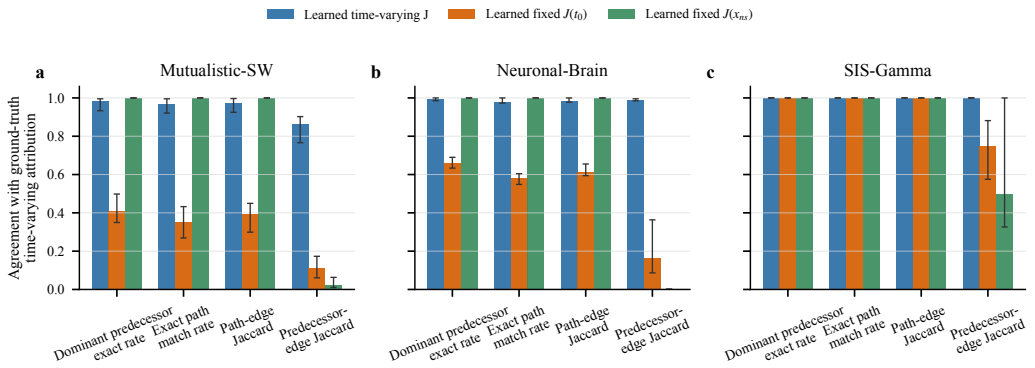

Figure 6: **Fixed-Jacobian path diagnostics at $r = 1\%$.** Bars and error bars denote medians and IQRs across sources. Exact dominant predecessor rate, exact path match rate, and path edge Jaccard are conditional on jointly resolved targets; predecessor-edge Jaccard compares the full recovered edge sets.

F Additional Network–Dynamics Combinations

We further evaluated the proposed framework across five additional dynamics–network pairings: Mutualistic–Gamma, Neuronal–SW, Neuronal–Gamma, SIS–SW, and SIS–Brain. These experiments followed the same training split, source-selection procedure, perturbation-amplitude sweep, arrival threshold, and predecessor construction used in the main experiments. Each pairing used 200 training, 10 validation, and 10 test trajectories, together with a fixed panel of 20 perturbation sources. For source-level metrics, we report the median and IQR across sources together with 95% percentile bootstrap confidence intervals from 5,000 source resamples. The number of finite source-level values is reported for each metric.

F.1 Learned Jacobian Fidelity

Table 10 reports pooled Jacobian recovery on the test trajectories. Neuronal–SW and SIS–Brain use 200 states sampled uniformly from 10 test trajectories; the other three combinations use all finite entries in their archived Jacobian evaluations. Structural-edge Pearson correlations range from 0.9915 to 0.9998, with relative L_2 errors between 1.79% and 7.13%. Diagonal relative errors range from 0.97% to 14.46%.

Table 10: **Jacobian fidelity for the additional network–dynamics combinations.** Entries are pooled point estimates over analytical and learned Jacobian entries. Neuronal–SW and SIS–Brain use 200 uniformly sampled test states; the remaining systems use all finite entries in the archived evaluations.

System	Edge r	Edge Rel. L_2 (%)	Diag. r	Diag. Rel. L_2 (%)
Mutualistic–Gamma	0.9965	5.20	0.9862	14.46
Neuronal–SW	0.9990	7.13	0.9966	3.13
Neuronal–Gamma	0.9998	1.79	0.9995	1.22
SIS–SW	0.9915	5.60	0.9972	2.96
SIS–Brain	0.9997	2.86	0.9171	0.97

F.2 Finite-Time Perturbation Inference

At $r = 1\%$, the learned time-varying Jacobians recover affected-node sets with median F_1 scores from 0.8715 to 1.0000. Arrival-order correlations range from 0.9890 to 1.0000, while peak-response correlations range from 0.7711 to 0.9999. Mutualistic–Gamma has the largest arrival-time error

among these combinations, whereas Neuronal–Gamma and both supplementary SIS systems recover the affected-node sets exactly.

Table 11: **Finite-time perturbation agreement at $r = 1\%$.** Values are source-level median [IQR] and 95% bootstrap confidence intervals. Arrival-time MAE is reported in system-specific time units.

Metric	Median [IQR]	95% CI	n_{valid}
Mutualistic–Gamma			
Affected F_1	0.8715 [0.7024, 0.9970]	[0.7256, 0.9767]	20
Arrival MAE	0.1541 [0.0698, 0.5304]	[0.0730, 0.5137]	20
Arrival ρ	0.9958 [0.9407, 0.9979]	[0.9662, 0.9975]	19
Peak ρ	0.9865 [0.8780, 0.9933]	[0.9098, 0.9932]	20
Neuronal–SW			
Affected F_1	0.9953 [0.9810, 1.0000]	[0.9848, 1.0000]	20
Arrival MAE	0.0080 [0.0054, 0.0094]	[0.0061, 0.0093]	20
Arrival ρ	0.9890 [0.9855, 0.9947]	[0.9860, 0.9939]	20
Peak ρ	0.8914 [0.8616, 0.9237]	[0.8694, 0.9147]	20
Neuronal–Gamma			
Affected F_1	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Arrival MAE	0.0016 [0.0012, 0.0089]	[0.0013, 0.0071]	20
Arrival ρ	0.9999 [0.9999, 0.9999]	[0.9999, 0.9999]	20
Peak ρ	0.9999 [0.9998, 0.9999]	[0.9998, 0.9999]	20
SIS–SW			
Affected F_1	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Arrival MAE	0.0024 [0.0017, 0.0040]	[0.0018, 0.0036]	20
Arrival ρ	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Peak ρ	0.7711 [0.6964, 0.8113]	[0.7227, 0.7970]	20
SIS–Brain			
Affected F_1	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Arrival MAE	0.0336 [0.0172, 0.0449]	[0.0183, 0.0420]	20
Arrival ρ	0.9906 [0.9688, 1.0000]	[0.9730, 1.0000]	14
Peak ρ	0.9303 [0.7882, 0.9631]	[0.8235, 0.9601]	20

The learned time-varying Jacobian gives the lowest response error in all five combinations. The difference is particularly large for the mutualistic and neuronal systems. The SIS fixed-Jacobian models remain numerically stable, but their response errors are higher and their recovered predecessor structures are less accurate at one or both freeze points.

Table 12: **Response errors in the fixed-Jacobian comparison at $r = 1\%$.** Values are source-level median [IQR] and 95% bootstrap confidence intervals for trajectory relative L_2 error.

Jacobian	Median [IQR]	95% CI	n_{valid}
Mutualistic–Gamma			
Time-varying $\boldsymbol{J}(t)$	0.4055 [0.2177, 0.7137]	[0.2332, 0.5990]	20
Fixed $\boldsymbol{J}(t_0)$	4.90×10^{11} [8.89×10^{10}, 1.78×10^{13}]	[9.44×10^{10}, 1.56×10^{13}]	20
Near-steady $\boldsymbol{J}$	0.9998 [0.9974, 1.0000]	[0.9990, 0.9999]	20
Neuronal–SW			
Time-varying $\boldsymbol{J}(t)$	0.0303 [0.0246, 0.0358]	[0.0258, 0.0342]	20
Fixed $\boldsymbol{J}(t_0)$	7.74×10^{9} [4.89×10^{9}, 1.17×10^{10}]	[5.04×10^{9}, 1.01×10^{10}]	20
Near-steady $\boldsymbol{J}$	0.6154 [0.5947, 0.6361]	[0.5993, 0.6280]	20
Neuronal–Gamma			
Time-varying $\boldsymbol{J}(t)$	0.0281 [0.0177, 0.0526]	[0.0192, 0.0497]	20
Fixed $\boldsymbol{J}(t_0)$	1.48×10^{43} [8.61×10^{42}, 1.96×10^{43}]	[1.02×10^{43}, 1.88×10^{43}]	20
Near-steady $\boldsymbol{J}$	1.0000 [0.9997, 1.0000]	[0.9998, 1.0000]	20
SIS–SW			
Time-varying $\boldsymbol{J}(t)$	0.0076 [0.0058, 0.0103]	[0.0061, 0.0095]	20
Fixed $\boldsymbol{J}(t_0)$	0.1814 [0.1261, 0.2098]	[0.1329, 0.2090]	20
Near-steady $\boldsymbol{J}$	0.1179 [0.0988, 0.1398]	[0.1029, 0.1390]	20
SIS–Brain			
Time-varying $\boldsymbol{J}(t)$	0.0074 [0.0044, 0.0148]	[0.0054, 0.0147]	20
Fixed $\boldsymbol{J}(t_0)$	0.0094 [0.0041, 0.0175]	[0.0056, 0.0166]	20
Near-steady $\boldsymbol{J}$	0.0330 [0.0207, 0.0689]	[0.0254, 0.0540]	20

Table 13: **Predecessor-edge agreement in the fixed-Jacobian comparison at $r = 1\%$.** Values are source-level median [IQR] and 95% bootstrap confidence intervals for edge-set Jaccard similarity with the analytical time-varying reference.

Jacobian	Median [IQR]	95% CI	n_{valid}
Mutualistic–Gamma			
Time-varying $\boldsymbol{J}(t)$	0.7523 [0.5173, 0.9732]	[0.5393, 0.9546]	20
Fixed $\boldsymbol{J}(t_0)$	0.2552 [0.0536, 0.7805]	[0.0706, 0.6641]	20
Near-steady $\boldsymbol{J}$	0.0146 [0.0040, 0.0754]	[0.0058, 0.0608]	20
Neuronal–SW			
Time-varying $\boldsymbol{J}(t)$	0.9266 [0.9032, 0.9683]	[0.9105, 0.9619]	20
Fixed $\boldsymbol{J}(t_0)$	0.0404 [0.0257, 0.0506]	[0.0278, 0.0472]	20
Near-steady $\boldsymbol{J}$	0.0000 [0.0000, 0.0000]	[0.0000, 0.0000]	20
Neuronal–Gamma			
Time-varying $\boldsymbol{J}(t)$	0.9965 [0.9931, 1.0000]	[0.9932, 0.9983]	20
Fixed $\boldsymbol{J}(t_0)$	0.9138 [0.7959, 0.9490]	[0.8070, 0.9405]	20
Near-steady $\boldsymbol{J}$	0.0000 [0.0000, 0.0000]	[0.0000, 0.0000]	20
SIS–SW			
Time-varying $\boldsymbol{J}(t)$	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Fixed $\boldsymbol{J}(t_0)$	1.0000 [0.7619, 1.0000]	[0.7778, 1.0000]	20
Near-steady $\boldsymbol{J}$	0.4495 [0.3563, 0.6036]	[0.3875, 0.5357]	20
SIS–Brain			
Time-varying $\boldsymbol{J}(t)$	1.0000 [0.8992, 1.0000]	[0.9201, 1.0000]	20
Fixed $\boldsymbol{J}(t_0)$	1.0000 [0.9223, 1.0000]	[0.9433, 1.0000]	20
Near-steady $\boldsymbol{J}$	0.9545 [0.7096, 1.0000]	[0.7220, 1.0000]	20

F.3 Arrival-Conditioned Edge Contributions and Propagation Structure

Arrival-conditioned edge contributions remain closely aligned with their analytical counterparts. Signed contribution correlations are at least 0.9970, and relative L_2 errors range from 0.205% to 6.30%. The resulting predecessor graphs attain median edge-set Jaccard similarities between 0.7523 and 1.0000. All five learned predecessor graphs retain source-connected routes for every evaluated source.

Table 14: **Arrival-conditioned edge-contribution and predecessor-graph agreement at $r = 1\%$.** Contribution metrics are evaluated at each node's respective variational arrival time. Values are source-level median [IQR] and 95% bootstrap confidence intervals.

Metric	Median [IQR]	95% CI	n_{valid}
Mutualistic–Gamma			
Signed Pearson r_C	0.9970 [0.9865, 0.9990]	[0.9913, 0.9989]	20
Contribution Rel. L_2 (%)	6.295 [3.855, 13.398]	[3.961, 10.862]	20
Predecessor-edge Jaccard	0.7523 [0.5173, 0.9732]	[0.5393, 0.9546]	20
Exact predecessor rate	0.6824 [0.5221, 0.9482]	[0.5294, 0.9327]	20
Predecessor-set Jaccard	0.7548 [0.5515, 0.9783]	[0.5597, 0.9607]	20
Neuronal–SW			
Signed Pearson r_C	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Contribution Rel. L_2 (%)	0.439 [0.415, 0.507]	[0.415, 0.500]	20
Predecessor-edge Jaccard	0.9266 [0.9032, 0.9683]	[0.9105, 0.9619]	20
Exact predecessor rate	0.9202 [0.8891, 0.9602]	[0.8930, 0.9514]	20
Predecessor-set Jaccard	0.9444 [0.9265, 0.9736]	[0.9322, 0.9701]	20
Neuronal–Gamma			
Signed Pearson r_C	1.0000 [0.9997, 1.0000]	[0.9998, 1.0000]	20
Contribution Rel. L_2 (%)	0.468 [0.392, 2.024]	[0.396, 1.317]	20
Predecessor-edge Jaccard	0.9965 [0.9931, 1.0000]	[0.9932, 0.9983]	20
Exact predecessor rate	0.9933 [0.9865, 1.0000]	[0.9866, 0.9983]	20
Predecessor-set Jaccard	0.9966 [0.9939, 1.0000]	[0.9942, 0.9987]	20
SIS–SW			
Signed Pearson r_C	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Contribution Rel. L_2 (%)	0.205 [0.163, 0.340]	[0.173, 0.331]	20
Predecessor-edge Jaccard	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Exact predecessor rate	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Predecessor-set Jaccard	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
SIS–Brain			
Signed Pearson r_C	1.0000 [1.0000, 1.0000]	[1.0000, 1.0000]	20
Contribution Rel. L_2 (%)	0.454 [0.130, 1.058]	[0.156, 0.847]	20
Predecessor-edge Jaccard	1.0000 [0.8992, 1.0000]	[0.9201, 1.0000]	20
Exact predecessor rate	1.0000 [0.7458, 1.0000]	[0.7679, 1.0000]	20
Predecessor-set Jaccard	1.0000 [0.9365, 1.0000]	[0.9484, 1.0000]	20

Matched removal experiments further evaluate the functional relevance of the learned predecessor structures. Across all five combinations, removing learned predecessor edges produces larger changes in target responses than removing matched random edge sets. The median differences are positive for peak suppression, arrival delay, and target-trajectory change.

Table 15: **Matched edge-removal intervention results.** Values are source-level median [IQR] differences between learned predecessor-edge removal and three matched random removals, with 95% bootstrap confidence intervals. Positive values indicate a larger effect under predecessor-edge removal; arrival delay uses system-specific time units. Holm adjustment is applied within each endpoint across the five systems and two predecessor-graph constructions.

Effect	Median [IQR]	95% CI	p_{adj}	n
Mutualistic–Gamma				
Peak suppression	1.003 [1.000, 1.045]	[1.000, 1.039]	1.91×10^{-5}	20
Arrival delay	2.270 [2.023, 2.860]	[2.091, 2.844]	1.91×10^{-5}	20
Target trajectory change	0.948 [0.880, 0.996]	[0.901, 0.996]	1.91×10^{-5}	20
Neuronal–SW				
Peak suppression	0.994 [0.990, 1.000]	[0.992, 1.000]	1.91×10^{-5}	20
Arrival delay	4.929 [4.924, 4.932]	[4.925, 4.931]	1.91×10^{-5}	20
Target trajectory change	0.902 [0.885, 0.924]	[0.888, 0.918]	1.91×10^{-5}	20
Neuronal–Gamma				
Peak suppression	0.928 [0.436, 1.084]	[0.458, 1.050]	1.91×10^{-5}	20
Arrival delay	14.953 [11.159, 15.115]	[14.685, 15.101]	1.91×10^{-5}	20
Target trajectory change	0.227 [0.068, 0.472]	[0.085, 0.462]	7.25×10^{-5}	20
SIS–SW				
Peak suppression	0.989 [0.982, 0.999]	[0.983, 0.999]	1.91×10^{-5}	20
Arrival delay	2.065 [2.058, 2.069]	[2.061, 2.069]	1.91×10^{-5}	20
Target trajectory change	0.985 [0.980, 0.996]	[0.982, 0.995]	1.91×10^{-5}	20
SIS–Brain				
Peak suppression	0.946 [0.935, 0.977]	[0.940, 0.965]	1.91×10^{-5}	20
Arrival delay	9.248 [9.033, 9.290]	[9.070, 9.285]	1.91×10^{-5}	20
Target trajectory change	0.939 [0.919, 0.976]	[0.922, 0.963]	1.91×10^{-5}	20

Taken together, the results across all evaluated systems support the central conclusion of this work: time-varying Jacobians learned from observed trajectories can recover finite-time perturbation responses, arrival-conditioned edge contributions, and source-connected propagation structures.

AI USE STATEMENT

In this study, we used AI tools only for grammar correction, language polishing, and troubleshooting LaTeX compilation errors. We manually reviewed all AI-edited text to ensure that the wording was accurate. All research data and figures were generated and prepared independently by the authors without the use of AI tools. All experiments were independently designed and conducted by us. We take full responsibility for the methods, experiments, results, discussion, and final content of this paper.