CAUSAL DISCOVERY VIA COMPLEXITY-GATED SPIKING NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

We introduce X-SPK (Cross-channel SPiKing), a fully unsupervised, gradient-free spiking method for causal discovery in multivariate time series. Each channel is encoded by an independent fixed leaky reservoir and converted to equal-rate events; directed evidence is accumulated in a shared $D \times D$ gate matrix via winner-take-all spike-timing plasticity and read out through the competition ratio W_{ij}/W_{jj}. The underlying mechanism is geometric: if $i \to j$, conditioning on events of i increases and concentrates j's near-future occupancy of its threshold-crossing region; this yields more threshold crossings, higher decayed trace mass, and a systematic WTA advantage in the causal direction. The method uses unified model settings with limited dataset-specific preprocessing and is evaluated on CausalRivers, Lorenz, Mackey-Glass, and TimeGraph (including non-stationary and confounded variants), reporting standard metrics against PCMCI+/GPDC, DYNOTEARS, VARLiNGAM, NeuralGC (Tank et al., 2021), CCM (Sugihara et al., 2012), and Rhino/RhinoLite across linear, nonlinear, and non-stationary regimes.

1 INTRODUCTION

Recovering directed graphs from nonlinear, non-stationary multivariate time series is difficult when identifiability depends on linear-Gaussian assumptions or stationarity constraints (Peters et al., 2017; Mooij et al., 2016). We target this setting with a fully unsupervised estimator that transforms each time series into an event-based internal representation: each variable is embedded by an independent fixed reservoir, converted to equal-rate spikes, and integrated by local winner-take-all (WTA) plasticity into a directed gate matrix.

Our central claim is geometric and causal: if $i \to j$, conditioning on events of i increases j's near-future occupancy of its threshold-crossing region. A trace-based WTA-STDP (spike-timing-dependent plasticity, STDP) rule turns this asymmetry into directional edge evidence and directly outputs a directed graph.

Contributions are: (i) a per-channel reservoir plus WTA-gate architecture for causal discovery, (ii) a formal bridge from pattern complexity to event occupancy, trace accumulation, and oriented weight accumulation, and (iii) comprehensive experiments across nonlinear synthetic DAGs, TimeGraph non-stationary/confounded settings, Lorenz chaos, and real CausalRivers data.

Related Work. Classical time-series causal discovery divides along two assumptions. DYNOTEARS and VARLiNGAM assume linear relationships (Pamfil et al., 2020; Hyvärinen et al., 2010); PCMCI+ is not restricted to linearity (paired with a nonlinear CI test such as GPDC, as we do below) but, like the others, assumes stationarity (Runge et al., 2019). Non-stationarity-aware alternatives — CD-NOD (Huang et al., 2020), LPCMCI, SVAR-FCI — remain constraint-based, inheriting CI-testing's combinatorial cost; we compare against CD-NOD directly in Appendix A.5. NeuralGC (Tank et al., 2021) and Rhino (Gong et al., 2023), both evaluated as baselines throughout this paper, instead replace CI-testing with gradient-based nonlinear fitting, trading its combinatorial cost for non-convex optimization. Reservoir computing provides expressive gradient-free temporal embeddings (Jaeger, 2001; Maass et al., 2002; Schrauwen et al., 2007), and STDP provides local temporal credit assignment (Gerstner et al., 1996; Song et al., 2000; Gilson et al., 2009; Izhikevich, 2006). Our causal framing follows Independence of Cause and Mechanism / Minimum Description Length (ICM/MDL) asymmetry (Janzing & Schölkopf, 2010; Peters et al., 2017; Schölkopf et al., 2021; Grünwald, 2007).

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.

Relative to closely related work, Zhao et al. (2024) require interventions, Krakovská (2021) rely on cross-predictions in reconstructed state spaces, and shadow-mapping approaches (Vowels et al., 2021; Sugihara et al., 2012) rely on manifold reconstruction and attractor stationarity. Existing spike-based causal studies often treat causality as auxiliary to prediction or control (Lansdell & Kording, 2023; Boedecker et al., 2012); X-SPK instead makes graph orientation the primary output, remaining fully observational throughout.

2 Architecture: Per-Channel Leaky Reservoirs and Winner-Take-All Causal Gate

Notation. D is the number of observed channels (the graph size); n_{res} is the reservoir's hidden dimension, a separate, per-channel quantity not to be confused with D. Each channel $i \in \{1, \dots, D\}$ has its own reservoir state $\mathbf{h}_i(t) \in \mathbb{R}^{n_{\text{res}}}$ and threshold θ_i (state norm and spike indicator are defined next, in Design summary); r_{target} is the shared target spike rate. $\mathbf{W} \in \mathbb{R}_{\geq 0}^{D \times D}$ is the gate matrix and $\boldsymbol{\rho} \in \mathbb{R}^D$ the trace vector, both shared across channels; $e_j \in \mathbb{R}^D$ is the standard basis vector with a 1 at coordinate j. γ, β, η, and λ_{L2} are, respectively, the trace-decay rate, self-bias, WTA learning rate, and gate weight-decay coefficient; θ_{comp} is the competition-ratio decision threshold and $\epsilon = 10^{-8}$ a numerical-stability constant.

Design summary. Each channel i drives an independent fixed leaky reservoir. Let $s_i(t) := \|\mathbf{h}_i(t)\|_2$ denote the channel-i state norm and $\delta_i(t) := \mathbb{I}[s_i(t) \geq \theta_i]$ the resulting spike indicator used throughout this section. Four stages follow: reservoir dynamics, eventization, gate update, and readout.

Reservoir dynamics. Input weights $\mathbf{w}_{\text{in}}^{(i)}$ and recurrent weights $\mathbf{W}_{\text{res}}^{(i)}$ are drawn once and then held fixed; the sparsity hyperparameter p applies to the recurrent matrix $\mathbf{W}_{\text{res}}^{(i)}$, while the spectral-radius hyperparameter $\lambda_{\max}$ rescales that same recurrent matrix after initialization. The reservoir update is

$$\mathbf{h}_i(t) = (1-\alpha)\,\mathbf{h}_i(t-1) + \alpha \tanh\big(\mathbf{w}_{\text{in}}^{(i)}\, x_i(t) + \mathbf{W}_{\text{res}}^{(i)}\, \mathbf{h}_i(t-1)\big), \tag{1}$$

with no gradient updates to reservoir weights. The spectral radius of $\mathbf{W}_{\text{res}}^{(i)}$ (set < 1) controls recurrent stability and memory horizon, while α is a separate leaky-integration coefficient (in code: `LEAKAGE`) that sets the per-step state interpolation; when $\alpha = 1$, Equation equation 1 reduces to the standard non-leaky update.

Eventization. Spikes are generated by per-channel quantile thresholding of norm trajectories at equal target rate r_{target}:

$$\theta_i = \text{Quantile}_{1-r_{\text{target}}}\big(\{s_i(t)\}\big). \tag{2}$$

In the current implementation, θ_i is estimated once from a post-washout calibration sweep of the norm trajectory and then held fixed during gate updates.

Gate update. We maintain a nonnegative gate matrix $\mathbf{W} \in \mathbb{R}_{\geq 0}^{D \times D}$ and trace vector $\boldsymbol{\rho}$. Both are initialised to zero at the start of training; across later epochs, gate weights are retained while traces are reset to zero at the start of each replay sweep. Reservoir hidden states are also reset and rewarmed for T_{wash} steps before each calibration or training pass. At each time step, traces first decay,

$$\rho_i \leftarrow \gamma\, \rho_i, \quad \forall\, i. \tag{3}$$

then, for each firing target j, we run WTA credit assignment ($\mathbf{e}_j$ as defined in Notation above):

$$\tilde{\boldsymbol{\rho}}_j = \boldsymbol{\rho} + \beta\, \mathbf{e}_j, \tag{4}$$

$$i^* = \arg\max_i\, \tilde{\rho}_{j,i}, \tag{5}$$

$$W_{i^* j} \leftarrow W_{i^* j} + \tfrac{\eta}{\sqrt{D}}, \tag{6}$$

$$W_{\cdot j} \leftarrow W_{\cdot j}\,(1 - \lambda_{L2}). \tag{7}$$

and finally traces are updated with the current step's spikes:

$$\rho_i \leftarrow \rho_i + \delta_i(t), \quad \forall\, i. \tag{8}$$

Readout. Edges are read from the competition ratio, where $\lambda_{L2} = 5 \times 10^{-4}$ is a fixed column-wise shrinkage coefficient that prevents unbounded gate growth. We do not sweep λ_{L2} in the Optuna search, so it is omitted from Table 6.

$$c_{ij} = \frac{W_{ij}}{W_{jj} + \epsilon}, \qquad \hat{A}_{ij} = \mathbb{I}[c_{ij} > \theta_{\text{comp}},\ i \neq j]\,. \tag{9}$$

We use a unified hyperparameter setting in main experiments and report sweep ranges, selected configurations, runtime, and ablations in the appendix (Sections A.8–A.9).

Implementation schedule. The equations above describe one replay sweep; training uses a fixed multi-stage schedule. (1) Calibrate per-channel quantile thresholds once on observed data after washout. (2) If $n_{\text{syn}} > 0$, pretrain $\mathbf{W}$ on n_{syn} synthetic DAG sequences. (3) If $n_{\text{pre}} > 0$, run self-gate pretraining on real data (update diagonal W_{jj} only). (4) Run n_{epochs} main sweeps on observed data; with `two_pass=True`, each sweep applies a residual pass-2 refinement (lower spike rate, shorter trace decay, cross-margin, self-inhibition, frozen self-gates). Table 9 evaluates removing pretraining or pass-2 refinement.

This architecture separates representation, eventization, and orientation. The reservoir performs nonlinear temporal lifting but stays fixed, so causal direction is never confounded with gradient-based fitting to specific labels. Quantile eventization enforces comparable spike rates across channels, so directional evidence depends on when events occur, not how many each channel emits. The gate then accumulates only local precedence evidence through WTA traces, giving a mechanism whose state is interpretable directly through W_{ij}/W_{jj} ratios.

3 From Event Patterns to Causal Direction

Core intuition. When variable i causes j, trajectory geometry in j's reservoir is systematically constrained by i's state: post-cause trajectories concentrate into threshold-crossing regions more often and more reliably than in the reverse direction. WTA-STDP detects this asymmetry by accumulating trace evidence over time; the direction that wins competition more consistently accumulates a larger gate weight.

Formal objective. This section makes precise how directed structure emerges from spike timing: asymmetric event occupancy creates an asymmetric trace, and WTA-STDP turns that into an asymmetric edge weight. Full proofs are in Appendix B.

3.1 Occupancy Rates and the Causal Asymmetry

Recall the spike indicator $\delta_i(t) := \mathbb{I}[s_i(t) \geq \theta_i]$ (Section 2); let $N_i^{(\Delta)}(t) := \sum_{\tau=1}^{\Delta} \delta_i(t+\tau)$ count i's spikes in the Δ-step window after anchor time t. By stationarity, each channel's marginal spike probability $\mathbb{P}(\delta_i(t){=}1)$ is well-defined, and quantile eventization (Equation equation 2) sets it to r_{target} for every channel — this is what removes raw firing rate as a confound, since a channel cannot look "predictive" merely because it fires more often. Define the conditioned occupancy rate

$$r_{i\to j} := \mathbb{P}\big(N_j^{(\Delta)}(t) > 0 \mid \delta_i(t){=}1\big), \tag{10}$$

the probability that j spikes within Δ steps of i spiking. Up to sign, $-\log r_{i\to j}$ is an idealized MDL codelength — the entropy proxy standard in ANM/IGCI-style causal discovery for the uncomputable Kolmogorov complexity (Janzing & Schölkopf, 2010; Peters et al., 2017; Grünwald, 2007) — but we work with $r_{i\to j}$ directly below, since it is what the algorithm measures.

Claim. If i causes j, $r_{i\to j} > r_{j\to i}$.

This is an instance of Independence of Cause and Mechanism (ICM) (Peters et al., 2017): when i causes j, j's dynamics admit the additive-noise decomposition

$$E_j = f(H_i) + N_j, \qquad N_j \perp H_i, \tag{11}$$

where H_i is i's trajectory state entering j's dynamics, f the (possibly nonlinear) coupling, and N_j collects j's locally generated dynamics, independent of i. No such decomposition need exist in reverse once f is non-degenerate — this is what breaks the symmetry between $r_{i\to j}$ and $r_{j\to i}$, and it is the mechanistic content a bare appeal to ICM does not by itself supply.

Lemma 3.1 (Signal-Noise Origin of Occupancy Asymmetry). *If $E_j = f(H_i) + N_j$ with $N_j \perp H_i$ and f excitatory — its variation systematically increases j's post-cause state norm, not merely non-constant — then $r_{i\to j} > r_{j\to i}$.*

Proof sketch. A spike (threshold excursion) in i is transmitted through f and, because f is excitatory, systematically pushes j's trajectory toward its own threshold region shortly after — these are *signal-driven* co-occurrences that inflate $r_{i\to j}$. A spike in j can also arise purely from N_j, independent of i; these *noise-driven* spikes carry no information about upstream timing in i and dilute $r_{j\to i}$, since j has no symmetric channel feeding back into i's dynamics. A purely inhibitory f (systematically suppressing j's state norm) pushes the opposite way and is not covered by this lemma as stated; we assume excitatory coupling throughout. Full derivation, including the step to conditional trace dominance, is in Appendix B. □

Neither channel is assumed to fire first by construction: either may spike first on any given occasion, and both occupancy rates above are computed over the same bidirectional data. The asymmetry comes from the *composition* of signal- versus noise-driven co-occurrences under Equation equation 11, not a timing convention.

Remark 3.2 (Falsifiable SNR prediction). Lemma 3.1 makes a testable prediction we do not evaluate in this paper: $r_{i\to j} - r_{j\to i}$ should shrink as the noise variance of N_j grows relative to f's gain, and vanish if f is degenerate. An SNR sweep on synthetic data confirming this trend would be the strongest available falsification test for this mechanism; we leave it to future work.

Relation to directed information and transfer entropy. Once quantile-normalized and evaluated at a fixed lag, $r_{i\to j}$ is structurally a trajectory-space instantiation of directed information and transfer entropy (Massey, 1990; Schreiber, 2000), motivated here via MDL/ICM rather than derived from scratch: the asymmetry measure itself is not new, and X-SPK's contribution is *how* it is computed and accumulated — over reservoir-trajectory threshold crossings, via WTA-STDP.

3.2 WTA STDP Selects the Causal Direction

When channel j fires at time t, the WTA rule (Equation 5) assigns credit to the variable i^* with the highest pre-synaptic trace:

$$i^*(t) = \arg\max_i \left[\gamma^{t-t_i^{\text{last}}} \cdot \rho_i^{\text{prev}} + \beta\, \mathbb{I}[i=j]\right], \qquad W_{i^*j} \leftarrow W_{i^*j} + \tfrac{\eta}{\sqrt{D}}, \tag{12}$$

where t_i^{last} is the last spike time of channel i before t, and ρ_i^{prev} denotes the trace value stored at that last spike time before the additional decay factor $\gamma^{t-t_i^{\text{last}}}$ is applied. This is an event-time rewriting of the stepwise implementation in Equations equation 3 and equation 8: in code, the trace is decayed once at every timestep and then used directly in WTA, so no double decay is applied. Because the trace is highest for the variable that fired most recently and consistently before j, the rule implements a winner-take-all analogue of LTP, reinforcing the most temporally proximate and frequent predecessor.

Theorem 3.3 (Asymmetric WTA Accumulation). *If i causes j, then*

$$\mathbb{E}[\Delta W_{ij}] > \mathbb{E}[\Delta W_{ji}]. \tag{13}$$

(Full proof in Appendix B.)

Scope of this result. Theorem 3.3 compares a true edge to its own reverse in isolation — it says nothing about non-edges (unrelated, confounded, or mediated pairs). Since WTA always crowns *some* winner on every target spike, a channel can receive credit even when the true parent isn't among the competitors; suppressing these false positives is therefore an empirical property of the thresholded competition ratio (Section 5, "Handling of (latent) confounders"; Section 4), not something this theorem guarantees.

The WTA update reinforces stable directional advantage. Since causal directions produce larger conditioned event occupancy and therefore larger conditioned traces (Appendix B), they win WTA competitions more often and accumulate larger expected weights.

3.3 WHY CYCLES ARE SUPPRESSED

Remark 3.4 (Cycle Instability Under an Ordering Criterion). If retained directed edges satisfy a strict global ordering criterion induced by directional margins, then directed cycles are unstable — true by definition of a strict ordering; the substantive, open question is whether X-SPK's learned margins actually satisfy it (Corollary 3.5).

(Proof in Appendix B.)

Corollary 3.5 (Conditional Acyclicity)**.** *If the retained directed edges' margins admit the strict global ordering hypothesized in Remark 3.4, then directed cycles cannot persist in the learned graph.*

Proof. By Theorem 3.3, causal edges receive larger expected weight updates than their reverse. If, in addition, the retained margins admit the strict global ordering required by Remark 3.4, directed cycles are unstable and cannot persist. □

The ordering premise itself is not derived from Theorem 3.3 alone (Appendix B gives the full argument and states exactly what is missing). Acyclicity in practice is threshold-conditional, not automatic: a targeted diagnostic (Appendix A.8) finds cyclic graphs in every run under this paper's main-result threshold and acyclic graphs in every run under a density-calibrated one, consistent with Corollary 3.5's conditional guarantee. A density-matched random-graph control (same appendix) rules out pure sparsity as the explanation under the practically relevant, dataset-agnostic version of that threshold ($p \approx 6 \times 10^{-5}$).

Interpretation. In one line: simpler dynamics → higher concentration (regularity) → more spikes → higher trace → WTA advantage. A cause leaves a simpler, more regular footprint in its effect's trajectory than the reverse, and that regularity is exactly what the reservoir-and-gate mechanism is built to pick up; in the anti-causal direction, source and mechanism are entangled, so no comparable footprint appears and no trace advantage accumulates. This is the dynamical-systems reading of a standard causal-discovery intuition — Independence of Cause and Mechanism, and its information-theoretic form, IGCI/MDL (Janzing & Schölkopf, 2010; Peters et al., 2017; Grünwald, 2007): the causal factorization is the simpler one, and here that simplicity shows up as event concentration rather than as a shorter code length. The STDP gate is a streaming version of this comparison, tallying precedence evidence spike by spike (the only offline step is the single quantile calibration pass over the observed data). Because each channel is measured through its own independent reservoir, the comparison is made per channel and combined by precedence into a directed graph.

4 EXPERIMENTS

We test three hypotheses: (H1) X-SPK recovers causal direction in nonlinear/chaotic regimes where linear methods fail; (H2) the STDP trace is the architecturally critical component; (H3) performance scales to large graphs. Our experimental setup combines datasets spanning nonlinear dynamics, real-world hydrological networks, and linear-Gaussian processes with varying stationarity and confounding patterns. We report four standard graph-recovery metrics throughout: F1 (edge-detection F1 score), AUROC and AUPRC (threshold-free ranking quality, area under the ROC and precision-recall curves), and SHD (Structural Hamming Distance, the edit distance between predicted and true graphs; lower is better, all others higher is better).

4.1 DATASETS

We evaluate on five benchmark suites with complementary properties: **Mackey–Glass** Mackey & Glass (1977) provides near-chaotic nonlinear synthetic DAGs; we use five cases: **A** (Collider+Ghost, confounding structure), **C** (Dense Hub), **F** (Gated Nonlinear), **G** (Step Functions), and **I** (Regime Switching), designed to stress linear/stationary assumptions while preserving controlled ground truth. **Lorenz Attractor** Lorenz (1963) is a 3D chaotic ODE with $\sigma = 10$, $\rho = 28$, $\beta = 8/3$ and dynamics $\dot{x} = \sigma(y - x)$, $\dot{y} = x(\rho - z) - y$, $\dot{z} = xy - \beta z$, isolating deterministic chaotic coupling (5 edges, 3 nodes). **TimeGraph** Ferdous et al. (2025) provides structural causal model (SCM) benchmarks with controlled non-stationarity, irregular sampling, and latent-confounder variants (A1/C1, A1C/C1C). **CausalRivers** (Stein et al., 2025) is a real hydrological benchmark from river-gauge discharge

with strong autocorrelation, nonlinear propagation, seasonal/nonstationary effects, and confounder scenarios; preprocessing includes rolling normalization, temporal aggregation, and elevation-based masking. **Tigramite Benchmark/Synthetic** Runge (2020) provides linear-Gaussian VAR ground truth under native linear/stationary assumptions.

4.2 Baseline Estimators

We compare X-SPK with seven baselines spanning linear and nonlinear causal discovery: **PCMCI+/PCMCI+/GPDC** (Runge et al., 2019) (PC with CI tests, GPDC for nonlinear dependencies), **DYNOTEARS** (Pamfil et al., 2020) and **RhinoLite** (Gong et al., 2023) (a linear Lasso-VAR proxy for Rhino, not the full nonlinear model; Appendix A.4), **VARLiNGAM** (Hyvärinen et al., 2010) (linear VAR with non-Gaussian identifiability), **NeuralGC** (Tank et al., 2021) (MLP Granger with group lasso), and **CCM** (Sugihara et al., 2012) (convergent cross-mapping). On nonlinear datasets (Lorenz, CausalRivers, Mackey–Glass), PCMCI+/GPDC, NeuralGC, and CCM replace linear counterparts; PCMCI+/GPDC is omitted at $D{>}8$ and CCM at $D{>}5$ for runtime. **Baselines were chosen for methodological diversity**, not recency. A nonstationarity-aware baseline (CD-NOD) and an established external benchmark (fMRI Netsim) are added in Appendix A.5.

4.3 Comprehensive Benchmark at $D{=}5$: All Estimators

We evaluate X-SPK against PCMCI+/GPDC, DYNOTEARS, NeuralGC, CCM, and RhinoLite across CausalRivers, Lorenz, TimeGraph, and five Mackey–Glass pathological topologies; for TG rows B1, B2 $D{=}6$, C1 $D{=}6$, and B2 $D{=}8$, PCMCI+/GPDC and NeuralGC replace linear variants, and for TG C2, B2C, and C1 $D{=}8$, NeuralGC replaces VARLiNGAM (PCMCI+/GPDC not run on C-type non-stationarity or $D{=}8$). X-SPK uses the unified configuration from Section A.8; baselines use defaults with $\tau_{\max}{=}2$. We report means over 10 seeds (Lorenz), 15 seeds (CausalRivers $D{=}5$), 10 seeds (CausalRivers $D{=}10$), and 5 seeds (TimeGraph). Results are in Table 1, with scalability in Appendix Table 3 and MG-pathology details in Table 2; the **Null** baseline (AUROC$=0.5$, AUPRC=edge density, F1$=0$, SHD$=\binom{D}{2}$) is not printed as rows. X-SPK is listed under a fixed and a permutation-null threshold (see caption); unless stated, text numbers are fixed-threshold.

Main result. X-SPK dominates nonlinear and chaotic regimes while remaining competitive as a ranker in linear regimes. Across all five Mackey–Glass topologies (Table 2), X-SPK leads F1 (0.519–0.667) and AUROC (0.650–0.791) under the default tuning/lag protocol used throughout this paper — Appendix A.4 shows this lead is not universal under equal tuning budget or extended lag; NeuralGC is the strongest nonlinear baseline (F1 0.331–0.533, AUROC 0.596–0.713). On the step-functions topology, CCM marginally leads AUROC (0.672 vs 0.650) while X-SPK leads F1 and AUPRC. On Lorenz it reaches AUROC/AUPRC of 1.000 with F1$=0.909$ (SHD$=1$), well above NeuralGC (AUROC 0.940) and CCM (AUROC 0.780). NeuralGC's advantages are confined to higher-dimensional scaling ($D{=}8$, TG C1) and a low-n uncertain condition (TG B2C, $n{=}3$); on the datasets that stress nonlinearity most heavily—chaotic dynamics (Lorenz, Mackey–Glass) and real-world river networks (CausalRivers)—X-SPK's AUROC margins over NeuralGC range from 0.06 to 0.28.

Regime dependence. On linear-Gaussian Tigramite Benchmark, PCMCI+, Rhino, and VAR-LiNGAM are perfect (F1$=1$, SHD$=0$), while X-SPK keeps strong ranking (AUROC$=0.975$) but lower thresholded F1 (0.400). On Tigramite Synthetic ($D{=}5$, $D{=}10$), DYNOTEARS/Rhino remain strongest (F1 $\approx$ 0.69–0.79), with X-SPK moderate. On CausalRivers $D{=}5$, X-SPK leads F1 and AUROC (0.554/0.884), well above the nonlinear baselines NeuralGC (0.237/0.570) and CCM (0.274/0.605), consistent with nonlinear, seasonal river dynamics; at $D{=}10$ X-SPK again leads AUROC (0.807) and F1 (0.323).

Distributional robustness. On TimeGraph B1 ($D{=}6$, pure nonlinear Gaussian), X-SPK's F1$=0.182$ is unchanged across all five seeds (zero std, driven by a near-complete predicted graph rather than sparse recovery — Section 4.4), while NeuralGC ($0.108{\pm}0.069$) and PCMCI+/GPDC ($0.100{\pm}0.224$) produce high-variance near-zero means. On TimeGraph B2 $D{=}6$, X-SPK leads AUROC (0.669) while NeuralGC (0.566) and linear baselines ($\approx$0.49–0.53) trail. At $D{=}8$, NeuralGC leads F1 and AUROC (0.146/0.615) marginally over X-SPK (0.133/0.594), though the gap lies within noise at $n{=}5$ seeds; Rhino achieves the lowest SHD (3.8) via sparse prediction. TimeGraph B2C ($D{=}6$, mixed noise with a hidden confounder) remains challenging: PCMCI+, DYNOTEARS, and Rhino

all predict F1$= 0$, while NeuralGC achieves mean F1$= 0.228$ (AUROC$= 0.751$) and X-SPK follows at F1$= 0.182$ (AUROC$= 0.626$); however, NeuralGC results are based on n=3 seeds with high variance (AUROC std 0.178), so the ordering should be treated with caution. On TimeGraph C2 (nonlinear + trend + seasonal, D=6), X-SPK leads at AUROC$= 0.788$ versus NeuralGC (0.632), Rhino (0.587), and PCMCI+ (0.433).

Nonlinear non-stationarity and large-graph scaling. TimeGraph C1 confirms partial regime-dependent failure: PCMCI+/GPDC (F1$= 0.088$, AUROC$= 0.398$) substantially outperforms linear PCMCI+ (AUROC$= 0.247$) due to its nonparametric CI test, but still trails X-SPK (AUROC$= 0.606$) and NeuralGC (0.535) under non-stationarity. At D=15–30 on CausalRivers (Appendix Table 3), X-SPK maintains the highest AUROC throughout; Rhino achieves the lowest SHD via sparse predictions, while PCMCI+ and DYNOTEARS partially time out at $D \geq 20$.

4.4 TimeGraph Benchmark: Nonlinear Gaussian and Non-Stationary SCMs

We report TimeGraph (Ferdous et al., 2025) results at D=6, $\tau_{\max}$=2 for **B1** (pure nonlinear Gaussian noise, 3 lagged edges) and **C1** (nonlinear non-stationary with polynomial trend and sinusoidal seasonality, 4 lagged edges), using 5 non-overlapping 1 000-sample windows from a T=5 000 series.

Nonlinear Gaussian (B1). On pure nonlinear Gaussian dynamics, DYNOTEARS scores F1$= 0$, while NeuralGC (0.108 ± 0.069) and PCMCI+/GPDC (0.100 ± 0.224) remain high-variance near zero. X-SPK's F1$= 0.182$ is identical across all five seeds (zero std) because it predicts the near-complete graph on every seed (SHD $= 27$ on a 3-edge, 6-node graph; Appendix A.8 quantifies this density mismatch and a mitigation) rather than achieving stable sparse recovery; its AUROC ties Rhino's mean (0.600) with lower variance (std 0.040 vs 0.162), and NeuralGC (0.503) and PCMCI+/GPDC (0.530) provide weaker, less stable ranking signal.

Non-stationary regime (C1). PCMCI+/GPDC (F1$= 0.088$, AUROC$= 0.398$) substantially recovers relative to linear PCMCI+ (AUROC$= 0.247$) through its nonparametric CI test, but remains well below X-SPK and NeuralGC. NeuralGC achieves AUROC$= 0.535$ and F1$= 0.163$ but still trails X-SPK, which leads with AUROC$= 0.606$ and F1$= 0.235$. Scaling to D=8, NeuralGC takes the lead (F1$= 0.214$, AUROC$= 0.685$) over X-SPK (F1$= 0.164$, AUROC$= 0.604$) and all linear baselines, with absolute performance degrading for all methods as dimension grows.

4.5 Confounder Scenarios: Mackey–Glass, TimeGraph, and CausalRivers

Multiple datasets include confounding scenarios. Mackey–Glass Case A exhibits collider/ghost structure with hidden confounding (results in appendix Table 2). Both TimeGraph and CausalRivers include latent-confounder variants discussed below. For TimeGraph, we evaluate the five linear-regime methods (PCMCI+, DYNOTEARS, Rhino, VARLiNGAM, X-SPK) on the confounder variants at D=4, $\tau_{\max}$=2: **A1C** (linear VAR + latent confounder) and **C1C** (nonlinear non-stationary + latent confounder). The latent variable U drives two observed variables contemporaneously (lag 0) and is absent from the lagged ground-truth graph, which retains the same 2-edge structure ($X_1 \to X_4$, lag 2; $X_3 \to X_2$, lag 1) as the base scenarios. CausalRivers confounders arise naturally from the hydrological system structure; these results are discussed together with the main CausalRivers findings above.

A1C. In the linear+confounder regime, linear methods remain strongest: PCMCI+ is best (F1$=0.787$, AUROC$= 0.929$), with Rhino second (F1$= 0.701$). X-SPK stays at the degenerate all-edges F1 (0.286), but AUROC$=0.621 \pm 0.203$ shows partial ranking signal.

C1C. In the nonlinear+nonstationary+confounder regime, all methods degrade sharply. X-SPK is the best ranker (AUROC 0.564, F1 0.286), Rhino minimizes SHD (4.2) via conservative edge prediction, and PCMCI+/DYNOTEARS/VARLiNGAM are near-random or worse in AUROC. This confounder stress test is valuable because it separates ranking robustness (X-SPK) from threshold calibration (Rhino) under latent confounding.

Detailed robustness analyses (scalability, hyperparameter sweeps, runtime, and component ablation) are provided in Appendix A. Component ablation (Table 9, Appendix A.9) shows that the STDP trace (γ) is the single critical component: removing it collapses AUROC to 0.5 on every dataset.

Table 1: Benchmark means over seeds (s.d. shown for ≥ 0.01); bold marks the best learned method. Nonlinear substitutions: NeuralGC on Lorenz/CausalRivers/all TG rows; PCMCI+/GPDC on TG B1, B2 $D{=}6$, C1 $D{=}6$, B2 $D{=}8$ (B-type, $D{\leq}8$); CCM only for $D{\leq}5$. TG B2C NeuralGC uses $n{=}3$. **Null**: AUROC$=0.5$, AUPRC$=|E|/D^2$, F1$=0$, SHD$=\binom{D}{2}$. X-SPK is shown under two readouts of one shared ranking (AUROC/AUPRC identical per seed): fixed global τ, and *perm. null* (70th percentile of a circular-shift surrogate null, $K{=}10$; quantile tuned on a held-out pool excluding Tigramite Benchmark). Runtime in Appendix A.3.

Dataset / Method		F1↑	AUROC↑	AUPRC↑	SHD↓
Lorenz $D{=}3$ ($D{=}3,\ \lvert E\rvert{\approx}5$)	PCMCI+/GPDC	$0.732_{\pm 0.143}$	$0.190_{\pm 0.032}$	0.456	$2.1_{\pm 1.0}$
	DYNOTEARS	0.333	0.150	0.494	4.0
	Rhino	0.333	0.600	0.644	4.0
	NeuralGC	$0.663_{\pm 0.100}$	$0.940_{\pm 0.046}$	$0.956_{\pm 0.037}$	$2.6_{\pm 0.7}$
	CCM	$0.600_{\pm 0.086}$	$0.780_{\pm 0.082}$	$0.822_{\pm 0.063}$	$3.4_{\pm 0.5}$
	X-SPK (fixed τ)	**0.909**	**1.000**	**1.000**	**1.0**
	X-SPK (perm. null)	0.750	**1.000**	**1.000**	2.0
CausalRivers $D{=}5$ ($D{=}5,\ \lvert E\rvert{\approx}4$)	PCMCI+/GPDC	$0.185_{\pm 0.143}$	$0.441_{\pm 0.057}$	$0.181_{\pm 0.029}$	$6.0_{\pm 2.1}$
	DYNOTEARS	$0.227_{\pm 0.157}$	$0.466_{\pm 0.110}$	$0.209_{\pm 0.078}$	$9.3_{\pm 4.5}$
	Rhino	$0.114_{\pm 0.201}$	$0.598_{\pm 0.090}$	$0.308_{\pm 0.137}$	$4.4_{\pm 0.9}$
	NeuralGC	$0.237_{\pm 0.137}$	$0.570_{\pm 0.133}$	$0.264_{\pm 0.107}$	$9.2_{\pm 2.5}$
	CCM	$0.274_{\pm 0.166}$	$0.605_{\pm 0.129}$	$0.298_{\pm 0.147}$	$8.7_{\pm 2.7}$
	X-SPK (fixed τ)	$0.554_{\pm 0.054}$	$\mathbf{0.884_{\pm 0.078}}$	$\mathbf{0.649_{\pm 0.188}}$	$6.3_{\pm 0.9}$
	X-SPK (perm. null)	$\mathbf{0.616_{\pm 0.107}}$	$\mathbf{0.884_{\pm 0.078}}$	$\mathbf{0.649_{\pm 0.188}}$	$\mathbf{4.3_{\pm 1.3}}$
CausalRivers $D{=}10$ ($D{=}10,\ \lvert E\rvert{\approx}9$)	PCMCI+	$0.167_{\pm 0.068}$	$0.510_{\pm 0.062}$	$0.146_{\pm 0.050}$	$21.3_{\pm 5.0}$
	DYNOTEARS	$0.185_{\pm 0.048}$	$0.522_{\pm 0.063}$	$0.159_{\pm 0.073}$	$40.3_{\pm 15.6}$
	Rhino	$0.114_{\pm 0.082}$	$0.583_{\pm 0.062}$	$0.190_{\pm 0.048}$	$\mathbf{10.8_{\pm 2.3}}$
	NeuralGC	$0.118_{\pm 0.064}$	$0.529_{\pm 0.082}$	$0.134_{\pm 0.049}$	$41.1_{\pm 5.1}$
	X-SPK (fixed τ)	$0.323_{\pm 0.023}$	$\mathbf{0.807_{\pm 0.070}}$	$\mathbf{0.329_{\pm 0.137}}$	$36.9_{\pm 1.4}$
	X-SPK (perm. null)	$\mathbf{0.331_{\pm 0.109}}$	$\mathbf{0.807_{\pm 0.070}}$	$\mathbf{0.329_{\pm 0.137}}$	$23.3_{\pm 4.6}$
TG B1 nonlin. Gauss. $D{=}6$ ($D{=}6,\ \lvert E\rvert{\approx}3$)	PCMCI+/GPDC	$0.100_{\pm 0.224}$	$0.530_{\pm 0.077}$	$0.144_{\pm 0.137}$	$\mathbf{3.0_{\pm 0.7}}$
	DYNOTEARS	0.000	0.500	0.083	**3.0**
	Rhino	0.000	$\mathbf{0.600_{\pm 0.162}}$	$\mathbf{0.254_{\pm 0.129}}$	**3.0**
	NeuralGC	$0.108_{\pm 0.069}$	$0.503_{\pm 0.097}$	$0.165_{\pm 0.138}$	$16.2_{\pm 1.5}$
	X-SPK (fixed τ)	0.182	$\mathbf{0.600_{\pm 0.040}}$	$0.217_{\pm 0.115}$	27.0
	X-SPK (perm. null)	$\mathbf{0.211_{\pm 0.046}}$	$\mathbf{0.600_{\pm 0.040}}$	$0.217_{\pm 0.115}$	$7.8_{\pm 1.9}$
TG C1 nonlin.+nonstat. $D{=}6$ ($D{=}6,\ \lvert E\rvert{\approx}4$)	PCMCI+/GPDC	$0.088_{\pm 0.106}$	$0.398_{\pm 0.109}$	$0.122_{\pm 0.017}$	$13.2_{\pm 1.5}$
	DYNOTEARS	$0.193_{\pm 0.068}$	$0.402_{\pm 0.051}$	0.108	$11.6_{\pm 0.9}$
	Rhino	$0.227_{\pm 0.075}$	$0.566_{\pm 0.052}$	$0.151_{\pm 0.011}$	$\mathbf{9.4_{\pm 1.1}}$
	NeuralGC	$0.163_{\pm 0.069}$	$0.535_{\pm 0.103}$	$0.148_{\pm 0.031}$	$15.8_{\pm 2.4}$
	X-SPK (fixed τ)	**0.235**	$\mathbf{0.606_{\pm 0.063}}$	$\mathbf{0.200_{\pm 0.058}}$	26.0
	X-SPK (perm. null)	$0.151_{\pm 0.102}$	$\mathbf{0.606_{\pm 0.063}}$	$\mathbf{0.200_{\pm 0.058}}$	$12.6_{\pm 1.1}$
TG B2 mixed noise $D{=}6$ ($D{=}6,\ \lvert E\rvert{\approx}3$)	PCMCI+/GPDC	$0.080_{\pm 0.179}$	$0.526_{\pm 0.068}$	$0.111_{\pm 0.062}$	$\mathbf{3.2_{\pm 0.4}}$
	DYNOTEARS	0.000	$0.494_{\pm 0.014}$	0.083	$3.4_{\pm 0.9}$
	Rhino	0.000	$0.391_{\pm 0.100}$	$0.091_{\pm 0.017}$	$\mathbf{3.2_{\pm 0.4}}$
	NeuralGC	$0.110_{\pm 0.136}$	$0.566_{\pm 0.106}$	$0.137_{\pm 0.043}$	$15.4_{\pm 2.3}$
	X-SPK (fixed τ)	0.182	$\mathbf{0.669_{\pm 0.068}}$	$\mathbf{0.178_{\pm 0.042}}$	27.0
	X-SPK (perm. null)	$\mathbf{0.198_{\pm 0.020}}$	$\mathbf{0.669_{\pm 0.068}}$	$\mathbf{0.178_{\pm 0.042}}$	$8.2_{\pm 1.1}$
TG C2 trend+seasonal $D{=}6$ ($D{=}6,\ \lvert E\rvert{\approx}3$)	PCMCI+	$0.180_{\pm 0.249}$	$0.433_{\pm 0.107}$	$0.093_{\pm 0.013}$	$\mathbf{5.4_{\pm 3.8}}$
	DYNOTEARS	$0.173_{\pm 0.031}$	$0.394_{\pm 0.026}$	0.080	$11.2_{\pm 1.9}$
	Rhino	$0.144_{\pm 0.132}$	$0.587_{\pm 0.124}$	$0.213_{\pm 0.139}$	$8.4_{\pm 3.0}$
	NeuralGC	$0.131_{\pm 0.051}$	$0.632_{\pm 0.033}$	$0.179_{\pm 0.048}$	$16.0_{\pm 1.4}$
	X-SPK (fixed τ)	0.182	$\mathbf{0.788_{\pm 0.057}}$	$\mathbf{0.329_{\pm 0.107}}$	27.0
	X-SPK (perm. null)	$\mathbf{0.255_{\pm 0.078}}$	$\mathbf{0.788_{\pm 0.057}}$	$\mathbf{0.329_{\pm 0.107}}$	$8.0_{\pm 1.0}$
TG B2C conf.+noise $D{=}6$ ($D{=}6,\ \lvert E\rvert{\approx}3$)	PCMCI+	0.000	$0.491_{\pm 0.013}$	0.083	$3.6_{\pm 0.9}$
	DYNOTEARS	0.000	$0.494_{\pm 0.013}$	0.083	$3.4_{\pm 0.9}$
	Rhino	0.000	$0.506_{\pm 0.127}$	$0.127_{\pm 0.063}$	$\mathbf{3.2_{\pm 0.4}}$
	NeuralGC	$\mathbf{0.228_{\pm 0.097}}$	$\mathbf{0.751_{\pm 0.178}}$	$\mathbf{0.313_{\pm 0.160}}$	$13.3_{\pm 2.1}$
	X-SPK (fixed τ)	0.182	$0.626_{\pm 0.097}$	$0.152_{\pm 0.038}$	27.0
	X-SPK (perm. null)	$0.137_{\pm 0.130}$	$0.626_{\pm 0.097}$	$0.152_{\pm 0.038}$	$8.8_{\pm 0.8}$
TG B2 mixed noise $D{=}8$ ($D{=}8,\ \lvert E\rvert{\approx}4$)	PCMCI+/GPDC	0.000	$0.495_{\pm 0.007}$	0.062	$4.6_{\pm 0.9}$
	DYNOTEARS	$0.067_{\pm 0.149}$	$0.523_{\pm 0.051}$	$0.084_{\pm 0.049}$	4.0
	Rhino	$0.080_{\pm 0.179}$	$0.490_{\pm 0.094}$	$\mathbf{0.121_{\pm 0.101}}$	$\mathbf{3.8_{\pm 0.4}}$
	NeuralGC	$\mathbf{0.146_{\pm 0.077}}$	$\mathbf{0.615_{\pm 0.175}}$	$0.117_{\pm 0.049}$	$27.6_{\pm 2.2}$
	X-SPK (fixed τ)	0.133	$0.594_{\pm 0.097}$	$0.117_{\pm 0.051}$	52.0
	X-SPK (perm. null)	$0.139_{\pm 0.054}$	$0.594_{\pm 0.097}$	$0.117_{\pm 0.051}$	$15.0_{\pm 1.6}$
TG C1 nonlin.+nonstat. $D{=}8$ ($D{=}8,\ \lvert E\rvert{\approx}5$)	PCMCI+	$0.048_{\pm 0.066}$	$0.369_{\pm 0.055}$	0.076	$16.0_{\pm 2.8}$
	DYNOTEARS	$0.132_{\pm 0.046}$	$0.426_{\pm 0.040}$	0.077	$18.2_{\pm 1.3}$
	Rhino	$0.159_{\pm 0.059}$	$0.523_{\pm 0.047}$	$0.098_{\pm 0.010}$	$\mathbf{14.8_{\pm 1.6}}$
	NeuralGC	$\mathbf{0.214_{\pm 0.021}}$	$\mathbf{0.685_{\pm 0.080}}$	$\mathbf{0.171_{\pm 0.035}}$	$27.8_{\pm 0.4}$
	X-SPK (fixed τ)	0.164	$0.604_{\pm 0.035}$	$0.163_{\pm 0.049}$	51.0
	X-SPK (perm. null)	$0.187_{\pm 0.020}$	$0.604_{\pm 0.035}$	$0.163_{\pm 0.049}$	$17.6_{\pm 2.1}$

5 Discussion

Mechanistic perspective. Independent per-variable reservoirs are a practical alternative to lag-augmented causal discovery: temporal complexity is absorbed into each node's own state space rather than an inflated graph, recovering a regular DAG without the combinatorial growth delay-augmented representations incur. Keeping reservoirs separate also aligns with Independent Causal Mechanisms

(Schölkopf et al., 2021), representing each channel's dynamics as an isolated unit before cross-channel competition. The geometric mechanism itself — why linear-Gaussian inputs collapse the directional footprint that nonlinear trajectory bending otherwise creates — is given in full in Section 3.

Handling of (latent) confounders. Architecturally, each reservoir encodes only its own channel, so a latent confounder can affect two channels only through observed temporal traces. Algorithmically, WTA-STDP is directional in time: near-synchronous co-driving gives symmetric trace advantages (neither direction dominates), while true directed influence produces repeated lead-lag asymmetry and a larger W_{ij}/W_{jj}. Within each main sweep, the two-pass refinement's first pass (Section 2) lets the self-gate absorb self-predictability before cross-links are evaluated in Pass-2 — distinct from the separate, optional self-gate *pretraining* stage (n_{pre}), which is off (n_{pre}=0) in `UNIFIED_CONFIG` and so plays no role in this robustness. Domain priors (e.g., elevation masks in CausalRivers) further remove physically impossible links and reduce confounded alternatives.

This robustness is strongest for contemporaneous or near-synchronous confounding (our A1C/C1C experiments, Section 4.5). A confounder with systematically different delays to each child can itself induce a stable lead-lag asymmetry that WTA-STDP may score as a direct edge — a failure mode our experiments do not probe (Appendix A.4); Section 3's scope statement already excludes such confounded pairs from the theory's guarantees, so this mitigation is empirical, not proven.

Limitations and scope. X-SPK is weakest in linear-VAR regimes, where causal direction leaves only a weak timing footprint and competition ratios can saturate above threshold, producing near-complete-graph readouts rather than well-calibrated ones. Nonlinear causes push effect trajectories into spike-prone regions earlier and more consistently, so true parents repeatedly win WTA competition; near-linear systems offer few such landmarks to separate true from spurious precedence, so methods like PCMCI+ and VARLiNGAM retain their structural advantage by modeling linear-Gaussian covariance directly.

On TimeGraph A1, X-SPK still shows ranking signal (AUROC $= 0.769$), but the default threshold can turn that ranking into an overly dense graph, so Appendix A.8 presents threshold recalibration only as a partial mitigation. A related consequence appears in Structural Hamming Distance (SHD) on sparse graphs: the globally-tuned θ_{comp}=0.383 is calibrated for denser nonlinear datasets and over-fires when the true graph is sparse, producing many false-positive edges that inflate SHD even when AUROC—a threshold-free ranking metric—remains competitive (see Appendix A.8). Table 1 reports a data-driven alternative requiring no density prior: thresholding at the 70th-percentile of a circular-shift permutation null (K=10 surrogates; quantile selected on a held-out pool, Appendix A.4). This improves SHD on 9 of 10 main-table rows (Lorenz, already near-optimal at SHD$= 1$, is the exception) without changing AUROC/AUPRC, since both readouts share one ranking — but F1 improves on only 7 of 10 rows and worsens on the other 3 (TG C1, TG B2C, Lorenz), so permutation-null calibration is a genuine alternative operating point, not a uniform improvement over the fixed global threshold.

6 Conclusion

Core result. X-SPK shows that causal direction can be recovered from local spike-timing competitions when temporal structure is lifted into high-dimensional reservoir geometry. Across nonlinear, non-stationary, and chaotic benchmarks, this mechanism yields strong ranking performance (AUROC/AUPRC) without supervision, gradient training, or stationarity-dependent conditional-independence tests, beating an untuned lead-lag control (Appendix A.10); thresholded recovery (F1/SHD) requires density-calibrating the threshold (Appendix A.8).

Biological interpretation. X-SPK's two mechanisms each have a recognizable cortical analogue. Winner-take-all competition via lateral inhibition is a standard model of local cortical microcircuit computation, and STDP is its literal, not merely analogous, synaptic learning rule here. Less obviously, the self-gate's routine participation in that competition mirrors a recognized role for self-generated expectation: cortical circuits are widely thought to evaluate incoming signals against a prediction formed from their own recent history, crediting new input only to the extent it improves on that baseline (Friston, 2010; Schultz et al., 1997). Because a genuine external driver arrives asynchronously relative to the target's own rhythm, it still wins a share of these comparisons over time even though self dominates any single instant (Appendix B, Assumption B.1).

AI USE STATEMENT

We used generative AI (Claude, Anthropic) throughout this project's research process, under the author's direction and with every contribution independently reviewed before incorporation. This included: (i) designing and implementing diagnostic experiments and analysis code; (ii) helping formalize theoretical arguments the author originated into precise mathematical statements and proofs (Appendix B); (iii) running and interpreting results; and (iv) manuscript editing, table/reference formatting, and citation and license verification. We did not use generative AI to design the core X-SPK architecture or to originate the hypotheses underlying this paper's theoretical claims. All AI-assisted work was reviewed, tested, and independently checked by the author before being incorporated, and we take responsibility for the final content of this work.

ETHICS STATEMENT

This work introduces a computational method for causal discovery from time-series data and does not involve human subjects, human-subjects data, or personally identifiable information. All benchmark datasets are either fully synthetic (Lorenz, Mackey–Glass, Tigramite, TimeGraph) or a public real-world hydrological dataset (CausalRivers, river-gauge discharge measurements with no personal data); dataset sources, licenses, and usage terms are documented in Appendix B.3. We are not aware of a direct pathway from this work to harmful applications; the generic dual-use risks of any causal-inference method applied to sensitive domains are not specific to this contribution, and we do not target or evaluate such applications here. The authors declare no conflicts of interest or sponsorship relevant to this work.

REPRODUCIBILITY STATEMENT

We report full experimental protocols (datasets, preprocessing, baselines, seeds, metrics) in Section 4 and Appendix A, the complete hyperparameter configuration used for every reported result (`UNIFIED_CONFIG`) in Table 6, and the full assumptions and derivations for every theoretical result in Appendix B. An anonymized code repository accompanies this submission as supplementary material, including the exact scripts that produced every table and figure, the frozen configuration files, and a document (`RUNNING_EXPERIMENTS.md`) that gives verified environment setup and the ordered run commands needed to reproduce each result end to end, including the post-hoc density-calibration, acyclicity, and lead-lag-control diagnostics introduced in Appendices A.8 and A.10, and the self-vs-cause win-rate check underlying Assumption B.1 (Appendix B). Per-asset dataset licenses and redistribution terms are listed in Appendix B.3.

REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiro Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2623–2631, 2019. URL https://arxiv.org/abs/1907.10902.

Joschka Boedecker, Oliver Obst, Joseph T Lizier, N Michael Mayer, and Minoru Asada. Information processing in echo state networks at the edge of chaos. *Theory in Biosciences*, 131(3):205–213, 2012. doi: 10.1007/s12064-011-0146-8.

Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. doi: 10.1109/4235.996017.

Muhammad Hasan Ferdous, Emam Hossain, and Md Osman Gani. Timegraph: Synthetic benchmark datasets for robust time-series causal discovery. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 5425–5435, 2025. doi: 10.1145/3711896.3737439. URL https://arxiv.org/abs/2506.01361.

Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2): 127–138, 2010. doi: 10.1038/nrn2787.

Wulfram Gerstner, Richard Kempter, J Leo van Hemmen, and Hermann Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383:76–81, 1996. doi: 10.1038/383076a0.

Matthieu Gilson, Anthony N Burkitt, David B Grayden, Doreen A Thomas, and J Leo van Hemmen. Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks. i. input selectivity–strengthening correlated input pathways. *Biological Cybernetics*, 101 (2):81–102, 2009. doi: 10.1007/s00422-009-0319-4.

Wenbo Gong, Joel Jennings, Cheng Zhang, and Nick Pawlowski. Rhino: Deep causal temporal relationship learning with history-dependent noise. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023. URL https://arxiv.org/abs/2210.14706.

Peter D. Grünwald. *The Minimum Description Length Principle*. MIT Press, Cambridge, MA, 2007. URL https://direct.mit.edu/books/monograph/3813/The-Minimum-Description-Length-Principle.

Biwei Huang, Kun Zhang, Jiji Zhang, Joseph Ramsey, Ruben Sanchez-Romero, Clark Glymour, and Bernhard Schölkopf. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research*, 21(89):1–53, 2020. URL https://www.jmlr.org/papers/v21/19-232.html.

Aapo Hyvärinen, Kun Zhang, Shohei Shimizu, and Patrik O Hoyer. Estimation of a structural vector autoregression model using non-gaussianity. *Journal of Machine Learning Research*, 11:1709–1731, 2010. URL https://www.jmlr.org/papers/v11/hyvarinen10a.html.

Eugene M Izhikevich. Polychronization: Computation with spikes. *Neural Computation*, 18(2): 245–282, 2006. doi: 10.1162/089976606775093882.

Herbert Jaeger. The "echo state" approach to analysing and training recurrent neural networks. *GMD Report 148, German National Research Center for Information Technology*, 2001. URL https://www.ai.rug.nl/minds/uploads/EchoStatesTechRep.pdf.

Dominik Janzing and Bernhard Schölkopf. Causal inference using the algorithmic markov condition. *IEEE Transactions on Information Theory*, 56(10):5168–5194, 2010. doi: 10.1109/TIT.2010.2060095. URL https://arxiv.org/abs/0804.3678.

Anna Krakovská. Cross-predictions in the search for effective connectivity in brain. In *2021 13th International Conference on Measurement*, pp. 2–5. IEEE, 2021. URL https://ieeexplore.ieee.org/document/9446774.

Benjamin James Lansdell and Konrad Paul Kording. Neural spiking for causal inference and learning. *PLOS Computational Biology*, 19(4):e1011005, 2023. doi: 10.1371/journal.pcbi.1011005. Preprint: bioRxiv 253351 (2018).

Edward N Lorenz. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2):130–141, 1963. URL https://journals.ametsoc.org/view/journals/atsc/20/2/1520-0469_1963_020_0130_dnf_2_0_co_2.xml.

Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560, 2002. URL https://direct.mit.edu/neco/article/14/11/2531/6650/Real-Time-Computing-Without-Stable-States-A-New.

Michael C Mackey and Leon Glass. Oscillation and chaos in physiological control systems. *Science*, 197(4300):287–289, 1977. doi: 10.1126/science.267326.

James L Massey. Causality, feedback and directed information. In *Proceedings of the 1990 International Symposium on Information Theory and Its Applications (ISITA-90)*, pp. 303–305, Waikiki, HI, 1990. URL https://www.isiweb.ee.ethz.ch/archive/massey_pub/pdf/BI532.pdf.

Joris M Mooij, Jonas Peters, Dominik Janzing, Jakob Zscheischler, and Bernhard Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research*, 17:1–102, 2016. URL https://jmlr.org/papers/v17/14-518.html.

Roxana Pamfil, Nisara Sriwattanaworachai, Shaan Desai, Philip Pilgerstorfer, Konstantinos Georgatzis, Paul Beaumont, and Bryon Aragam. DYNOTEARS: Structure learning from time-series data. *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020. URL https://proceedings.mlr.press/v108/pamfil20a.html.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017. URL https://mitpress.mit.edu/9780262037310/elements-of-causal-inference/.

Jakob Runge. Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series datasets. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 1388–1397, 2020. URL https://proceedings.mlr.press/v124/runge20a.html.

Jakob Runge, Peter Nowack, Marlene Kretschmer, Seth Flaxman, and Dino Sejdinovic. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5 (11):eaau4996, 2019. doi: 10.1126/sciadv.aau4996.

Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. doi: 10.1109/JPROC.2021.3058954. URL https://arxiv.org/abs/2102.11107.

Benjamin Schrauwen, David Verstraeten, and Jan Van Campenhout. An overview of reservoir computing: theory, applications and implementations. *Proceedings of the European Symposium on Artificial Neural Networks*, 2007. URL https://www.esann.org/sites/default/files/proceedings/legacy/es2007-8.pdf.

Thomas Schreiber. Measuring information transfer. *Physical Review Letters*, 85(2):461–464, 2000. doi: 10.1103/PhysRevLett.85.461. URL https://arxiv.org/abs/nlin/0001042.

Wolfram Schultz, Peter Dayan, and P Read Montague. A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599, 1997. doi: 10.1126/science.275.5306.1593.

Stephen M Smith, Karla L Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for FMRI. *NeuroImage*, 54(2):875–891, 2011. doi: 10.1016/j.neuroimage.2010.08.063.

Sen Song, Kenneth D Miller, and L F Abbott. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3(9):919–926, 2000. URL https://www.nature.com/articles/nn0900_919.

Gideon Stein, Maha Shadaydeh, Jan Blunk, Niklas Penzel, and Joachim Denzler. Causalrivers – scaling up benchmarking of causal discovery for real-world time-series. In *International Conference on Learning Representations (ICLR)*, 2025. URL https://arxiv.org/abs/2503.17452.

George Sugihara, Robert May, Hao Ye, Chih-hao Hsieh, Ethan Deyle, Michael Fogarty, and Stephan Munch. Detecting causality in complex ecosystems. *Science*, 338(6106):496–500, 2012. doi: 10.1126/science.1227079.

Alex Tank, Ian Covert, Nicholas Foti, Ali Shojaie, and Emily B. Fox. Neural granger causality. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4267–4279, 2021. doi: 10.1109/TPAMI.2021.3065601. URL https://arxiv.org/abs/1802.05842.

Matthew J. Vowels, Necati Cihan Camgoz, and Richard Bowden. Shadow-mapping for unsupervised neural causal discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 1740–1743, 2021. doi: 10.1109/CVPRW53098.2021.00190. URL https://arxiv.org/abs/2104.08183.

Jintong Zhao, Zhongxue Gan, Ruixi Huang, Chun Guan, Jifan Shi, and Siyang Leng. Detecting dynamical causality via intervened reservoir computing. *Communications Physics*, 7:232, 2024. doi: 10.1038/s42005-024-01730-6.

A Additional Experimental Details

A.1 Mackey–Glass Pathological Benchmarks

Table 2 reports results for all five Mackey–Glass $D{=}5$ topologies using nonlinear baselines (PCMCI+/GPDC, NeuralGC, CCM; 10 seeds each) *at library-default settings and* $\tau_{\max}{=}2$, the protocol used throughout the rest of this paper. Under that protocol, X-SPK leads F1 and AUROC on four of five topologies; on step functions, CCM marginally leads AUROC (0.672 vs 0.650) while X-SPK leads F1 (0.519 vs 0.436) and AUPRC. NeuralGC is the strongest nonlinear baseline overall (F1 0.331–0.533, AUROC 0.596–0.713), well above PCMCI+/GPDC (AUROC 0.315–0.377, F1≈ 0–0.151). CCM achieves the lowest SHD on gated and step-functions topologies; NeuralGC minimizes SHD on collider/ghost and dense hub. **This ranking is protocol-dependent, not fixed.** Appendix A.4 reports two further diagnostics under equal conditions: given the same 150-trial tuning budget X-SPK receives, NeuralGC (ensembled) beats X-SPK outright on three of these five topologies (collider+ghost, gated nonlinear, step functions); and given $\tau_{\max}{\geq}5$ instead of the default 2, NeuralGC beats X-SPK on dense hub as well. Only `mg_regime_switching` remains an uncontested X-SPK win across all conditions tested.

Table 2: Mackey–Glass pathological benchmarks ($D{=}5$, 10 seeds): all five nonlinear topologies. Null row omitted; reference values: AUROC$=0.5$, AUPRC=edge density, F1$=0$, SHD$=\binom{D}{2}$. Notation as Table 1.

Dataset / Method		F1↑	AUROC↑	AUPRC↑	SHD↓
Mackey-Glass pathological scenarios ($D{=}5$ nonlinear baselines)					
MG collider ghost $D{=}5$ ($D{=}5,\ \lvert E\rvert{\approx}9$)	PCMCI+/GPDC	0.000	$0.331_{\pm 0.016}$	0.360	$9.4_{\pm 0.5}$
	NeuralGC	$0.533_{\pm 0.131}$	$0.683_{\pm 0.082}$	$0.535_{\pm 0.090}$	$\mathbf{9.3}_{\pm \mathbf{2.8}}$
	CCM	$0.332_{\pm 0.121}$	$0.635_{\pm 0.075}$	$0.499_{\pm 0.072}$	$10.1_{\pm 1.3}$
	X-SPK	**0.621**	$\mathbf{0.758}_{\pm \mathbf{0.057}}$	$\mathbf{0.640}_{\pm \mathbf{0.085}}$	11.0
MG regime switching $D{=}5$ ($D{=}5,\ \lvert E\rvert{\approx}10$)	PCMCI+/GPDC	$0.034_{\pm 0.071}$	$0.327_{\pm 0.021}$	$0.395_{\pm 0.010}$	$10.2_{\pm 0.6}$
	NeuralGC	$0.513_{\pm 0.112}$	$0.697_{\pm 0.076}$	$0.610_{\pm 0.115}$	$9.9_{\pm 2.3}$
	CCM	$0.432_{\pm 0.108}$	$0.617_{\pm 0.078}$	$0.573_{\pm 0.091}$	$\mathbf{9.1}_{\pm \mathbf{1.6}}$
	X-SPK	**0.667**	$\mathbf{0.791}_{\pm \mathbf{0.024}}$	$\mathbf{0.744}_{\pm \mathbf{0.046}}$	10.0
MG dense hub $D{=}5$ ($D{=}5,\ \lvert E\rvert{\approx}10$)	PCMCI+/GPDC	$0.036_{\pm 0.077}$	$0.330_{\pm 0.029}$	0.395	$10.1_{\pm 0.9}$
	NeuralGC	$0.530_{\pm 0.115}$	$0.713_{\pm 0.090}$	$0.618_{\pm 0.100}$	$\mathbf{9.1}_{\pm \mathbf{2.1}}$
	CCM	$0.256_{\pm 0.139}$	$0.541_{\pm 0.037}$	$0.456_{\pm 0.059}$	$11.7_{\pm 1.6}$
	X-SPK	**0.667**	$\mathbf{0.770}_{\pm \mathbf{0.062}}$	$\mathbf{0.697}_{\pm \mathbf{0.069}}$	10.0
MG gated nonlinearity $D{=}5$ ($D{=}5,\ \lvert E\rvert{\approx}8$)	PCMCI+/GPDC	$0.151_{\pm 0.136}$	$0.315_{\pm 0.040}$	$0.304_{\pm 0.015}$	$10.0_{\pm 1.1}$
	NeuralGC	$0.331_{\pm 0.119}$	$0.596_{\pm 0.064}$	$0.409_{\pm 0.058}$	$11.7_{\pm 1.5}$
	CCM	$0.390_{\pm 0.068}$	$0.567_{\pm 0.050}$	$0.379_{\pm 0.032}$	$\mathbf{9.3}_{\pm \mathbf{1.2}}$
	X-SPK	**0.571**	$\mathbf{0.695}_{\pm \mathbf{0.047}}$	$\mathbf{0.488}_{\pm \mathbf{0.050}}$	12.0
MG step functions $D{=}5$ ($D{=}5,\ \lvert E\rvert{\approx}7$)	PCMCI+/GPDC	$0.149_{\pm 0.165}$	$0.377_{\pm 0.039}$	$0.270_{\pm 0.011}$	$\mathbf{7.0}_{\pm \mathbf{0.8}}$
	NeuralGC	$0.425_{\pm 0.119}$	$0.656_{\pm 0.088}$	$0.479_{\pm 0.106}$	$9.5_{\pm 2.2}$
	CCM	$0.436_{\pm 0.085}$	$\mathbf{0.672}_{\pm \mathbf{0.070}}$	$0.460_{\pm 0.108}$	$9.1_{\pm 1.8}$
	X-SPK	**0.519**	$0.650_{\pm 0.065}$	$\mathbf{0.519}_{\pm \mathbf{0.111}}$	13.0

A.2 CausalRivers Scalability: $D \in \{15, 20, 25, 30\}$

Table 3 shows the full five-method comparison at $D \in \{15, 20, 25, 30\}$. PCMCI+ and DYNOTEARS complete all 5 seeds through $D{=}20$; at $D{=}25$, PCMCI+ times out on 1/5 seeds; at $D{=}30$, PCMCI+ and DYNOTEARS each time out on 1/5 seeds within the 600 s budget. X-SPK maintains the best AUROC across all sizes ($0.789 \to 0.771 \to 0.785 \to 0.780$), with a consistent $\approx$0.15–0.20 margin over the next-best method. X-SPK's SHD grows steeply with D because the unified threshold θ_{comp} is calibrated for smaller graphs; per-size recalibration is a natural extension. Rhino achieves the lowest SHD at all sizes via conservative sparse-linear regression, at the cost of lower AUROC.

A.3 Runtime Scaling with Graph Size

Table 4 reports mean wall-clock time per run across the CausalRivers scalability benchmark. All methods are single-threaded on the same machine (Intel Core i7, 16 GB RAM, no GPU).

NeuralGC is the fastest nonlinear baseline at 5.0–10.6 s, growing modestly with D (one MLP per target, fixed architecture). CCM requires 224 s at $D{=}5$ and is omitted for $D{>}5$ due to exponential manifold-search cost; PCMCI+/GPDC requires 1317 s at $D{=}5$ and is omitted for $D{>}8$. Among the linear-regime methods, the scaling profiles diverge sharply above $D{=}10$: PCMCI+ begins timing out at $D{=}25$ (1/5 seeds), while X-SPK completes all seeds in 25–41 seconds across $D \in \{10, 15, 20, 25\}$. This near-flat empirical scaling in D over the tested range is not X-SPK's asymptotic complexity, which is superlinear (reservoir computation $O(T \cdot D \cdot n_{\text{res}})$, gate-matrix update $O(T \cdot D^2)$) — at $D \leq 30$, constant per-step overhead dominates measured runtime, flattening the empirical curve well short of where the $O(D^2)$ term would visibly dominate. In contrast, PCMCI+ scales roughly as $O(D^3 \cdot T)$ due to conditional independence testing for each triple, and DYNOTEARS' continuous optimization has $O(D^3)$ matrix operations per gradient step. Rhino completes in $<$1 s at all sizes because RhinoLite uses Lasso VAR without recurrent inference. At $D{=}25$, PCMCI+ is the only method that fails to complete all 5 seeds within the time budget; X-SPK's precision degrades severely despite completing every run (Section A.2).

A.4 Methodological Parity and Baseline Fidelity

RhinoLite is not the full Rhino model. Throughout this paper, results attributed to "Rhino" are produced by **RhinoLite**, a linear Lasso-VAR approximation without recurrent inference (Appendix B.3), not Gong et al.'s (2023) full deep generative, history-dependent-noise model. This substitution likely understates Rhino's true capability on the nonlinear, chaotic benchmarks (Lorenz, Mackey–Glass) reported here; readers should not read RhinoLite's results as the full model's ceiling.

Tuning-budget asymmetry. `UNIFIED_CONFIG` was selected via a 150-trial, 13-parameter NSGA-II Optuna sweep (Section A.8); every baseline is run with library defaults, never retuned. A 14th X-SPK parameter — the permutation-null significance quantile used for the alternative readout in Table 1 — was subsequently grid-searched over $\{65, 70, \ldots, 99\}$ on a held-out tuning pool (CausalRivers $n{=}6$ + Tigramite static $D{=}3$, excluding Tigramite Benchmark to prevent leakage into that diagnostic's own evaluation set), the same held-out-pool discipline as `UNIFIED_CONFIG` itself. This asymmetry is disclosed in Section 4.2 as a deliberate design decision to preserve cross-benchmark comparability and avoid coupling baselines to X-SPK's own model-selection protocol, but it does not equalize optimization effort. To quantify how much of X-SPK's margin this explains, we retuned two gradient-sensitive baselines under the identical protocol (150-trial NSGA-II, same tuning pool) and re-evaluated at the paper's seeds. **DYNOTEARS** improves modestly (mean AUROC $0.399 \rightarrow 0.439$) but remains far below X-SPK (0.730); its improved SHD under tuning is a sparsity

Table 3: CausalRivers scalability at $D \in \{15, 20, 25, 30\}$ (5 seeds unless noted). † 4/5 completed ($D{=}25$, PCMCI+); ‡ 4/5 completed ($D{=}30$, PCMCI+/DYNOTEARS). Notation as Table 1.

Dataset / Method		F1↑	AUROC↑	AUPRC↑	SHD↓
Large-graph scalability (CausalRivers)					
CausalRivers $D{=}15$ ($D{=}15$, $\lvert E\rvert \approx 14$)	PCMCI+	$0.221_{\pm 0.101}$	$0.591_{\pm 0.092}$	$0.146_{\pm 0.058}$	$36.0_{\pm 11.8}$
	DYNOTEARS	$0.174_{\pm 0.045}$	$0.557_{\pm 0.053}$	$0.107_{\pm 0.033}$	$83.6_{\pm 46.3}$
	Rhino	$0.130_{\pm 0.088}$	$0.619_{\pm 0.046}$	$0.126_{\pm 0.033}$	$\mathbf{18.4_{\pm 6.0}}$
	VARLiNGAM	$0.137_{\pm 0.061}$	$0.532_{\pm 0.085}$	$0.112_{\pm 0.039}$	$77.4_{\pm 21.0}$
	Random/Null	0.000	0.500	0.063	105.0
	X-SPK	$\mathbf{0.232_{\pm 0.013}}$	$\mathbf{0.789_{\pm 0.056}}$	$\mathbf{0.211_{\pm 0.095}}$	$91.6_{\pm 1.3}$
CausalRivers $D{=}20$ ($D{=}20$, $\lvert E\rvert \approx 19$)	PCMCI+	$\mathbf{0.202_{\pm 0.122}}$	$0.596_{\pm 0.085}$	$\mathbf{0.120_{\pm 0.073}}$	$42.8_{\pm 10.9}$
	DYNOTEARS	$0.139_{\pm 0.029}$	$0.585_{\pm 0.057}$	$0.085_{\pm 0.048}$	$147.0_{\pm 70.0}$
	Rhino	$0.061_{\pm 0.064}$	$0.667_{\pm 0.067}$	$\mathbf{0.123_{\pm 0.036}}$	$\mathbf{23.8_{\pm 1.8}}$
	VARLiNGAM	$0.123_{\pm 0.051}$	$0.560_{\pm 0.080}$	$0.087_{\pm 0.030}$	$120.6_{\pm 45.2}$
	Random/Null	0.000	0.500	0.048	190.0
	X-SPK	$0.180_{\pm 0.012}$	$\mathbf{0.771_{\pm 0.036}}$	$\mathbf{0.124_{\pm 0.036}}$	$171.8_{\pm 1.6}$
CausalRivers $D{=}25$ ($D{=}25$, $\lvert E\rvert \approx 25$)	PCMCI+†	$0.139_{\pm 0.089}$	$0.548_{\pm 0.064}$	$0.068_{\pm 0.047}$	$50.5_{\pm 6.2}$
	DYNOTEARS	$0.130_{\pm 0.030}$	$0.599_{\pm 0.062}$	$0.055_{\pm 0.010}$	$168.4_{\pm 93.2}$
	Rhino	$0.089_{\pm 0.028}$	$0.648_{\pm 0.078}$	$\mathbf{0.113_{\pm 0.039}}$	$\mathbf{28.2_{\pm 2.5}}$
	VARLiNGAM	$0.119_{\pm 0.028}$	$0.593_{\pm 0.025}$	$0.076_{\pm 0.011}$	$161.4_{\pm 63.5}$
	Random/Null	0.000	0.500	0.040	300.0
	X-SPK	$\mathbf{0.148_{\pm 0.008}}$	$\mathbf{0.785_{\pm 0.036}}$	$0.103_{\pm 0.025}$	$276.8_{\pm 1.6}$
CausalRivers $D{=}30$ ($D{=}30$, $\lvert E\rvert \approx 29$)	PCMCI+‡	$0.138_{\pm 0.048}$	$0.560_{\pm 0.037}$	$0.058_{\pm 0.031}$	$71.0_{\pm 5.0}$
	DYNOTEARS‡	$\mathbf{0.159_{\pm 0.079}}$	$0.617_{\pm 0.043}$	$0.056_{\pm 0.016}$	$169.3_{\pm 90.8}$
	Rhino	$0.143_{\pm 0.099}$	$0.662_{\pm 0.054}$	$\mathbf{0.138_{\pm 0.050}}$	$\mathbf{32.4_{\pm 3.4}}$
	VARLiNGAM	$0.109_{\pm 0.036}$	$0.597_{\pm 0.051}$	$0.063_{\pm 0.018}$	$213.2_{\pm 64.0}$
	Random/Null	0.000	0.500	0.033	435.0
	X-SPK	$0.123_{\pm 0.006}$	$\mathbf{0.780_{\pm 0.039}}$	$0.091_{\pm 0.021}$	$406.8_{\pm 2.5}$

Table 4: Mean wall-clock time (seconds per run) on CausalRivers subgraphs. NeuralGC, CCM, and PCMCI+/GPDC measured on 10 seeds at $D \in \{5, 10\}$ (nonlinear benchmark); remaining rows from the scalability benchmark (5 seeds). Bold: fastest per D among methods with data. ‡: 1/5 seeds timed out (PCMCI+ at D=25); mean over completed seeds shown. —: not benchmarked at this size.

D	PCMCI+/GPDC	NeuralGC	CCM	PCMCI+	DYNOTEARS	Rhino	VARLiNGAM	X-SPK
5	1317	**5.0**	224	0.8	1.0	**0.0**	1.5	13.3
10	—	**10.6**	—	1.8	1.3	**0.0**	0.6	32.0
15	—	—	—	80.5	96.4	**0.1**	5.4	23.6
20	—	—	—	122.5	158.7	**0.1**	8.7	**26.6**
25	—	—	—	138.6‡	191.9	**0.2**	21.6	40.7

effect of the selected config, not higher accuracy. **NeuralGC**, ensembled at K=5 to match X-SPK's own N_{ENSEMBLE}=5 (score-averaged across 5 independent initializations; a single initialization is not representative, since the 150-trial tuning objective itself used only one), closes substantially more of the gap: mean AUROC $0.604 \rightarrow 0.674$ across the 14 datasets shared with Table 1/Table 2 (`tg_b2c_n6` excluded, not in this sweep's pool), versus X-SPK's 0.730 on the same 14. Critically, this residual gap is *not* uniform across datasets: tuned NeuralGC outright beats X-SPK on 5 of 14 — three Mackey–Glass topologies (collider+ghost: 0.785 vs. 0.758; gated nonlinear: 0.708 vs. 0.695; step functions: 0.776 vs. 0.650) and two TimeGraph nonlinear sets (B2 D=8: 0.710 vs. 0.594; C1 D=8: 0.666 vs. 0.604) — while X-SPK retains a clear, wide lead on CausalRivers (both sizes), Lorenz, `mg_regime_switching`, and two further TimeGraph rows (B2 D=6, C2 D=6, the latter by 0.503 vs. 0.788). Selection used `select_auroc`, the paper's documented Pareto-selection rule (Section A.8).

Lag-budget asymmetry. All baselines are restricted to $\tau_{\max}$=2 by default, while X-SPK's leaky reservoirs integrate a continuous, effectively unbounded history. We reran DYNOTEARS, RhinoLite, PCMCI+, and NeuralGC at $\tau_{\max} \in \{2, 5, 10\}$ (and additionally $\{20, 50\}$ on Mackey–Glass). DYNOTEARS, RhinoLite, and PCMCI+ show no meaningful benefit from the larger lag budgets tested, so X-SPK's margin over these three is not a lag-budget artifact. NeuralGC does benefit, specifically on Mackey–Glass: at $\tau_{\max} \geq 5$, NeuralGC beats X-SPK on `mg_dense_hub` in all 5 seeds we ran. Combined with the equal-tuning-budget result immediately above, Section A.1's claim that X-SPK leads all five MG topologies holds only at the baseline $\tau_{\max}$=2, default-baseline-tuning protocol used throughout the rest of this paper — under either an increased lag budget or an equalized tuning budget, X-SPK's lead over NeuralGC is not universal across MG topologies (4 of 5 are contested by one or the other), though it remains undisputed on CausalRivers and Lorenz throughout.

We report both asymmetries here, alongside the existing density-calibration and lead-lag-control diagnostics (Appendices A.8, A.10), so that the paper's ranking claims are read against a fully characterized evaluation protocol rather than an implicitly level one.

A.5 An Established External Benchmark: Netsim, and a Non-Stationarity-Aware Baseline

All datasets above are ours (TimeGraph, Mackey–Glass, Tigramite-generated) or a river-flow benchmark (CausalRivers). To test X-SPK against an independently generated, widely used benchmark, we add **Netsim** (Smith et al., 2011), simulated resting-state fMRI BOLD time series (hemodynamic-response-convolved, TR= 3s) with a known ground-truth directed network, and **CD-NOD** (Huang et al., 2020), a constraint-based causal discovery method built to detect nonstationarity/heterogeneity. We use Netsim sim1 (D=5, chain plus a skip edge) and sim2 (D=10), 10 of 50 available subjects (seeds) each, 200 timepoints/subject. Ground-truth lag is not well defined for BOLD-convolved fMRI (continuous-time neural dynamics, discretized by a 3s TR), so we evaluate directed-skeleton recovery only (does edge $i \rightarrow j$ exist, collapsed over whichever lag a method uses to detect it); lag-specific columns are not reported for this benchmark. CD-NOD is not natively identifiable on purely contemporaneous bivariate pairs (no v-structure to exploit), so we embed each variable at lags $0..\tau_{\max}$=1 as separate nodes and use `BackgroundKnowledge` to forbid present-to-past edges, restoring the same past-precedes-present identifiability that time-lagged PC/FCI variants rely on; the time index is CD-NOD's nonstationarity surrogate c.

Table 5: Netsim (Smith et al., 2011), 10 subjects/seeds, directed-skeleton recovery only (lag not evaluated; see text). Bold: best per column per dataset.

Dataset	Method	F1↑	AUROC↑	AUPRC↑	SHD↓
sim1 (D=5, $\|E\|$=5)	CD-NOD	0.000	$0.378_{\pm 0.014}$	0.200	$5.3_{\pm 0.5}$
	PCMCI+	$0.312_{\pm 0.319}$	$0.474_{\pm 0.145}$	$0.236_{\pm 0.052}$	$\mathbf{4.8_{\pm 2.1}}$
	Rhino	$0.399_{\pm 0.189}$	$\mathbf{0.873_{\pm 0.016}}$	$\mathbf{0.633_{\pm 0.026}}$	$6.5_{\pm 2.1}$
	X-SPK	**0.400**	$0.676_{\pm 0.073}$	$0.353_{\pm 0.050}$	15.0
sim2 (D=10, $\|E\|$=11)	CD-NOD	$0.064_{\pm 0.083}$	$0.463_{\pm 0.026}$	0.109	$11.3_{\pm 0.9}$
	PCMCI+	$0.244_{\pm 0.206}$	$0.521_{\pm 0.077}$	$0.128_{\pm 0.026}$	$\mathbf{10.7_{\pm 2.4}}$
	Rhino	$\mathbf{0.297_{\pm 0.158}}$	$\mathbf{0.918_{\pm 0.019}}$	$\mathbf{0.531_{\pm 0.049}}$	$21.7_{\pm 6.8}$
	X-SPK	0.218	$0.557_{\pm 0.027}$	$0.159_{\pm 0.014}$	79.0

Results match the paper's already-disclosed linear-regime limitation. RhinoLite (linear) wins clearly — expected, not an anomaly: Netsim's generative process is close to a linear SEM at the neural level, and Smith et al. (2011) themselves report regression-based methods as strongest on this benchmark. X-SPK's SHD equals $D(D-1) - |E|$ exactly on both sizes ($15=20-5$; $79=90-11$) — it predicts the complete off-diagonal graph on every seed, the same dense-graph pattern already documented for linear-regime data (Section A.8, Appendix A.8), now confirmed on an independent, real-world-derived benchmark rather than our own synthetic configurations.

CD-NOD sits at or near chance on both sizes: 200 timepoints and one context dimension leave too little genuine distributional heterogeneity for its nonstationarity test to exploit, and the past-precedes-present background knowledge we impose is conservative (CD-NOD uses its library-default α=0.05, untuned, like every other baseline). This is a negative result for CD-NOD in this specific low-context, short-window regime, not evidence against nonstationarity-aware discovery in general.

A.6 Pass-2 Refinement: Formal Definition

Section 2 describes Pass-2 refinement (`two_pass=True`) only in prose ("lower spike rate, shorter trace decay, cross-margin, self-inhibition, frozen self-gates"). We give the exact update rule here. Let $\tilde{\rho}_j(t)$ be the biased trace vector at target j's firing time t (Equation equation 4), $i^* = \arg\max_k \tilde{\rho}_{j,k}(t)$ the Pass-1 winner, and write $s := \tilde{\rho}_{j,j}(t)$ (self-trace) and $w := \tilde{\rho}_{j,i^*}(t)$ (winner's trace).

Cross-margin. A cross-channel winner ($i^* \neq j$) is only credited if it beats the self-trace by more than a fixed margin $c \geq 0$; otherwise credit reverts to self:

$$i^{**} = \begin{cases} i^* & \text{if } i^* = j \text{ or } w > s + c \\ j & \text{otherwise (reverts to self)} \end{cases} \tag{14}$$

Self-inhibition. When a genuine cross-link survives the margin test ($i^{**} \neq j$), the update magnitude is scaled continuously by how much the winner's trace exceeds self, rather than applying a fixed step:

$$\lambda = \frac{\max(0,\, w - s)}{w + \epsilon} \in [0, 1), \qquad (\lambda = 1 \text{ only in the limit } s \to -\infty) \tag{15}$$

Frozen self-gates. If $i^{**} = j$ (self wins, whether originally or via margin reversion), the diagonal gate W_{jj} receives *no* update in Pass-2 (it is frozen at its Pass-1 value); otherwise the cross-link update is applied at the scaled rate:

$$W_{i^{**},j} \leftarrow \begin{cases} W_{i^{**},j} & \text{if } i^{**} = j \text{ (frozen)} \\ W_{i^{**},j} + \lambda \cdot \eta / \sqrt{D} & \text{if } i^{**} \neq j \end{cases} \tag{16}$$

The column-wise $L2$ weight decay (Equation equation 7) is applied to $W_{\cdot,j}$ unconditionally in both cases. Pass-2 also uses a reduced target spike rate ($r_{\text{target}} \cdot$ `pass2_target_rate_mult`) and faster trace decay ($\gamma \cdot$ `pass2_trace_decay_mult`) than Pass-1, both fixed multiplicative factors on the Pass-1 values (Table 6 reports the selected multipliers).

A.7 Initialization, Pretraining, and Tie-Breaking Details

Reservoir weight initialization. Input weights $\mathbf{w}_{\text{in}}^{(i)}$ are drawn i.i.d. from $\mathcal{N}(0, 0.5^2)$; the dense entries of the recurrent matrix $\mathbf{W}_{\text{res}}^{(i)}$ are drawn i.i.d. from the standard normal $\mathcal{N}(0, 1)$ before the sparsity mask (p) and spectral-radius rescaling ($\lambda_{\max}$) of Section 2 are applied.

WTA tie-breaking. Ties go to the lowest channel index (the underlying `argmax` convention). This matters only when competing traces are exactly equal. In practice, whenever $\beta > 0$ (as in `UNIFIED_CONFIG`, β=0.081), the one genuine tie — every trace exactly zero, at the very first credited spike of each replay sweep — already resolves toward the self-slot j via the β bonus rather than an arbitrary index; every later comparison is between generically distinct accumulated real values, where exact ties occur with probability zero.

Synthetic DAG pretraining protocol ($n_{\text{syn}} > 0$). Each of the n_{syn} pretraining sequences is a fresh, randomly generated linear VAR(1) process at the same dimensionality D and (at least) the same active length as the real dataset being pretrained for: a random acyclic graph with a guaranteed confluence motif (the middle node receives two parents, matching the physical merging topology of CausalRivers) plus additional random edges (probability 0.25 per remaining upper-triangular pair), edge weights drawn uniformly at random in $[0.25, 0.50]$ for confluence parents and $[0.15, 0.40]$ otherwise (random sign), rescaled so the system's spectral radius is ≤ 0.7 for stability, and simulated as $X_t = X_{t-1}A + \varepsilon_t$ with $\varepsilon_t \sim \mathcal{N}(0, 0.5^2)$ i.i.d. noise. The gate matrix $\mathbf{W}$ (not the reservoir state) carries over from one synthetic sequence to the next and then into real-data training.

A.8 Hyperparameter Search

Single unified configuration. *All* results in this paper use one configuration, `UNIFIED_CONFIG` (Table 6), selected from a single joint Optuna sweep (Akiba et al., 2019) over a mixed pool of CausalRivers subgraphs (n=6, 4 subgraphs) and Tigramite datasets (static D=3, benchmark D=5). We use NSGA-II (Deb et al., 2002) with population 15, 150 trials, and four objectives: maximize F1, maximize AUROC, minimize SHD, maximize AUPRC. The Pareto trial with highest AUROC was selected. Table 6 lists all 13 swept parameters and their search ranges.

Fairness and leakage control. The Optuna sweep was run once and frozen; `UNIFIED_CONFIG` is then transferred zero-shot to all remaining benchmarks (Lorenz, Mackey–Glass, TimeGraph, and larger CausalRivers), with no per-dataset retuning of X-SPK. Within CausalRivers and Tigramite, tuning and final reporting use disjoint runs/subgraphs/windows, so no reported evaluation sequence appears in the sweep objective. We also intentionally avoid retuning baselines on this development split to prevent coupling them to our model-selection protocol and to preserve cross-benchmark comparability under canonical baseline defaults.

Table 6: Left: UNIFIED_CONFIG values used for all benchmarks in this paper. Right: Optuna search ranges for the 13 swept hyperparameters.

Parameter	Value	Parameter	Value
n_{res}	72	γ (trace decay)	0.688
p (sparsity)	0.159	η (WTA LR)	0.028
$\lambda_{\max}$ (spectral radius)	0.789	n_{syn}	5
α (leakage)	0.820	n_{pre}	0
T_{wash}	17	n_{epochs}	5
r_{target}	0.172	β (self-bias)	0.081
		θ_{comp}	0.383

Parameter	Range	Type
n_{res}	[20, 80]	int
p (sparsity)	[0.05, 0.25]	float
$\lambda_{\max}$ (spectral radius)	[0.70, 0.98]	float
α (leakage)	[0.20, 1.00]	float
T_{wash}	[10, 80]	int
r_{target}	[0.05, 0.25]	float
γ (trace decay)	[0.70, 0.98]	float
η (WTA LR)	[0.01, 0.15]	float
n_{syn}	[0, 5]	int
n_{pre}	[0, 3]	int
n_{epochs}	[1, 5]	int
β (self-bias)	[0.00, 0.25]	float
θ_{comp}	[0.15, 0.60]	float

Parameter interpretation. n_{res} controls representational capacity; $\lambda_{\max}$ and α set the reservoir memory timescale; r_{target} governs spike sparsity; γ accumulates STDP temporal credit; and θ_{comp} is

the primary precision-recall trade-off knob. The selected config (α=0.82, n_{syn}=5, θ_{comp}=0.383) balances fast-dynamics datasets (Tigramite) with slow-dynamics datasets (CausalRivers): fast leakage allows quick response while generous synthetic warm-up stabilises the self-gate.

Linear-regime limitation. X-SPK's performance on linear-VAR data (Tigramite, TimeGraph A1/B2) has two distinct causes, and they do not affect all linear-regime datasets equally. On the sparser linear benchmarks — Tigramite Benchmark (D=5, density 25%) and TimeGraph A1 (D=6, density 13%) — the same threshold-density mismatch diagnosed in Section A.8 recurs even under `LINEAR_CONFIG`'s stricter θ_{comp}=0.553: X-SPK still saturates toward the complete graph, SHD is worse than the null baseline, and every predicted graph is cyclic. A density-calibrated threshold on the identical ranking closes this gap (Table 7): SHD beats the null baseline by 60–87% and every predicted graph is acyclic. On these two datasets, the limitation is calibration, not architecture. On the denser, larger random linear graphs (Tigramite Synthetic, D=5 and D=10, density ≈ 35–38%), density calibration only partially helps and acyclicity is not restored, because the ranking itself is weaker and noisier there (AUROC as low as 0.51 on individual seeds, versus 0.71–0.98 on the sparser datasets above). Here the architectural account from Section 3 still applies — linear-Gaussian dynamics leave no concentrated geometric footprint for the WTA gate to exploit — while dedicated methods such as PCMCI+ and VARLiNGAM pursue sparse maximum-likelihood solutions over continuous covariance structure, directly aligned with linear-Gaussian generative assumptions, and hold a structural advantage here regardless of X-SPK's threshold. In practice, purely linear-Gaussian dynamics represent an academic limiting case: real-world causal systems characteristically exhibit threshold effects, multiplicative noise, saturation, or environmental non-stationarity that generate the nonlinear geometry X-SPK is designed to exploit. AUROC = 0.769 on TimeGraph A1 (and 0.91–0.98 on Tigramite Benchmark) confirms that relative edge rankings survive on the sparser linear-regime datasets; `LINEAR_CONFIG`'s stricter θ_{comp} sharpens precision relative to `UNIFIED_CONFIG` but, like the global threshold, is still one fixed value applied across datasets of different density, so it does not by itself resolve the mismatch (Table 7).

Table 7: Same density-calibration test as Table 8, applied to the four linear-regime datasets under `LINEAR_CONFIG` (3 seeds/row for Tigramite Benchmark/A1; 3 independent random-graph draws/row for Tigramite Synthetic). AUROC range is across seeds/draws within each dataset. "DAG rate" is the fraction of runs producing a fully acyclic graph, global → density-oracle.

Dataset	Density	SHD (global)	SHD (oracle)	SHD (12% prior)	SHD (Null)
Tigramite Benchmark (D=5)	25%	15.0	2.0	2.0	10.0
TimeGraph A1 (D=6)	13%	26.0	6.0	6.0	15.0
Tigramite Synthetic (D=5)	35%*	13.0	10.0	8.0	10.0
Tigramite Synthetic (D=10)	38%*	55.0	39.3	35.3	45.0

*Mean over 3 independent random-graph draws; density varies per draw.

DAG rate (global → density-oracle) and AUROC range per seed/draw: Tigramite Benchmark 0% → 100% (AUROC 0.98, fixed graph); TimeGraph A1 0% → 100% (AUROC 0.63–0.77); Tigramite Synthetic D=5 0% → 33% (AUROC 0.51–0.87); Tigramite Synthetic D=10 0% → 0% (AUROC 0.58–0.70) — density calibration restores acyclicity exactly where it also fixes SHD, and fails to on exactly the two datasets whose ranking is already weak.

SHD sensitivity to threshold calibration. Structural Hamming Distance counts the minimum number of edge insertions, deletions, and reversals needed to convert the predicted graph into the ground truth (Peters et al., 2017)—it therefore penalises each spurious edge equally and is highly sensitive to the chosen decision threshold. X-SPK's globally-tuned θ_{comp}=0.383 is optimised on a held-out mix of nonlinear datasets with moderate edge density ($|E|/\binom{D}{2} \approx 0.4$). On sparse graphs this threshold over-fires: competition ratios for nearly all $D(D-1)$ directed pairs exceed it, so X-SPK predicts a near-complete graph and accumulates a large false-positive count. AUROC is unaffected because it measures only the ranking of true edges above spurious ones, independently of where the threshold falls; the same model can maintain high AUROC while producing poor SHD whenever its operating threshold is mis-matched to the true graph density.

Concrete examples from our benchmarks illustrate the gap. On TimeGraph B1 (D=6, $|E|$=3, density 10%), X-SPK predicts approximately 27 directed edges, yielding SHD ≈ 27, whereas Rhino—which pursues a sparse structural-equation solution—achieves SHD = 3.0 by predicting a near-empty

graph with few false positives. Notably, a trivial *predict-nothing* baseline (empty graph) would also achieve SHD $= 3$ on this topology, because with only three true edges the only error is three missed insertions; it would outperform X-SPK on SHD while having AUROC $= 0.5$ and F1 $= 0$. On CausalRivers D=10, X-SPK SHD $= 36.9$ versus Rhino SHD $= 10.8$, reflecting the same density mismatch. A natural mitigation is per-dataset threshold calibration: setting θ_{comp} to the $(1-\hat{p})$-quantile of the competition-ratio distribution, where $\hat{p}$ is the estimated or prior-known edge density, would match prediction density to the ground truth and convert X-SPK's ranking advantage into competitive SHD without altering the core spike-timing mechanism.

Testing this mitigation directly. We reran X-SPK (identical `UNIFIED_CONFIG`, reservoirs, and STDP dynamics; only the post-hoc binarization changes) on the eight main-text TimeGraph configurations, replacing the fixed threshold with a density-quantile threshold: keep the top $\hat{p} \cdot D(D-1)$ competition ratios, using either the true edge density (oracle) or a fixed, dataset-agnostic 12% prior (non-oracle, no ground truth used). Table 8 reports means over 3 seeds. Because both thresholds act on the identical ranking, AUROC is unchanged to four decimal places in every configuration, confirming the ranking itself is not the problem. Under the global threshold, X-SPK predicts the complete graph ($D(D-1)$ of $D(D-1)$ possible directed edges) in all 24 runs, giving SHD roughly double the null baseline; density calibration — including the non-oracle 12% prior — cuts SHD by 60–85% and beats the null baseline in every configuration.

Table 8: Fixed global threshold vs. density-quantile threshold, same competition-ratio ranking (3 seeds/row; AUROC identical between columns to 4 decimals, omitted for space). "Null" is the reference SHD $= \binom{D}{2}$ used throughout Table 1.

Dataset	SHD (global θ_{comp})	SHD (density, oracle)	SHD (density, 12% prior)	SHD (Null)
TG B1 (D=6)	27.0	4.7	5.7	15.0
TG B2 (D=6)	27.0	6.0	6.3	15.0
TG B2 (D=8)	52.0	7.3	10.3	28.0
TG B2C (D=6)	27.0	6.0	6.3	15.0
TG C1 (D=6)	26.0	6.0	6.0	15.0
TG C1 (D=8)	51.0	8.0	10.0	28.0
TG C2 (D=6)	27.0	4.0	5.0	15.0
TG C2 (D=8)	52.0	6.7	7.0	28.0

The same diagnostic settles the acyclicity question raised in Section 3: under the global threshold, every run (24/24) yields a graph containing the maximum possible number of 2-cycles ($\binom{D}{2}$ bidirectional pairs) — not acyclic. Under either density-calibrated threshold, every run (24/24) yields a fully acyclic graph. Acyclicity in practice is therefore real, but conditional on the same density calibration that fixes SHD, not a property of the global-threshold configuration used for this paper's headline numbers.

Is 24/24 acyclic just sparsity? A D-node graph with only $k \ll D(D-1)$ retained edges is acyclic with non-trivial probability by chance alone, regardless of which edges are chosen — the sharper version of the tautology concern a skeptical reading of Corollary 3.5 invites. We test this directly: for each of the eight configurations and both thresholds, we compute the fraction of uniformly random k-edge directed subgraphs (same D, same k; 200,000 Monte Carlo draws) that are acyclic. Under the oracle-density threshold, per-configuration null rates are already high (76–89%, since k is small relative to $D(D-1)$), giving $\mathbb{P}(\text{all } 24/24 \text{ acyclic by chance}) \approx 2.2\%$ — distinguishable from pure sparsity at conventional significance, but not decisively. Under the non-oracle, dataset-agnostic 12%-prior threshold — the practically relevant case, since it uses no ground-truth density information — null rates are lower (54–76%) and $\mathbb{P}(\text{all } 24/24 \text{ by chance}) \approx 6 \times 10^{-5}$: here the acyclicity finding is not explained by sparsity alone.

A.9 Component Ablation

To identify which architectural components drive X-SPK's performance, we evaluate six configurations on four datasets (Mackey–Glass: Collider+Ghost, Dense Hub, Gated Nonlinear; CausalRivers D=5), each with 5 seeds and n_{ens}=3. Each condition disables exactly one component of the full model:

- **full**: complete X-SPK (UNIFIED_CONFIG)
- **no-two-pass**: `two_pass=False`, removes the residual Pass-2 refinement stage
- **no-trace**: γ=0, removes STDP temporal credit (WTA sees only instantaneous spikes)
- **tiny-res**: n_{res}=1, removes representational capacity of reservoir
- **no-self-bias**: β=0, self-gate competes without the β bonus
- **no-pretrain**: n_{syn}=0, removes the synthetic causal warm-start (n_{pre} is already 0 in `UNIFIED_CONFIG`, so real-data self-gate pretraining is off in **full** too — this ablation isolates n_{syn} alone, not two mechanisms)

STDP trace is the single critical component. The **no-trace** condition (γ=0) collapses to F1= 0.000 and AUROC= 0.500 on every dataset: without temporal credit accumulation the WTA gate has no directional signal and reduces to random prediction. All other ablations preserve F1 (binary threshold is robust), so AUROC is the discriminative metric.

Reservoir capacity matters for chaotic dynamics. Reducing to n_{res}=1 (**tiny-res**) drops mean AUROC from 0.772 to 0.675, driven entirely by the Mackey–Glass datasets (AUROC 0.525–0.616 vs. 0.706–0.775 for full). On CausalRivers D=5 (simpler linear dynamics) tiny-res AUROC is 0.914, actually exceeding full (0.871), suggesting the 72-unit reservoir adds noise on low-complexity graphs.

Synthetic warm-start helps most on complex topologies. The **no-pretrain** condition loses −0.116 AUROC on the gated nonlinear topology (0.590 vs. 0.706) while remaining near-full elsewhere. Since real-data self-gate pretraining is already off in **full** (n_{pre}=0), this drop is attributable to the synthetic-causal warm-start alone, cleanly, not to two conflated mechanisms: it accelerates separation between self-predictability and cross-channel evidence where the true autocorrelation structure is hardest to distinguish.

Two-pass refinement and self-bias are secondary. Removing Pass-2 (**no-two-pass**) slightly reduces mean AUROC (−0.008), with the largest drop on the gated topology (−0.057) but a marginal gain on collider+ghost (+0.047). Removing the self-gate bonus (**no-self-bias**) is effectively neutral (+0.002 mean AUROC), confirming that the two-pass architecture provides the same discrimination as an explicit scalar bias.

Table 9: Ablation results (mean across 5 seeds): effect of removing each X-SPK component. The "full" row is UNIFIED_CONFIG. Bold: best per column (excluding no-trace). n_{ens}=3, $\tau_{\max}$=2.

	MG: Collider+Ghost		MG: Dense Hub		MG: Gated Nonlin.		CausalRivers (D=5)		Mean	
Condition	F1	AUROC	F1	AUROC	F1	AUROC	F1	AUROC	F1	AUROC
full	0.621	0.735	**0.667**	0.775	**0.571**	**0.706**	**0.571**	0.871	0.608	0.772
no-two-pass	0.621	**0.782**	**0.667**	0.764	**0.571**	0.649	**0.571**	0.862	0.608	0.764
no-trace	0.000	0.500	0.000	0.500	0.000	0.500	0.000	0.500	0.000	0.500
tiny-res	0.621	0.525	**0.667**	0.644	**0.571**	0.616	**0.571**	**0.914**	0.608	0.675
no-self-bias	0.621	0.729	**0.667**	**0.785**	**0.571**	0.682	**0.571**	0.900	0.608	**0.774**
no-pretrain	**0.625**	0.690	**0.667**	0.713	0.563	0.590	**0.571**	0.893	**0.607**	0.722

A.10 Is WTA-STDP Necessary, or Does Reservoir Preprocessing Alone Explain the Gains?

To test this directly, we replace trace-decay and winner-take-all credit assignment with the paper's own theoretical target quantity computed directly from the same event stream: the conditioned occupancy rate $r_{i\to j}$ (Equation 10), estimated empirically over a fixed 5-step window, with no competition, no learning rate, no self-bias, no pretraining, and no hyperparameter search. Everything upstream — reservoir construction, washout, per-channel quantile eventization, seeds — is identical to X-SPK's own pipeline, so any difference isolates the contribution of the trace+WTA mechanism itself.

Table 10 reports AUROC and density-calibrated SHD/acyclicity for both, on the same 8 TimeGraph configurations and seeds as Table 8. X-SPK's ranking (AUROC) beats this untrained control on 7 of 8 datasets, by up to 0.19 (TG C2, D=6) — a real, fairly consistent advantage for the trace+WTA mechanism over a naive lead-lag statistic on identical features. Once both methods are given the same oracle-density threshold, however, the discretized graph-recovery advantage (SHD) is far weaker: X-SPK wins on only 3 of 8 datasets, ties on 2, and loses narrowly on 3; both methods reach 100% acyclicity on 5 of 8 datasets under this calibration.

We read this as: WTA-STDP measurably improves the continuous ranking over a naive precedence statistic on identical features, consistent with this paper's central ranking claims, but its advantage for thresholded recovery specifically is smaller and less consistent than the ranking gap suggests once both are calibrated to the same density. The control used no hyperparameter search at all, unlike `UNIFIED_CONFIG`'s 150-trial sweep — so its near-parity in SHD is, if anything, the stronger signal: an untuned baseline should not come this close if trace+WTA were the primary driver of graph-recovery quality specifically.

Table 10: X-SPK vs. a direct lead-lag control on identical reservoir features (mean over 3 seeds; same TimeGraph configurations as Table 8). "Control" replaces trace+WTA with the raw conditioned occupancy rate $r_{i\to j}$ over a 5-step window; no tuning, no competition, no self-bias. SHD/DAG columns use oracle-density thresholding for both.

Dataset	AUROC (X-SPK)	AUROC (control)	SHD (X-SPK)	SHD (control)	DAG (X-SPK)	DAG (control)
TG B1 (D=6)	**0.613**	0.552	**4.7**	6.0	**100%**	67%
TG B2 (D=6)	**0.653**	0.598	6.0	**5.3**	**100%**	67%
TG B2 (D=8)	**0.619**	0.549	**7.3**	8.0	100%	100%
TG B2C (D=6)	**0.684**	0.667	6.0	**5.3**	100%	100%
TG C1 (D=6)	0.648	**0.702**	6.0	**5.0**	100%	100%
TG C1 (D=8)	**0.625**	0.624	**8.0**	9.3	100%	100%
TG C2 (D=6)	**0.828**	0.643	**4.0**	4.0	100%	100%
TG C2 (D=8)	**0.782**	0.627	6.7	6.7	**100%**	67%

A.11 Extended CausalRivers Seeds and TimeGraph D-Series (MCAR)

Extended CausalRivers seeds. The main-body CausalRivers results (Table 1) use 15 seeds at D=5 and 10 seeds at D=10. X-SPK retains the highest AUROC at both sizes (0.769 at D=5; 0.728 at D=10) and the highest AUPRC at both sizes. Variance across seeds is substantial, reflecting natural variability in which CausalRivers subgraph is sampled; the qualitative ranking (X-SPK highest AUROC/AUPRC) is stable across it.

TimeGraph D-series: bidirectional feedback and MCAR imputation. The D-series scenarios augment the existing A1/B2/C2 framework with two simultaneous challenges absent from other TimeGraph scenarios: (i) *bidirectional lagged feedback* — at D=6 the causal graph contains a true lag-1 cycle ($X_4(t-1) \to X_5$ and $X_5(t-1) \to X_4$, creating a genuine feedback oscillation), and (ii) *MCAR imputation* — data was generated with Missing Completely At Random dropout then forward-filled before method evaluation. D1 uses linear equations (like A1), D2 nonlinear polynomial (like B2), and D3 nonlinear with trend and seasonality (like C2 / D3-equivalent). Each condition is evaluated on a single n=1,000-sample window; results are therefore exploratory with no variance estimate. D2 and D3 simulation files were not present in the benchmark distribution and are therefore omitted.

Table 11 reports D1 at D=4 and D=6 on the imputed-complete data and, for D=6, also on the MCAR30 forward-filled missing variant. VARLiNGAM fails at D=4 (near-singular covariance on the small imputed sample) and on the D=6 complete variant, but succeeds on MCAR30 missing.

Key findings. *(1) Bidirectional feedback and linear dynamics together suppress X-SPK.* At D=4, X-SPK AUROC (0.429) falls below random (0.500), confirming that timing-based direction inference is unreliable when a true feedback oscillation is present in a linear system. At D=6, X-SPK AUROC (0.742) recovers and matches Rhino, suggesting the richer reservoir geometry at higher D provides partial compensation. Notably, X-SPK achieves the highest AUPRC at D=6 (0.513 vs PCMCI+ 0.507), indicating it ranks the true positives relatively well in the aggregate even when threshold-based F1 is weak.

(2) PCMCI+ and DYNOTEARS are most sensitive to MCAR imputation. Under 30% MCAR with forward-fill at D=6, PCMCI+ AUROC drops from 0.848 to 0.711 (-0.137) and DYNOTEARS from 0.813 to 0.648 (-0.165). These methods rely on linear Gaussian conditional independence assumptions; forward-fill introduces artificial autocorrelation that inflates partial correlations. X-SPK degrades more moderately ($0.742 \to 0.656$, -0.086), while Rhino *improves* ($0.742 \to 0.828$, $+0.086$), suggesting sparse-regression objectives are less sensitive to smoothing artifacts introduced by imputation.

Table 11: TimeGraph D1 (linear + bidirectional feedback, MCAR imputed). Single-seed results (n=1,000); no variance reported. †VARLiNGAM failed (near-singular covariance). Bold: best per block (excluding null). Null: AUROC= 0.5, AUPRC= $|E|/D(D-1)$, SHD= $\binom{D}{2}$.

		D1 D=6 — complete (imputed)				*D1 D=6 — MCAR30 missing (ffill)*			
	Method	F1↑	AUROC↑	AUPRC↑	SHD↓	F1↑	AUROC↑	AUPRC↑	SHD↓
D1 D=6 (|E|=4)	PCMCI+	**0.750**	**0.848**	0.507	**2**	**0.750**	0.711	0.209	**2**
	DYNOTEARS	0.545	0.813	0.386	5	0.545	0.648	0.174	5
	Rhino	0.600	0.742	0.268	4	**0.750**	**0.828**	0.312	**2**
	VARLiNGAM	†—	—	—	—	0.667	0.656	0.174	3
	Null	0.000	0.500	0.133	15	0.000	0.500	0.133	15
	X-SPK	0.235	0.742	**0.513**	26	0.235	0.656	**0.388**	26

B PROOFS

Proof roadmap. The proof chain is direct: occupancy asymmetry (Lemma 3.1, main text) → conditional trace dominance over every *external* competitor (Lemma B.3) → monotonicity of the resulting competition (Lemma B.6), applied to the game restricted to externals once self's own, independent draw is conditioned out → asymmetric expected weight accumulation (Theorem 3.3), which combines this residual win-share asymmetry with Assumption B.1's comparison of how self-predictable each endpoint is. A1 and A2 are structural. Assumption B.4 is the only premise needed to turn a trace gap in expectation into a win-probability gap against any number of competitors at once; Assumption B.1 is the only further premise, and it is a comparison between two autocorrelation-driven quantities rather than a requirement that the causal signal itself outweigh the target's autocorrelation.

B.1 ASSUMPTIONS USED IN PROOFS

A1 (stationarity and mixing). The post-washout spike process $\{\delta_i(t)\}_{t\geq 0}$ is strictly stationary and strongly mixing for every channel i. Under A1, time averages and conditional expectations are well-defined and stationary.

A2 (equal spike rates). Quantile thresholding (Equation equation 2) enforces $\mathbb{P}(\delta_i(t)=1) = r_{\text{target}}$ for every channel i. Consequently, for all i, j, T:

$$\mathbb{E}[N_i(T)] = \mathbb{E}[N_j(T)] =: N(T). \tag{17}$$

A2 is a direct consequence of the architectural choice in Equation equation 2.

Assumption B.1 (Comparable Self-Predictability)**.** Let $q_k := \mathbb{P}(\text{winner} = k \mid \delta_k(t)=1)$ denote channel k's own *self-win rate* — how often its own trace wins the WTA competition at its own firings. The target is not more self-predictable than the cause:

$$q_j \;\leq\; q_i. \tag{18}$$

This compares two autocorrelation-driven quantities against each other — j's and i's own self-predictability — rather than requiring the causal signal to outweigh the target's autocorrelation outright. It is substantially weaker than requiring i's decay-weighted elevation alone to exceed j's full self-trace, and is symmetric in spirit: it fails only when the target is a systematically stronger self-predictor than the source (e.g. a physically inertial target driven by a comparatively fast or noisy source), not whenever autocorrelation is merely present.

Intuition. Self winning a given draw costs the true cause nothing beyond that instant. A genuine external driver's own activity is not synchronized to the target's decay clock, so it arrives asynchronously relative to self's own rhythm and still gets its turn whenever self's trace happens to be momentarily weak. Self's routine participation is then not noise to be tolerated but a genuine signal-to-noise improvement: it is precisely what stops a merely-correlated, non-causal channel from accumulating credit on the strength of the target's own predictable activity, since only a channel that beats the target's *own* recent history — not just an arbitrary rival — is ever credited. This is an informal picture, not a further formal claim: the assumption above only needs the target to not out-predict itself by more than the source does, not that self be beaten outright.

Remark B.2 (Where This Comes From). Self's own reservoir weights and exogenous noise are drawn independently of every other channel's, by the same architectural fact Remark B.5 already uses for non-causal competitors (i.i.d. weight draws, independent per-channel noise) — there is nothing structurally special about channel j competing for its own credit versus any other channel doing so. Unconditionally, then, self-win rate is governed by the same generic architecture as every other channel's, and $q_i \approx q_j$ absent a genuine difference in the underlying data's autocorrelation structure — the same kind of approximation already accepted for Assumption B.4 (architecture symmetry gets us most of the way; the residual gap is a property of the data, not the mechanism). Measuring q_i, q_j directly on real spike trains (`analyze_conditional_win_rate.py`) supports this: $q_j \leq q_i$ holds on the majority of tested edges, and where it fails the gap $q_j - q_i$ is small and centered near zero rather than large and systematic.

B.2 Full Proofs

Lemma B.3 (Conditional Trace Dominance)**.** *Let $\tilde{\rho}_{j,k}(t) := \rho_k(t^-) + \beta\,\mathbf{1}[k=j]$ denote the biased pre-spike trace of channel k at a spike time of target j, where $\rho_k(t^-) = \sum_{\Delta'\geq 1} \gamma^{\Delta'} \delta_k(t-\Delta')$ (Equations equation 3–equation 8). If i causes j (Lemma 3.1), then $\mathbb{E}[\tilde{\rho}_{j,i}(t) \mid \delta_j(t){=}1] > \mathbb{E}[\tilde{\rho}_{j,k}(t) \mid \delta_j(t){=}1]$ for every genuinely non-causal external competitor $k \neq i, j$. (Self, $k = j$, still competes throughout — it is simply not one of the external rivals compared here; Theorem 3.3's proof shows how self's own participation is handled.)*

Proof. Taking expectations of the trace definition and applying Bayes' rule under the equal-rate property A2, for $k \neq j$,

$$\mathbb{E}[\tilde{\rho}_{j,k}(t) \mid \delta_j(t){=}1] = \sum_{\Delta'\geq 1} \gamma^{\Delta'}\, \mathbb{P}(\delta_j(t){=}1 \mid \delta_k(t-\Delta'){=}1). \tag{19}$$

For $k = i$, the signal-noise decomposition (Equation equation 11) elevates each per-lag term above baseline within f's effective memory horizon (assuming, as throughout, that this elevation holds lag-by-lag rather than only on aggregate — true whenever f's influence decays smoothly, without cancelling across lags). A genuinely non-causal channel $k \neq i, j$ has no channel into j at all, so every per-lag term sits at baseline. Hence i's decay-weighted sum strictly exceeds every such rival's. □

Why the next step must be true, before the algebra. Every competing trace at a firing event is built the same way — a decaying sum of that channel's own recent spikes — and channels differ only in how likely those spikes were, not in the *shape* of the randomness around that likelihood. So relabeling which channel is which cannot change the odds of any outcome; only a genuine gap in average level can tip a competition. That is already enough to show the channel with the largest average wins more than its fair share, and every channel without such an advantage wins less. The lemma below makes this precise for any number of correlated competing traces at once.

Assumption B.4 (Exchangeable Fluctuations)**.** Write each competitor's biased trace at a conditioning event as $\tilde{\rho}_k = \mu_k + \xi_k$, its mean plus a residual fluctuation. The fluctuation vector $(\xi_1, \ldots, \xi_D)$ is exchangeable: its joint law is unchanged by permuting channel labels. This holds whenever every channel's trace is built from the same decay recipe (Equations equation 3–equation 8) and channels differ only in their conditional spike probabilities — exactly the case here, since μ_k already absorbs every channel-specific difference.

Remark B.5 (Where This Comes From). Not fully derivable from first principles, but three structural facts bound how far from true Assumption B.4 can be.

Unconditionally, exchangeability is automatic, no data assumption needed. Every channel runs the identical pipeline: reservoir weights drawn i.i.d. across channels, identical calibration to r_{target} (A2), and identical per-channel z-scoring to mean 0/variance 1 before any reservoir processing (`standardize()`, the first line of `run_xspk`, applied uniformly in every experiment here). Before conditioning on any target's firing, the D channels' fluctuations are exchangeable by construction.

Non-causal competitors are independent, not merely exchangeable. Weights are drawn independently across channels, and exogenous noise terms are independent across nodes under the no-confounding setup already assumed elsewhere (Equation equation 11). Two channels with no causal link to

the target are therefore functions of disjoint independent randomness — fully independent, strictly stronger than exchangeable. That is all Corollary B.7's part (ii) ever needs.

The one real approximation is second-order. Only the true cause i's history is genuinely correlated with the target's firing, so only i's conditioning is not a pure relabeling symmetry. Under a Gaussian idealization, conditioning on a variable correlated at level ρ shrinks its variance to $\sigma^2(1-\rho^2)$; since z-scoring already equalizes every channel's unconditional σ^2, this ρ-dependent shrinkage is the only asymmetry left — and it is $O(\rho^2)$, versus the $O(\rho)$ mean elevation Theorem 3.3 actually needs. In the weak-signal regime this architecture targets, the first-order effect being proven dominates the second-order approximation being made.

Lemma B.6 (Monotonicity of Competition)**.** *Under Assumption B.4, if $\mu_a \geq \mu_b$ then $\mathbb{P}(a \text{ wins}) \geq \mathbb{P}(b \text{ wins})$, strictly whenever $\mu_a > \mu_b$ and ties occur with probability zero.*

Proof. Fix a, b and let W be the best score among every *other* competitor — untouched by relabeling a and b. Exchangeability of (ξ_a, ξ_b) lets us swap them without changing any probability, so

$$\mathbb{P}(b \text{ wins}) = \mathbb{P}\big(\underbrace{\xi_a - \xi_b > \mu_a - \mu_b}_{\text{beats } a},\ \underbrace{\mu_b + \xi_a > W}_{\text{beats the rest}}\big), \tag{20}$$

while $\mathbb{P}(a \text{ wins})$ is the identical expression with both thresholds relaxed: $\mu_a - \mu_b$ replaced by its negative, and μ_b replaced by $\mu_a \geq \mu_b$. Relaxing a threshold can only enlarge an event, so $\{b \text{ wins}\} \subseteq \{a \text{ wins}\}$ pointwise, giving the claim; the inclusion is strict with positive probability whenever $\mu_a > \mu_b$ and the fluctuation distribution has no atom exactly cancelling the gap. □

Corollary B.7. *If one competitor's mean strictly exceeds every other's, it wins with probability $> 1/D$; any competitor at or below the common baseline wins with probability $\leq 1/D$.*

Proof. Apply Lemma B.6 pairwise against every competitor and sum the D win probabilities to 1. □

Full proof of Theorem 3.3. Fix an ordered pair (i, j) with i causing j. Let $N_j(T)$ and $N_i(T)$ denote spike counts up to time T. Define winner indicators

$$I_{ij}^{(n)} := \mathbf{1}\{i \text{ wins WTA at } j\text{'s } n\text{-th spike}\}, \quad I_{ji}^{(m)} := \mathbf{1}\{j \text{ wins WTA at } i\text{'s } m\text{-th spike}\}. \tag{21}$$

From Equation equation 6,

$$W_{ij}(T) - W_{ij}(0) = \frac{\eta}{\sqrt{D}} \sum_{n=1}^{N_j(T)} I_{ij}^{(n)}, \qquad W_{ji}(T) - W_{ji}(0) = \frac{\eta}{\sqrt{D}} \sum_{m=1}^{N_i(T)} I_{ji}^{(m)}. \tag{22}$$

Under A1, the marks $I_{ij}^{(n)}$ attach to j's stationary spike epochs, so linearity of expectation over them follows from Campbell's theorem for stationary marked point processes — not Wald's identity, which needs i.i.d. summands or a stopping time, neither of which applies to a temporally dependent point process observed to a fixed horizon T:

$$\mathbb{E}[\Delta W_{ij}(T)] = \frac{\eta}{\sqrt{D}}\, \mathbb{E}[N_j(T)]\, p_{ij}, \qquad \mathbb{E}[\Delta W_{ji}(T)] = \frac{\eta}{\sqrt{D}}\, \mathbb{E}[N_i(T)]\, p_{ji}, \tag{23}$$

where $p_{ij} := \mathbb{P}(i \text{ wins} \mid j \text{ fires})$ and $p_{ji} := \mathbb{P}(j \text{ wins} \mid i \text{ fires})$. By A2, $\mathbb{E}[N_i(T)] = \mathbb{E}[N_j(T)] = N(T)$, so the comparison reduces to p_{ij} versus p_{ji}.

Conditioning on self. A self-win credits no external channel, so $p_{ij} = (1 - q_j)\, \pi_{ij}$ and $p_{ji} = (1 - q_i)\, \pi_{ji}$, where q_k is channel k's self-win rate (Assumption B.1) and $\pi_{ij} := \mathbb{P}(i \text{ wins} \mid j \text{ fires},\ j\text{'s self-slot does not win})$ is i's share of j's *residual* firings once self is excluded (not to be confused with the trace ρ; π denotes a win probability throughout this proof). This is an exact identity, not an approximation: winning implies self did not win.

The residual share is asymmetric. Conditional on self not winning, the overall winner is exactly the $\arg\max$ among the $D-1$ external competitors (removing a non-maximal candidate cannot change the $\arg\max$ of the rest). Self's own draw is independent of every other channel's (Remark B.2), so conditioning on it does not privilege any particular external over another. Lemma B.3 makes i's mean

trace the strict maximum among those $D-1$ externals at j's firings, so by Lemma B.6 applied to this $(D-1)$-competitor game, $\pi_{ij} > 1/(D-1)$. Symmetrically, j carries no elevation at i's firings — the causal edge runs $i \to j$, not $j \to i$ — so $\pi_{ji} \leq 1/(D-1)$. Hence $\pi_{ij} > \pi_{ji}$, using nothing beyond what Lemma B.3 already establishes.

Combining. Assumption B.1 gives $q_j \leq q_i$, i.e. $1 - q_j \geq 1 - q_i > 0$, so

$$p_{ij} = (1 - q_j)\pi_{ij} \;\geq\; (1 - q_i)\pi_{ij} \;>\; (1 - q_i)\pi_{ji} = p_{ji}. \tag{24}$$

Writing $p_{ij} - p_{ji} =: \delta_{ij} > 0$ and applying equation 23,

$$\mathbb{E}[\Delta W_{ij}(T)] - \mathbb{E}[\Delta W_{ji}(T)] = \frac{\eta}{\sqrt{D}}\, N(T)\,(p_{ij} - p_{ji}) = \frac{\eta}{\sqrt{D}}\, N(T)\,\delta_{ij} > 0. \tag{25}$$

This is the claimed asymmetric WTA accumulation. □

Full proof of Remark 3.4. Assume by contradiction that a directed cycle is stable:

$$i_1 \to i_2 \to \cdots \to i_k \to i_1. \tag{26}$$

Assume the retained orientation satisfies the theorem's ordering criterion: there exists a strict scalar potential Ψ over nodes such that any retained edge $u \to v$ requires $\Psi_u > \Psi_v$. Then stability of each edge $i_m \to i_{m+1}$ implies

$$\Psi_{i_m} > \Psi_{i_{m+1}}, \qquad m = 1, \ldots, k, \tag{27}$$

with $i_{k+1} := i_1$. Chaining gives

$$\Psi_{i_1} > \Psi_{i_2} > \cdots > \Psi_{i_k} > \Psi_{i_1}, \tag{28}$$

an impossibility. Therefore no directed cycle can be stable. □

Full proof of Corollary 3.5 (Conditional Acyclicity). By Theorem 3.3, each true orientation has positive expected update margin over its reverse orientation, and since empirical gate updates are averages of bounded increments, concentration implies the true orientation is retained with probability tending to 1 as observation length grows. *If*, in addition, the retained edges' margins admit the strict global ordering hypothesized in Remark 3.4, then by that remark no directed cycle can be stable, so the learned graph is acyclic. We do not derive the ordering hypothesis from Theorem 3.3 alone — it is an additional, unverified structural condition on the joint distribution of all pairwise margins, and Corollary 3.5 is conditional on it holding. □

Remark B.8 (What the Ordering Hypothesis Actually Requires of the Data). This is not a claim we prove, or need to prove, about the competition-ratio score function's own transitivity. Theorem 3.3 only ever compares a genuine edge to its own reverse at its own target — nothing here implies the full pairwise ranking is transitive as a property of the estimator, and indeed it empirically is not (a full-tournament check across all $\binom{D}{2}$ pairs on the density-calibration runs of Appendix A.8 finds cyclic triples in 23 of 24 runs, averaging 15.6% of triples among those). This is expected: mediator attenuation is exactly what Lemma 3.1's decomposition predicts, since a target's onward influence on a further variable depends on its *total* activity, noise included, not the isolated fraction its own cause produced.

What the hypothesis actually requires is a condition on the *data-generating process*, in the spirit of standard causal-sufficiency assumptions. For a true chain $A \to B \to C$: no alternative channel into A from C — confounding, a genuine reverse pathway, or simply C's dynamics being more associated with A's history than the doubly-attenuated true path is — may make the local, target-anchored competition between A and C favor $C \to A$. WTA reasons about one pair and one target at a time, with no notion of paths, so it cannot detect or veto such a channel; a spurious $C \to A$ can close a cycle even when $A \to B$ and $B \to C$ are each individually correct. Acyclicity is conditional on the generative process having no such channel, not on any transitivity property of the learned ranking.

B.3 Asset Licenses and Usage Terms

Table 12 documents external assets used in this work, with citation/source, stated license status at source, and usage/redistribution notes for this submission.

Table 12: Per-asset license and usage-term summary for code/data assets used in this paper.

Asset	Source / Citation	License status at source	Usage and redistribution in this submission
CausalRivers benchmark data	CausalRivers benchmark release (Stein et al., 2025)	Data License Germany – Attribution – Version 2.0 (dl-de/by-2-0)	Used for evaluation only; we do not redistribute raw benchmark files in this manuscript
TimeGraph synthetic benchmarks (A/B/C-series, confounder variants)	TimeGraph generator/release (Ferdous et al., 2025)	Creative Commons Attribution 4.0 International (CC BY 4.0)	Used to generate/evaluate benchmark runs; no third-party raw data redistributed
Tigramite benchmark/synthetic datasets	Tigramite benchmark framework (Runge, 2020; Runge et al., 2019)	GNU General Public License v3.0 (GPLv3)	Used for benchmark generation/evaluation under upstream usage terms
Mackey–Glass and Lorenz datasets	Generated from published dynamical equations (Mackey & Glass, 1977; Lorenz, 1963)	Procedurally generated from equations (no third-party dataset license required)	Generated locally by script; reproducible from equations and seeds
Baseline method implementations: PCMCI+/GPDC (via tigramite) (Runge et al., 2019)	Public code release	GNU General Public License v3.0 (GPLv3)	Used for comparison only; no upstream code is redistributed in our artifact
Baseline method implementation: DYNOTEARS (via causalnex) (Pamfil et al., 2020)	Public code release	Apache License 2.0	Used for comparison only; no upstream code is redistributed in our artifact
Baseline method implementation: VARLiNGAM (via lingam) (Hyvärinen et al., 2010)	Public code release	MIT License	Used for comparison only; no upstream code is redistributed in our artifact
Baseline method implementation: Rhino/RhinoLite (via causica) (Gong et al., 2023)	Public code release	MIT License	Used for comparison only; no upstream code is redistributed in our artifact
NeuralGC implementation	Tank et al. (2021) and accompanying public code	MIT License	Custom re-implementation with group-lasso training; no upstream code redistributed
CCM via pyEDM	Sugihara et al. (2012); pyEDM library	Custom UC San Diego research/education-use license (non-commercial; not OSI-approved BSD)	Used via pyEDM API for non-commercial research comparison only, per license terms