CYCLIC STRUCTURAL HAWKES PROCESSES FOR LEARNING CAUSAL STRUCTURE FROM DISCRETE-TIME EVENT SEQUENCES

Anonymous authors
Paper under double-blind review

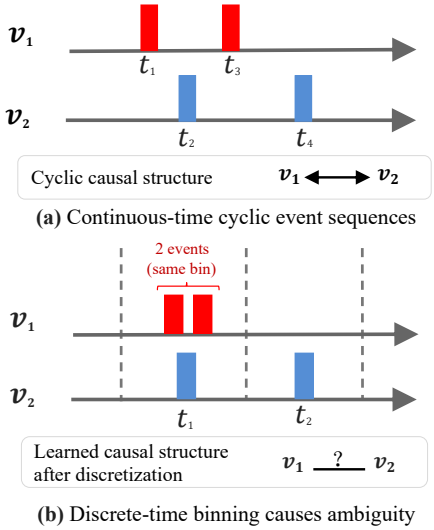
ABSTRACT

Learning causal structure from discrete-time event sequences is particularly challenging because the exact temporal order of events within the same time interval is unobserved. As a result, causal relations that occur within one interval cannot be captured by methods that rely only on strictly lagged dependence. To tackle this issue, existing approaches such as Structural Hawkes Processes (SHPs) explicitly model causal interactions within the same interval, but require the corresponding causal structure to be acyclic to ensure a well-defined recursive generative process. Such an assumption can be restrictive, since events within one interval may repeatedly excite each other and naturally form directed cycles. In this work, we propose Cyclic Structural Hawkes Processes (C-SHPs) for learning cyclic causal structures from discrete-time event sequences. By connecting the observed event counts to an underlying continuous-time Hawkes process through the probability generating functional (PGFL), we develop theoretical tools for characterizing self-loops, adjacency, and edge directions from temporally aggregated observations. We further establish the identifiability of the directed causal structure in the presence of self-loops, reciprocal edges, and arbitrary directed cycles. Based on these results, we develop a practical causal discovery algorithm that operates directly on discrete-time event counts. Experiments on cyclic graphs demonstrate robust performance across varying time resolutions, model parameters, sample sizes, graph dimensions, and graph densities.

1 INTRODUCTION

Learning causal structures among event types from event sequences is an important task in many domains, including social systems, finance, and network operations (Zhou et al., 2013; Bacry et al., 2015; Cai et al., 2022). When event records are sufficiently accurate, event sequences are particularly informative for causal discovery because their temporal ordering helps distinguish a potential cause from its response, given that an event can only excite events that occur later. Multivariate Hawkes processes encode this temporal mechanism directly through their excitation kernels, and the support of these kernels therefore provides a directed causal graph among event types (Hawkes, 1971; Xu et al., 2016). The usefulness of this temporal information, however, heavily depends on recording event times at a sufficiently high resolution.

Such high-resolution recording is often expensive or practically impossible because of limited logging and storage, temporal aggregation, or imperfect synchronization across devices. Consequently, multiple continuous-time events may be recorded in the same observation interval, producing low-resolution event sequences in which their within-interval order is unavailable. Figure 1 illustrates the resulting ambiguity for a reciprocal causal relation. In Figure 1(a), events of types v_1 and v_2 occur in an alternating order. The sequence therefore contains temporal evidence for both directions of the reciprocal relation $v_1 \leftrightarrow v_2$. After low-resolution recording, as shown in Figure 1(b), several of these events are placed in the same interval. Some order across intervals may remain, but the within-interval ordering is missing, creating *instantaneous effect* issues. Recovering these effects requires using dependence within an observation interval, rather than relying only on temporal order between intervals. Previous methods such as Structural Hawkes Processes (SHPs) do so by allowing the count of one event type to depend on counts of other types in the same interval (Qiao et al., 2023). SHP specifies these same-interval relations through conditional distributions and combines them recursively into a joint model for the current count vector.



(a) Continuous-time cyclic event sequences

(b) Discrete-time binning causes ambiguity

Figure 1: Toy example of the continuous-time and discrete-time event sequences under the cyclic causal relationship.

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.

However, for this construction to work, the instantaneous parents must be generated before the child, which in turn requires a directed acyclic graph (DAG) assumption for providing exactly such an ordering through its topological order.

The cyclic structure shown in Figure 1, however, violates this requirement, and the same issue arises for longer directed cycles and self-loops. While in continuous time, a cycle among event types is well defined, we therefore retain a cyclic continuous-time Hawkes process as the generative model and investigate the low-resolution effect, aiming to answer whether the cyclic excitation structure can still be identified from the joint distribution of the observed interval counts from the continuous-time Hawkes process.

Under this formulation, the key difficulty is to characterize the count distribution induced by low-resolution observation. A given interval count can correspond to many continuous-time event configurations, while the unobserved event times determine how excitation propagates within and across intervals. The latent event times therefore need to be integrated out without losing the dependence of the observed distribution on the underlying excitation structure. We address this through the probability generating functional (PGFL) of the Hawkes process. The PGFL characterizes the continuous-time point-process distribution, and with test functions that are constant on each observation interval, it reduces to the PGFL of the observed counts. This provides a direct link from the cyclic continuous-time process to the low-resolution count distribution.

Based on this connection, we derive an exact finite-horizon representation of the count PGFL for cyclic Hawkes processes. For shared-decay exponential kernels, the representation further reduces to a finite-dimensional ODE. To retain excitation accumulated before a later observation window, we also derive an affine conditional-window log-PGFL that separates the empty-history contribution from the entrance-state contribution. Its derivatives reveal local signals of self-loops, adjacency, and edge direction. We then establish support identifiability from the complete marginal pgfl in the presence of self-loops and general directed cycles, developing Cyclic Structural Hawkes Processes (C-SHPs) for causal discovery from low-resolution event sequences.

Our contributions are summarized as follows.

1. We derive a PGFL-based representation that connects the underlying continuous-time Hawkes process to the distribution of discrete-time event counts. For shared-decay exponential kernels, this representation admits an exact finite-dimensional ODE formulation, providing a tractable basis for analyzing cyclic causal structures under temporal aggregation.
2. We characterize how self-loops, adjacency, and edge directions are reflected in the distribution of discrete-time counts using PGFL, and prove that the directed causal structure is identifiable even in the presence of self-loops, reciprocal edges, and arbitrary directed cycles.
3. We develop C-SHP, a causal discovery procedure that combines complementary structural information from discrete-time graph fitting, PGFL-based scores, and marginal factorial cumulants, without requiring reconstruction of latent within-interval event times.

2 PRELIMINARIES

We begin with the definition of the Hawkes process. A Hawkes process is a self- and mutually-exciting point process in which past events can increase the occurrence intensity of future events. Its excitation structure therefore provides a natural representation of directed interactions among event types.

Definition 1 (Multivariate Hawkes Process)**.** For types $V = \{1, \dots, d\}$, let $N_v(t)$ be the type-v counting process. The Hawkes intensity is

$$\lambda_v(t \mid \mathcal{H}_t) = \mu_v + \sum_{v' \in V} \int_0^t \phi_{v'v}(t-s)\, \mathrm{d}N_{v'}(s), \tag{1}$$

where $\mathcal{H}_t = \{N_u(s) : 0 \le s < t,\ u \in V\}$ denotes the event history before time t, and $\mu_v > 0$ is the immigration intensity of type v. Throughout, the process starts from empty history at time zero.

In this work, we specialize the excitation kernel to the shared-decay exponential form $\phi_{v'v}(t) = \alpha_{v'v} e^{-\beta t}$ with $t > 0$, where $\alpha_{v'v} \ge 0$ denotes the excitation strength from type v' to type v, and $\beta > 0$ is the shared decay rate. Exponential kernels constitute a standard and widely used specification of Hawkes processes, while also providing a tractable setting for theoretical analysis. In particular, the exponential form admits a finite-dimensional Markov representation of the event history, and its convolution structure is preserved under differentiation. These properties allow the PGFL of the discretized counts to be represented exactly through a finite-dimensional ODE, providing the analytical foundation for the identifiability results developed below. Writing $\boldsymbol{\mu} = (\mu_v)_{v \in V}$ and $A = (\alpha_{uv})_{u,v \in V}$, we define the target causal graph $G = (V, E)$ by $E = \{(u, v) : \alpha_{uv} > 0\}$. This support convention permits self-loops, reciprocal pairs, and longer directed cycles.

Assumption 1 (Causal Sufficiency and Model Specification)**.** The set V contains all event types relevant to the joint excitation mechanism among the observed types and no unobserved event type jointly excites two or more types in V. The counting process on V is correctly specified by the d-variate Hawkes model in Equation (1)

with exponential kernels, such that latent confounding and model misspecification are outside the scope of our identifiability results.

For the parameter $\theta = (\boldsymbol{\mu}, A, \beta)$, structural identifiability asks whether the joint distribution of the observed discrete-time counts uniquely determines the directed support of the excitation matrix A, and hence the causal graph G. Establishing this result requires characterizing how the underlying continuous-time Hawkes process induces a distribution over temporally aggregated event counts after the exact event times within each observation interval are marginalized out. The probability generating functional (PGFL) provides a convenient representation for this purpose.

Definition 2 (PGFL of a Poisson Process)**.** Let $\boldsymbol{N} = (N_v)_{v\in V}$ be independent inhomogeneous Poisson processes with locally integrable intensities $\nu_v : [0,T] \to [0,\infty)$. For any measurable $\boldsymbol{z}(t) = (z_v(t))_{v\in V} \in [0,1]^d$, their finite-horizon PGFL is

$$G_T[\boldsymbol{z}] = \mathbb{E}\left[\prod_{v\in V}\prod_{t\in N_v\cap(0,T]} z_v(t)\right] = \exp\left\{\sum_{v\in V}\int_0^T \nu_v(t)\big[z_v(t)-1\big]\,\mathrm{d}t\right\}. \tag{2}$$

The PGFL assigns a factor $z_v(t)$ to each type-v event occurring at time t. This representation can be further applied to construct the PGFL of the Hawkes process using a cluster representation in which immigrant events arrive according to Poisson processes and each event generates offspring through type-specific Poisson processes. When $\boldsymbol{z}(t)$ is chosen to be constant within each observation interval, the PGFL depends on an event sequence only through the corresponding interval counts. It therefore reduces to the multivariate probability generating function of the observed count vector, providing a direct connection between the continuous-time Hawkes process and its discrete-time observations.

3 IDENTIFIABILITY

Structural identifiability asks whether two different directed graph structures can generate the same distribution of the observed data. To study this question for count processes, we represent the distribution through its probability generating functional (PGFL), which uniquely characterizes the distribution of the observed event counts, and provides a convenient characterization of the distribution while retaining the branching structure of the Hawkes process. We first derive the PGFL of the underlying continuous-time process. By restricting the test functions to be constant within each observation window, we then obtain the PGFL of the discretely observed counts, which will form the basis of the identifiability analysis.

3.1 PGFL AND DISCRETE-TIME REPRESENTATION

We begin with the PGFL of the underlying continuous-time Hawkes process. Unlike an acyclic model, the excitation graph may contain directed cycles, so the process cannot in general be constructed recursively according to a topological ordering of the event types. We therefore use the Poisson-cluster representation of the Hawkes process, under which each immigrant event initiates an independent cluster consisting of the immigrant itself and all of its descendants. This representation applies without requiring the excitation graph to be acyclic. For a test function $\boldsymbol{z} = (z_v(t))_{v\in V}$ with $z_v(t) \in [0,1]$, let $F_{v'}(t;\boldsymbol{z})$ denote the finite-horizon PGFL of the cluster rooted at a type-v' event at time t, and let $G_T[\boldsymbol{z}]$ denote the PGFL of the full Hawkes process. The following theorem gives the corresponding cluster and process-level PGFLs.

Theorem 1 (Continuous-Time Hawkes Cluster and Global PGFLs)**.** *For the Hawkes process on* $[0,T]$*, the cluster and global PGFLs satisfy*

$$F_{v'}(t;\boldsymbol{z}) = z_{v'}(t)\exp\left\{\sum_{v\in V}\alpha_{v'v}\int_t^T e^{-\beta(\tau-t)}\big[F_v(\tau;\boldsymbol{z})-1\big]\,\mathrm{d}\tau\right\}, \tag{3}$$

$$G_T[\boldsymbol{z}] = \exp\left\{\sum_{v'\in V}\mu_{v'}\int_0^T \big[F_{v'}(t;\boldsymbol{z})-1\big]\,\mathrm{d}t\right\}. \tag{4}$$

Theorem 1 provides the starting point for the distributional analysis. In (3), the factor $z_{v'}(t)$ accounts for the root event, and the exponential term accounts for all descendant clusters generated from it. Equation (4) then combines the independent clusters initiated by the Poisson immigrant processes.

Under discrete observation, however, the timestamps and the ordering of events within each window are lost, and only the window counts remain. Their joint PGF is obtained by evaluating the continuous-time PGFL at test functions that are constant on each observation interval.

Consider a partition $0 = t_0 < t_1 < \cdots < t_K = T$, with $I_k = (t_{k-1}, t_k]$ and $t_k - t_{k-1} = \Delta$. Let $Y_{v,k} := N_v(I_k)$, suppressing the record index, and set $z_v(t) = z_{v,k}$ for $t \in I_k$. Writing $\boldsymbol{Z} = (z_{v,k})_{v\in V,\,1\le k\le K}$, evaluating the

continuous-time PGFL at this stepwise-constant test function gives the joint PGFL of the discrete-time processes. Specifically, we define $G_K(\boldsymbol{Z})$ and denote its logarithm by $\ell_K(\boldsymbol{Z})$:

$$G_K(\boldsymbol{Z}) := \mathbb{E}\left[\prod_{k=1}^{K}\prod_{v\in V} z_{v,k}^{Y_{v,k}}\right], \qquad \ell_K(\boldsymbol{Z}) := \log G_K(\boldsymbol{Z}). \tag{5}$$

Thus, structural identifiability from discrete observations can be studied through $\ell_K(\boldsymbol{Z})$. To use this representation, however, we still need to express the discrete-count PGFL explicitly in terms of Theorem 1 is defined through convolution integrals over future descendants. For the exponential kernel, these convolutions admit a finite-dimensional state representation, which allows the discrete-count PGFL to be characterized by an ordinary differential equation.

Theorem 2 (PGFL of Discrete Observation Windows)**.** *Under the partition and stepwise-constant PGFL argument defined above, let $F_v(\tau;\boldsymbol{Z})$ denote the cluster PGFL of Theorem 1 evaluated at the argument specified by $\boldsymbol{Z}$. For each interval I_k, set $s = t_k - t \in [0,\Delta]$ and define*

$$U_v^{(k)}(s) = \int_{t_k-s}^{T} e^{-\beta[\tau-(t_k-s)]}\left[F_v(\tau;\boldsymbol{Z})-1\right]\mathrm{d}\tau. \tag{6}$$

This state summarizes the future cluster contribution seen from time $t_k - s$. Substituting the cluster PGFL and differentiating with respect to s gives

$$\begin{aligned}
\dot{U}_{v'}^{(k)}(s) &= -\beta U_{v'}^{(k)}(s) + z_{v',k}\exp\left\{\sum_{v\in V}\alpha_{v'v}U_v^{(k)}(s)\right\} - 1,\\
\dot{J}_{v'}^{(k)}(s) &= z_{v',k}\exp\left\{\sum_{v\in V}\alpha_{v'v}U_v^{(k)}(s)\right\}, \qquad J_{v'}^{(k)}(0) = 0.
\end{aligned} \tag{7}$$

Here $J_{v'}^{(k)}$ accumulates the cluster contribution over I_k. The states propagate backward from T using $U_{v'}^{(K)}(0) = 0$ and $U_{v'}^{(k)}(0) = U_{v'}^{(k+1)}(\Delta)$ for $k < K$. For the empty-history process, the joint log-PGFL is

$$\ell_K(\boldsymbol{Z}) = \sum_{k=1}^{K}\sum_{v\in V}\mu_v\left[J_v^{(k)}(\Delta)-\Delta\right]. \tag{8}$$

A later window inherits excitation from preceding events. Define $X_v(t) := \int_{(0,t)} e^{-\beta(t-\tau)}\,\mathrm{d}N_v(\tau)$. Now fix one interval I_k and introduce the d-dimensional argument $\boldsymbol{z} = (z_v)_{v\in V}$. It is the argument of this individual window's PGFL: in the joint argument $\boldsymbol{Z}$, set $z_{v,k} = z_v$ for each v and set the arguments of all other windows to 1. Solve the ODE afresh for this single window, using z_v in place of $z_{v,k}$ and $U_v^{(k)}(0) = J_v^{(k)}(0) = 0$. In the following two formulas, $U^{(k)}$ and $J^{(k)}$ refer to this single-window solve, not to the kth segment of the joint backward solve above. Conditional on the entrance state, the single-window log-PGFL is

$$\ell_k(\boldsymbol{z}\mid X(t_{k-1})) := \log\mathbb{E}\left[\prod_{v\in V} z_v^{Y_{v,k}}\,\middle|\, X(t_{k-1})\right] = \sum_{v\in V}\mu_v\left[J_v^{(k)}(\Delta)-\Delta\right] + X(t_{k-1})^\top A\boldsymbol{U}^{(k)}(\Delta). \tag{9}$$

Marginalizing over the entrance state gives

$$\ell_k(\boldsymbol{z}) := \log\mathbb{E}\left[\prod_{v\in V} z_v^{Y_{v,k}}\right] = \sum_{v\in V}\mu_v\left[J_v^{(k)}(\Delta)-\Delta\right] + \log\mathbb{E}\left[e^{X(t_{k-1})^\top A\boldsymbol{U}^{(k)}(\Delta)}\right]. \tag{10}$$

For $k = 1$, $X(0) = \mathbf{0}$, so the conditional and marginal formulas coincide. For $\boldsymbol{Z} \in [0,1]^{d\times K}$, the joint ODE solution is bounded and unique.

Theorem 2 connects the continuous-time process to the distribution of the observed discrete counts through an exact finite-dimensional representation for any Δ. For a single window, the effect of all preceding events is summarized by $X(t_{k-1})$, while the evolution within the window is governed by the same ODE system.

This conditional representation separates the influence of the pre-window history from the events generated within the current window, making it possible to examine how individual edges are reflected in the count distribution. We therefore first study derivatives of the conditional log-PGF to characterize self-loops, adjacency, and direction, and then return to the marginal PGF to establish identifiability directly from the observed count distribution.

3.2 Motivating Example

The conditional log-PGFL preserves complementary information about the underlying causal structure at different arguments. Given a fixed window-entrance state, derivatives at $\boldsymbol{z} = \mathbf{1}$ are conditional factorial cumulants that can reveal adjacency and directional asymmetry. Derivatives at $\boldsymbol{z} = \mathbf{0}$ isolate low-count configurations and provide self-loop and direction signatures. The following examples show these local signals and where self-excitation requires a correction.

Discovering adjacency using $\partial_{z_a}\partial_{z_b}\ell|_{\boldsymbol{z}=\mathbf{1}}$. Consider an isolated two-type Hawkes model with a single directed edge $a \to b$ and no self-loops, so $\alpha_{ab} > 0$ and $\alpha_{ba} = \alpha_{aa} = \alpha_{bb} = 0$. Fix any observation window I_k and condition on its entrance state $X_k := X(t_{k-1}) = (X_a, X_b) \geq \mathbf{0}$. The conditional-window formula of Theorem 2 reduces to

$$\ell_k(z_a, z_b \mid X_k) = \mu_a \int_0^\Delta [z_a \exp\{\alpha_{ab} c_\beta(s)(z_b - 1)\} - 1]\, ds + \mu_b \Delta(z_b - 1) + X_a \alpha_{ab} c_\beta(\Delta)(z_b - 1), \tag{11}$$

where $c_\beta(s) = (1 - e^{-\beta s})/\beta$, which is positive for $s > 0$. The generating variables z_a and z_b mark the counts of types a and b, respectively. Their roles in Equation (11) are asymmetric: z_a marks the type-a root event, whereas z_b also appears inside the exponential term describing its type-b descendants. The final term represents excitation inherited from earlier type-a events and is linear in z_b and independent of z_a. Differentiating once with respect to each generating variable gives

$$K_{ab}^{(2)}(X_k) := \partial_{z_a}\partial_{z_b}\ell_k(\boldsymbol{z} \mid X_k)|_{\boldsymbol{z}=\mathbf{1}} = \mu_a \alpha_{ab} \int_0^\Delta c_\beta(s)\, ds > 0. \tag{12}$$

Thus the mixed second derivative detects dependence between the two types. However, mixed derivatives commute, so $K_{ab}^{(2)}(X_k) = K_{ba}^{(2)}(X_k)$. Therefore this second-order derivative identifies adjacency but cannot determine its direction.

Discovering direction using $\partial_{z_a}\partial_{z_b}^2\ell\big|_{\boldsymbol{z}=\mathbf{1}}$. The asymmetry visible in Equation (11) becomes directional when one generating variable is differentiated twice. For the edge $a \to b$,

$$\begin{aligned} K_{abb}^{(3)}(X_k) &:= \partial_{z_a}\partial_{z_b}^2\ell_k(\boldsymbol{z} \mid X_k)\big|_{\boldsymbol{z}=\mathbf{1}} = \mu_a \alpha_{ab}^2 \int_0^\Delta c_\beta(s)^2\, ds > 0, \\ K_{baa}^{(3)}(X_k) &:= \partial_{z_b}\partial_{z_a}^2\ell_k(\boldsymbol{z} \mid X_k)\big|_{\boldsymbol{z}=\mathbf{1}} = 0. \end{aligned} \tag{13}$$

The first quantity is positive because a type-a root can generate multiple type-b descendants, so repeated differentiation with respect to z_b detects this descendant structure. By contrast, $K_{baa}^{(3)}(X_k) = 0$ because type b has no mechanism for generating additional type-a events. Hence

$$a \to b \quad \Longrightarrow \quad K_{abb}^{(3)}(X_k) > 0, \qquad K_{baa}^{(3)}(X_k) = 0,$$

while reversing the edge exchanges the two roles. In this single-edge model, the entrance-state term is linear in z_b and independent of z_a, so it contributes neither to the mixed second derivative nor to the repeated-index third derivatives. Consequently, these quantities coincide with their empty-history counterparts. This simplification need not hold with self-loops, reciprocal edges, or additional event types.

Discovering self-loops using $\partial_{z_b}^2\ell\big|_{\boldsymbol{z}=\mathbf{0}}$. The PGFL also provides a direct signature of self-excitation. For a general observation window, write X_k for its entrance state and define $D_b(X_k) := \partial_{z_b}^2\ell_k(\boldsymbol{z} \mid X_k)\big|_{\boldsymbol{z}=\mathbf{0}}$. At $\boldsymbol{z} = \mathbf{0}$, the second derivative selects two type-b events with no other observed types. Taking the derivative of the *log*-PGFL removes contributions from independent immigrant clusters. Within one remaining cluster, an indirect return such as $b \to c \to b$ cannot contribute because it also contains a type-c event. Two type-b events in that cluster therefore require the direct self-loop $b \to b$. The conditional PGFL gives

$$D_b(X_k) = 2\alpha_{bb} \int_0^\Delta \nu_b(\Delta - s; X_k) E_b^0(s) v_b(s)\, ds, \tag{14}$$

where $\nu_b(t; X_k) = \mu_b + (A^\top X_k)_b e^{-\beta t} > 0$ is the conditional type-b root rate, $E_b^0(s) > 0$ is the cluster exponential factor at the PGFL origin, and $v_b(s)$ is the corresponding first-order response to the generating variable z_b, positive for $s > 0$. Their explicit forms are given in Appendix B.4. Since all factors inside the integral are positive, $D_b(X_k) = 0 \iff \alpha_{bb} = 0$ and $D_b(X_k) > 0 \iff \alpha_{bb} > 0$. Thus the second derivative of the conditional log-PGFL at the origin distinguishes the absence and presence of the self-loop $b \to b$. Other edges may change the positive weighting in Equation (14), but not this zero–positive criterion.

Discovering direction using $\partial_{z_a}\partial_{z_b}^2\ell\big|_{\boldsymbol{z}=\mathbf{0}}$. Using the same entrance-state notation, define $T_{abb}(X_k) := \partial_{z_a}\partial_{z_b}^2\ell_k(\boldsymbol{z} \mid X_k)\big|_{\boldsymbol{z}=\mathbf{0}}$.

First suppose that the target b has no self-loop. In the isolated single-edge case $a \to b$, with $\alpha_{aa} = \alpha_{ba} = \alpha_{bb} = 0$, the conditional PGFL gives

$$T_{abb}(X_k) = \alpha_{ab}^2 \int_0^\Delta \nu_a(t; X_k) E_a^0(\Delta - t) v_b(\Delta - t)^2\, dt > 0, \qquad T_{baa}(X_k) = 0. \tag{15}$$

Since the integral is strictly positive, T_{abb} is positive exactly because the forward coefficient α_{ab} is present, whereas T_{baa} vanishes when the reverse edge is absent. Reversing the edge exchanges the two quantities. More generally,

when $\alpha_{bb} = 0$, every nonzero contribution to T_{abb} contains α_{ab} as a factor: $T_{abb}(X) = 0 \iff \alpha_{ab} = 0$ and $T_{abb}(X) > 0 \iff \alpha_{ab} > 0$.

When the target b has a self-loop, the same derivative contains one additional structural contribution. In general,

$$T_{abb}(X_k) = \underbrace{\alpha_{ab}\mathcal{Q}_{ab}(X_k)}_{\text{contribution of } a\to b} + \underbrace{2\alpha_{ba}\alpha_{bb}\mathcal{C}_{ab}(X_k)}_{\text{joint contribution of } b\to a \text{ and } b\to b} \tag{16}$$

where $\mathcal{Q}_{ab}(X_k) > 0$ and $\mathcal{C}_{ab}(X_k) > 0$ are positive ODE terms whose explicit forms are given in Appendix B.4.

The second term explains why the raw T_{abb} may remain positive even when $a \to b$ is absent: the reverse edge $b \to a$ together with the self-loop $b \to b$ can generate the same (a, b, b) response. Subtracting this term leaves $\alpha_{ab}\mathcal{Q}_{ab}(X_k)$, which is zero exactly when $\alpha_{ab} = 0$ and strictly positive exactly when $\alpha_{ab} > 0$. Thus the conditional PGFL provides a directional zero–positive criterion in both cases; a target self-loop only requires separating its additional contribution.

3.3 Structural Identifiability

The examples above show that different derivatives of the PGFL encode complementary local features of the causal structure. We now formalize these observations and state the conditions under which each structural signature is valid. The results below characterize self-loops, adjacency, and edge direction locally; global identifiability of a general cyclic support is established subsequently from the complete marginal PGFL.

Theorem 3 (Local PGFL Structural Signatures)**.** *Fix a window I_k and condition on an arbitrary entrance-state realization $X_k = X(t_{k-1}) \geq \mathbf{0}$. Its conditional log-PGFL is $\ell_k(\boldsymbol{z} \mid X_k)$.*

(i) Self-loop. *For any $b \in V$, $D_b(X_k) = \partial_{z_b}^2 \ell_k(\boldsymbol{z} \mid X_k)\big|_{\boldsymbol{z}=\mathbf{0}}$. Then*

$$D_b(X_k) = 0 \iff \alpha_{bb} = 0, \qquad D_b(X_k) > 0 \iff \alpha_{bb} > 0. \tag{17}$$

Thus D_b identifies the presence of the direct self-loop $b \to b$.

(ii) Adjacency. *Consider an isolated two-type model on $\{a, b\}$ with positive baselines, unrestricted self-loops, and the fixed entrance state X_k. $K_{ab}^{(2)}(X_k) = \partial_{z_a}\partial_{z_b}\ell_k(\boldsymbol{z} \mid X_k)|_{\boldsymbol{z}=\mathbf{1}}$. Then*

$$K_{ab}^{(2)}(X_k) = 0 \iff \alpha_{ab} = \alpha_{ba} = 0, \qquad K_{ab}^{(2)}(X_k) > 0 \iff \alpha_{ab} > 0 \text{ or } \alpha_{ba} > 0. \tag{18}$$

Hence $K_{ab}^{(2)}(X_k)$ identifies adjacency but not its direction.

(iii) Direction without self-loops. *Consider an isolated two-type model with positive baselines and fixed entrance state X_k, $\alpha_{aa} = \alpha_{bb} = 0$ and $\alpha_{ab}\alpha_{ba} = 0$. $K_{abb}^{(3)}(X_k) = \partial_{z_a}\partial_{z_b}^2\ell_k(\boldsymbol{z} \mid X_k)\big|_{\boldsymbol{z}=\mathbf{1}}$ and $K_{baa}^{(3)}(X_k) = \partial_{z_b}\partial_{z_a}^2\ell_k(\boldsymbol{z} \mid X_k)\big|_{\boldsymbol{z}=\mathbf{1}}$. Then*

$$\alpha_{ab} > 0 \iff K_{abb}^{(3)}(X_k) > 0,\ K_{baa}^{(3)}(X_k) = 0, \qquad \alpha_{ba} > 0 \iff K_{baa}^{(3)}(X_k) > 0,\ K_{abb}^{(3)}(X_k) = 0. \tag{19}$$

Under these single-edge assumptions, both repeated-index mixed derivatives are independent of X_k; this does not extend to a general cyclic graph.

(iv) Direction. *For $a \neq b$, $T_{abb}(X_k) = \partial_{z_a}\partial_{z_b}^2\ell_k(\boldsymbol{z} \mid X_k)\big|_{\boldsymbol{z}=\mathbf{0}}$. If the target b has no self-loop, $\alpha_{bb} = 0$, then*

$$T_{abb}(X_k) = \alpha_{ab}\left[\alpha_{ab}\mathcal{R}_{ab}(X_k) + 2\alpha_{ba}\mathcal{S}_{ab}(X_k)\right], \qquad \mathcal{R}_{ab}(X_k), \mathcal{S}_{ab}(X_k) > 0, \tag{20}$$

and consequently

$$T_{abb}(X_k) = 0 \iff \alpha_{ab} = 0, \qquad T_{abb}(X_k) > 0 \iff \alpha_{ab} > 0. \tag{21}$$

If the target has a self-loop, $\alpha_{bb} > 0$, then

$$T_{abb}(X_k) = \alpha_{ab}\mathcal{Q}_{ab}(X_k) + 2\alpha_{ba}\alpha_{bb}\mathcal{C}_{ab}(X_k), \qquad \mathcal{Q}_{ab}(X_k), \mathcal{C}_{ab}(X_k) > 0. \tag{22}$$

The second term is the additional contribution generated jointly by $b \to a$ and $b \to b$. Define the corrected quantity $T_{abb}^{\mathrm{corr}}(X_k) := T_{abb}(X_k) - 2\alpha_{ba}\alpha_{bb}\mathcal{C}_{ab}(X_k)$. Then

$$T_{abb}^{\mathrm{corr}}(X_k) = \alpha_{ab}\mathcal{Q}_{ab}(X_k), \qquad T_{abb}^{\mathrm{corr}}(X_k) = 0 \iff \alpha_{ab} = 0, \qquad T_{abb}^{\mathrm{corr}}(X_k) > 0 \iff \alpha_{ab} > 0. \tag{23}$$

The positive ODE terms $\mathcal{R}_{ab}$, $\mathcal{S}_{ab}$, $\mathcal{Q}_{ab}$, and $\mathcal{C}_{ab}$ are given explicitly in Appendix B.4.

Parts (i) and (iv) follow from the conditional PGFL analysis in Appendix B.4, while Parts (ii) and (iii) are proved in Appendix B.3. These are population criteria at a fixed entrance state; the algorithm does not explicitly estimate $\mathcal{C}_{ab}$. To move from the conditional local signatures of Theorem 3 to global support identifiability, we examine the marginal log-PGFL along negative coordinate axes. For each model and a fixed window I_k, define

$$L_b(q) := \ell_k(q\boldsymbol{e}_b), \qquad H_{a|b}(q) := \partial_{z_a}\ell_k(\boldsymbol{z})\big|_{\boldsymbol{z}=q\boldsymbol{e}_b}, \qquad q < 0, \quad a \neq b, \tag{24}$$

where $\boldsymbol{e}_b$ is the bth canonical basis vector. When $\boldsymbol{z} = q\boldsymbol{e}_b$ with $q < 0$, these functions use the unique analytic extension established in Appendix B.5 and are therefore determined by the marginal window PGFL. The following theorem shows that their behavior as $q \to -\infty$ distinguishes self-loops and ordered edges, so equality of the marginal window PGFL implies equality of the complete directed support. Together, these criteria determine the complete directed support, leading to the following global identifiability theorem.

Theorem 4 (Global Causal-Structure Identifiability)**.** *Consider two d-variate exponential Hawkes models $\theta = (\boldsymbol{\mu}, A, \beta)$ and $\widetilde{\theta} = (\widetilde{\boldsymbol{\mu}}, \widetilde{A}, \widetilde{\beta})$ satisfying Assumption 1, strictly positive finite baselines and decay rates, and finite nonnegative excitation matrices. For fixed k and $\Delta > 0$, suppose $\ell_k(\boldsymbol{z}; \theta) = \ell_k(\boldsymbol{z}; \widetilde{\theta})$ for all $\boldsymbol{z} \in [0,1]^d$.*

For either model, the self-loop status is characterized by

$$\alpha_{bb} = 0 \iff \lim_{q \to -\infty} \frac{L_b(q)}{|q|} < 0, \qquad \alpha_{bb} > 0 \iff \lim_{q \to -\infty} \frac{L_b(q)}{|q|} = 0. \tag{25}$$

If $\alpha_{bb} = 0$, then, for finite constants $C_+, C_- > 0$,

$$H_{a|b}(q) \xrightarrow[q \to -\infty]{} \begin{cases} C_+, & \alpha_{ab} = 0,\ \alpha_{ba} = 0, \\ -\infty \text{ linearly in } |q|, & \alpha_{ab} = 0,\ \alpha_{ba} > 0, \\ 0, & \alpha_{ab} > 0,\ \alpha_{ba} = 0, \\ -C_-, & \alpha_{ab} > 0,\ \alpha_{ba} > 0. \end{cases} \tag{26}$$

These cases distinguish no edge, $b \to a$, $a \to b$, and $a \leftrightarrow b$, respectively. If $\alpha_{bb} > 0$, the edge $a \to b$ is distinguished by

$$\begin{cases} \lim_{q \to -\infty} \log|q|\, |H_{a|b}(q)| = 0, & \alpha_{ab} > 0, \\ \liminf_{q \to -\infty} \log|q|\, |H_{a|b}(q)| > 0, & \alpha_{ab} = 0. \end{cases} \tag{27}$$

Equality of the window-count laws implies equality of these witnesses by analytic continuation. Applying the classifications to all nodes and ordered pairs therefore gives $\operatorname{supp}(A) = \operatorname{supp}(\widetilde{A})$, including self-loops, reciprocal edges, and longer directed cycles.

The analytic-continuation argument, the treatment of the random entrance state, and the derivation of the negative-axis asymptotic classes are given in Appendix B.5.

4 Algorithm

The observation grid is fixed independently of the realized events. For a width-Δ partition, let $Y_{v,rk}$ denote the number of type-v events in window k of independent record r, including zero-count windows, and write $\mathcal{X} = \{Y_{v,rk}\}$ for the observed count data. Raw timestamps are used only to assign events to windows; C-SHP then discards all within-window offsets and event order and operates on $\mathcal{X}$ over the common discrete grid. The observable distribution is therefore the joint law of these window counts. Appendix C specifies the fixed decay and calibration parameters, so Algorithm 1 lists only $\mathcal{X}$ as input. The fitted parameters are $\Theta = (\boldsymbol{\mu}, A)$, and the graph may contain self-loops, reciprocal edges, and longer directed cycles.

Let m be the lag between two bins and $x_\Delta = \beta\Delta$. Temporal aggregation conceals the exact occurrence times of events within each bin, so the continuous-time delay t in the exponential excitation kernel is not observable from the count data. To retain the expected excitation carried by each bin pair, we adopt a working uniform distribution for the unobserved within-bin event positions and average the continuous-time kernel $\phi_{uv}(t) = \alpha_{uv} e^{-\beta t}$. This yields the discrete-time excitation kernel used for fitting,

$$\phi_{uv}^{\mathrm{DT}}(m) = \alpha_{uv} \begin{cases} c_0(x_\Delta), & m = 0, \\ a(x_\Delta) e^{-m x_\Delta}, & m \geq 1, \end{cases} \quad a(x_\Delta) = \frac{(e^{x_\Delta} - 1)(1 - e^{-x_\Delta})}{x_\Delta^2}, \quad c_0(x_\Delta) = \frac{1}{x_\Delta}\left(1 - \frac{1 - e^{-x_\Delta}}{x_\Delta}\right). \tag{28}$$

Within the same observation bin, only source events that occur before the target may contribute; for $u = v$, the target event itself is excluded, while earlier events of the same type still contribute through self-excitation. The resulting discrete-time criterion is

$$\mathcal{L}(\Theta; \mathcal{X}) = \sum_v \sum_{(r,k) \in \mathcal{I}_v} Y_{v,rk} \log \bar{\lambda}_{v,rk} - T_\Sigma \sum_v \mu_v - \sum_{u,v} c_u \alpha_{uv}, \tag{29}$$

where $Y_{v,rk}$ is the type-v count in bin k of record r, $\mathcal{I}_v$ indexes the retained target bins, T_Σ is the total observation time, and $\bar{\lambda}_{v,rk} = \mu_v + \sum_u \alpha_{uv} s_{u,rk}^{(v)}$. The regressor $s_{u,rk}^{(v)}$ sums the discrete-time kernel over source counts, while c_u is the corresponding exposure; both are specified in Appendix C.1. This working criterion is distinct from the exact count likelihood characterized by the PGFL. Following the responsibility construction of SHP (Qiao et al., 2023), the fixed-graph updates at MM iteration h are

$$\mu_v^{(h+1)} = \frac{\mu_v^{(h)}}{T_\Sigma} \sum_{(r,k) \in \mathcal{I}_v} \frac{Y_{v,rk}}{\bar{\lambda}_{v,rk}^{(h)}}, \qquad \alpha_{uv}^{(h+1)} = \frac{\alpha_{uv}^{(h)}}{c_u} \sum_{(r,k) \in \mathcal{I}_v} \frac{Y_{v,rk} s_{u,rk}^{(v)}}{\bar{\lambda}_{v,rk}^{(h)}}. \tag{30}$$

These are EM-style MM updates for $\mathcal{L}$, rather than exact EM updates for latent event times. Each candidate graph is refitted and evaluated by

$$S_{\mathrm{BIC}}(G) = \frac{\mathcal{L}_G}{N_{\mathrm{ev}}} - \omega_{\mathrm{BIC}} \frac{(d + |E_G|) \log N_{\mathrm{ev}}}{2N_{\mathrm{ev}}}, \tag{31}$$

where $\mathcal{L}_G$ is the maximized criterion under G, E_G is its directed edge set, N_{ev} is the total event count, and $\omega_{\mathrm{BIC}} > 0$ is fixed.

The population PGFL quantities in Section 3 cannot be evaluated exactly from finite discrete-time observations because the entrance state and within-bin event locations are unobserved. For the conditional D/T coordinates, we use the count-based approximation

$$\widetilde{X}_{u,r,k+1} = e^{-\beta\Delta} \widetilde{X}_{u,r,k} + \frac{1 - e^{-\beta\Delta}}{\beta\Delta} Y_{u,rk}, \qquad \widetilde{X}_{u,r,1} = 0. \tag{32}$$

It averages the unknown contributions of events within each bin and generally differs from the exact state X_k. The remaining PGFL coordinates are estimated directly from the observed count distributions. These finite-sample approximations guide the search process rather than strictly replicating population-level theoretical proofs; Appendix C details their estimation, calibration, and acceptance rules. In Algorithm 1, the "hat" notation denotes the finite-sample estimate of the corresponding quantity discussed in the main text. Specifically, the repeated-index derivative $K^{(3)}_{abb}(X_k)$ in Theorem 3 is a conditional factorial cumulant of the window counts given the entrance state X_k. Because this latent state is not observed, the algorithm applies the same factorial-cumulant functional to the empirical marginal distribution of the window counts. The resulting statistic $\widehat{K}^{(3)}_{abb}$ is therefore the observable marginal counterpart of the conditional quantity in Theorem 3 and is used as a finite-sample orientation signal.

Algorithm 1: Learning Causal Structure Using C-SHP

Input: Dataset $\mathcal{X}$.
Output: $G^\star, \Theta^\star$.
$Y_{v,rk} \leftarrow N_{v,r}(I_k)$; $G \leftarrow$ empty graph.
repeat
  for every $\{a, b\}$ in decreasing $|\widehat{K}^{(2)}_{ab}|$ **do**
    $G'_{\{a,b\}} \in \{0, a \to b, b \to a, a \leftrightarrow b\}$.
    Update $\Theta_{G'}$ by Eq. (30); record $S_{\mathrm{BIC}}(G')$.
    $G \leftarrow G'$ with the largest strict improvement.
  end for
until no state changes or the search budget is exhausted.
for every $b \in V$ **do**
  Evaluate $\widehat{D}_b$ and $\widehat{\alpha}_{bb}/\beta$; update G_{bb}.
end for
for every $a \leftrightarrow b$ in G **do**
  Evaluate $\widehat{T}_{abb}, \widehat{T}_{baa}, \widehat{H}_{a|b}, \widehat{H}_{b|a}$ and S_{BIC}.
  Refine $G_{\{a,b\}}$ only when the evidence agrees.
  Evaluate $\widehat{K}^{(3)}_{abb}$ and $\widehat{K}^{(3)}_{baa}$; orient or retain $a \leftrightarrow b$.
end for
Refit $\Theta^\star$ by Eq. (30); **return** $(G^\star, \Theta^\star)$.

The off-diagonal elements of the initial graph are all zero, while the diagonal elements are temporarily active during the initial fitting phase. The empirical quantity K_2 determines only the order in which data pairs are examined. Each candidate must be refitted using Equation (30) and must pass a stability test. Evidence regarding D, T, H, and the repeated-index quantity K_3 is applied conservatively: if the comparison is inconclusive, the existing reciprocal state is retained. Appendix C specifies the relevant estimators, bootstrap calibration procedures, acceptance rules, and numerical thresholds.

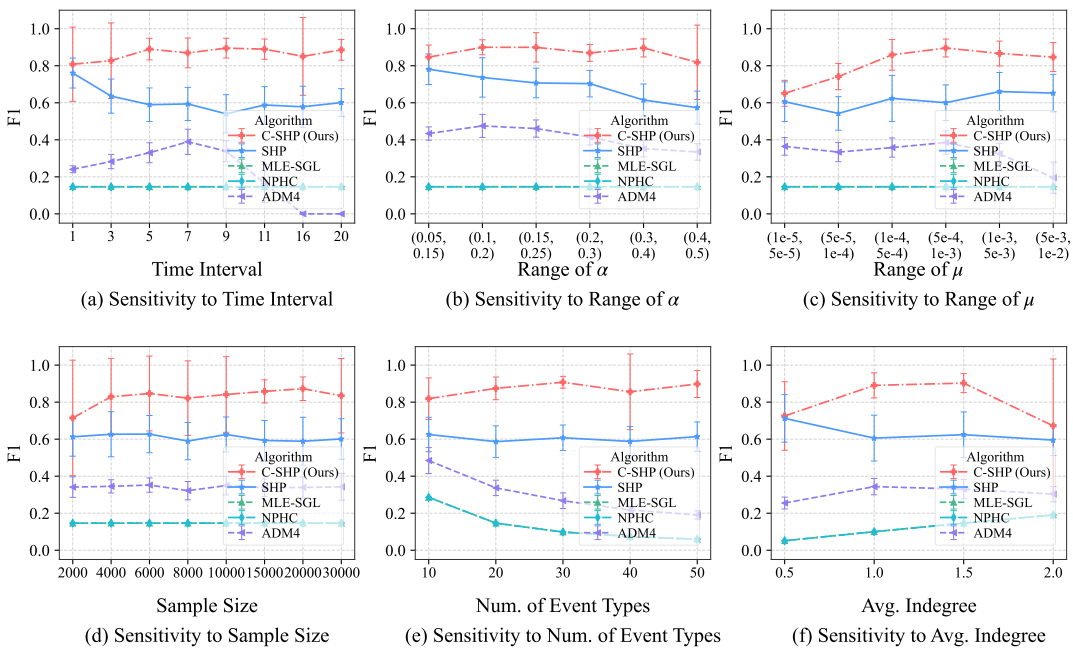
5 Experiments

We compare C-SHP with SHP (Qiao et al., 2023), MLE-SGL (Xu et al., 2016), NPHC (Achab et al., 2017), and ADM4 (Zhou et al., 2013) in both the synthetic and real-world experiments.

Synthetic experiments. We design the synthetic experiments to systematically examine causal structure recovery under different degrees of temporal aggregation, Hawkes process parameters, sample sizes, and graph complexities. For each random seed, we generate a cyclic directed graph by introducing a randomly selected 2-cycle or 3-cycle and additionally assign self-loops to $\max\{1, \mathrm{round}(0.2d)\}$ randomly selected nodes. We then simulate an exponential Hawkes process and aggregate the events into intervals of width Δ.

The default configuration contains 20,000 events and 20 event types, with interval width $\Delta = 5$, off-diagonal average indegree 1.5, excitation strengths $\alpha_{uv} \in [0.3, 0.5]$ for present edges, and immigration intensities $\mu_v \in [10^{-4}, 5 \times 10^{-4}]$. We vary one experimental factor at a time while keeping the others fixed, as shown in Figure 2. Each configuration is evaluated over 20 random seeds, and we report the mean off-diagonal directed F1 score with across-seed error bars.

Figure 2 shows that C-SHP achieves the highest mean F1 in all settings. It remains stable as wider intervals reduce temporal information (a) and performs consistently across excitation and immigration ranges (b–c), improving as higher immigration alleviates data sparsity. It also leads across sample sizes despite greater variation in the smallest-sample regime (d), and degrades only moderately as the number of event types increases (e). Although denser graphs are more challenging, C-SHP retains the best mean F1 at each tested average indegree (f). These results demonstrate robustness across all six factors in cyclic graphs with self-loops.



(a) Sensitivity to Time Interval (b) Sensitivity to Range of α (c) Sensitivity to Range of μ

(d) Sensitivity to Sample Size (e) Sensitivity to Num. of Event Types (f) Sensitivity to Avg. Indegree

Figure 2: F1 in the Sensitivity Experiments.

Table 1: F1 in the Real-World Experiment.

Method	Precision ↑	Recall ↑	F1 ↑	SHD ↓
C-SHP (Ours)	**0.8889**	0.8889	**0.8889**	**2**
SHP	0.0000	0.0000	0.0000	10
ADM4	0.5000	0.4444	0.4706	9
NPHC	0.4500	**1.0000**	0.6207	11
MLE-SGL	0.4500	**1.0000**	0.6207	11

Real-world experiments. We consider a dataset from the NeurIPS 2023 competition on event sequences (CSL Competition Organizing Team, 2023). To examine cyclic structure recovery, we retain five alarm types whose induced reference graph contains directed cycles. The data contain 18,873 alarm records across 100 devices; the selected reference graph has nine directed edges and no self-loops. We treat each device as an independent event sequence and bin alarm start times at width $\Delta = 9$ in the dataset's timestamp units. C-SHP uses decay $\beta = 0.1$.

Table 1 compares directed-edge recovery on the retained dataset. Precision, recall, and F1 use directed edges, and SHD counts disagreements between the estimated and reference adjacency matrices.

C-SHP recovers eight of the nine reference edges with one false positive, giving F1 $= 0.8889$ and SHD $= 2$. NPHC and MLE-SGL attain full recall but introduce more false positives.

6 Conclusion

In this work, we study how to learn cyclic causal structures from discrete-time event sequences when the exact temporal ordering of events within the same observation interval is unavailable. We propose Cyclic Structural Hawkes Processes, which retain the cyclic excitation mechanism of an underlying continuous-time Hawkes process while enabling causal discovery from discretely observed event counts. By exploiting the probability generating functional, we derive a series of theoretical tools to characterize the local structural signatures of self-loops, adjacency, and edge direction. Based on these results, we develop a practical causal discovery procedure. The proposed C-SHP framework extends causal structure learning for low-resolution event sequences beyond the acyclic setting and provides a principled way to recover cyclic interactions from temporally aggregated observations. Experiments on cyclic graphs demonstrate that the algorithm exhibits robust performance across varying time resolutions, model parameters, sample sizes, graph dimensions, and graph densities.

REFERENCES

Massil Achab, Emmanuel Bacry, Stéphane Gaïffas, Iacopo Mastromatteo, and Jean-François Muzy. Uncovering causality from multivariate Hawkes integrated cumulants. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1–10, 2017. URL https://proceedings.mlr.press/v70/achab17a.html.

Emmanuel Bacry, Iacopo Mastromatteo, and Jean-François Muzy. Hawkes processes in finance. *Market Microstructure and Liquidity*, 1(1):1550005, 2015.

Ruichu Cai, Siyu Wu, Jie Qiao, Zhifeng Hao, Keli Zhang, and Xi Zhang. THPs: Topological Hawkes processes for learning causal structure on event sequences. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. doi: 10.1109/TNNLS.2022.3175622.

Ruichu Cai, Zhiyi Huang, Wei Chen, Zhifeng Hao, and Kun Zhang. Causal discovery with latent confounders based on higher-order cumulants. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 3380–3407, 2023. URL https://proceedings.mlr.press/v202/cai23a.html.

CSL Competition Organizing Team. CSL-ESPK Dataset 4. NeurIPS 2023 Causal Structure Learning from Event Sequences and Prior Knowledge Competition, 2023. URL https://github.com/gcastle-hub/dataset/tree/master/csl_espk/dataset_4.

Alan G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971. doi: 10.1093/biomet/58.1.83.

Matthias Kirchner. Hawkes and INAR(∞) processes. *Stochastic Processes and their Applications*, 126(8):2494–2525, 2016. doi: 10.1016/j.spa.2016.02.008.

Yuequn Liu, Ruichu Cai, Wei Chen, Jie Qiao, Yuguang Yan, Zijian Li, Keli Zhang, and Zhifeng Hao. TNPAR: Topological neural poisson auto-regressive model for learning Granger causal structure from event sequences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 20491–20499, 2024. doi: 10.1609/aaai.v38i18.30033.

Joris M. Mooij and Tom Claassen. Constraint-based causal discovery using partial ancestral graphs in the presence of cycles. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence*, volume 124 of *Proceedings of Machine Learning Research*, pp. 1159–1168, 2020. URL https://proceedings.mlr.press/v124/m-mooij20a.html.

Jie Qiao, Ruichu Cai, Siyu Wu, Yu Xiang, Keli Zhang, and Zhifeng Hao. Structural Hawkes processes for learning causal structure from discrete-time event sequences. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 5702–5710, 2023. doi: 10.24963/ijcai.2023/633.

Jie Qiao, Yu Xiang, Zhengming Chen, Ruichu Cai, and Zhifeng Hao. Causal discovery from poisson branching structural causal model using high-order cumulant with path analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 20524–20531, 2024. doi: 10.1609/aaai.v38i18.30037.

Yu Xiang, Jie Qiao, Zhefeng Liang, Zihuai Zeng, Ruichu Cai, and Zhifeng Hao. On the identifiability of poisson branching structural causal model using probability generating function. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Hongteng Xu, Mehrdad Farajtabar, and Hongyuan Zha. Learning Granger causality for Hawkes processes. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1717–1726, 2016. URL https://proceedings.mlr.press/v48/xuc16.html.

Xiufan Yu, Karthikeyan Shanmugam, Debarun Bhattacharjya, Tian Gao, Dharmashankar Subramanian, and Lingzhou Xue. Hawkesian graphical event models. In *Proceedings of the 10th International Conference on Probabilistic Graphical Models*, volume 138 of *Proceedings of Machine Learning Research*, pp. 569–580, 2020. URL https://proceedings.mlr.press/v138/yu20a.html.

Ke Zhou, Hongyuan Zha, and Le Song. Learning social infectivity in sparse low-rank networks using multi-dimensional hawkes processes. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, pp. 641–649, 2013.

A Related Work

Continuous-time event-sequence discovery. Hawkes-process methods encode directed influence through excitation kernels and learn Granger-causal structure from accurately ordered events (Xu et al., 2016). Related approaches recover integrated influence from second- and third-order cumulants (Achab et al., 2017), represent temporal dependence with graphical event models (Yu et al., 2020), or incorporate a known topological network when event sequences are dependent (Cai et al., 2022). These methods exploit temporal information that is only partially available after events are aggregated into discrete observation intervals. We instead study whether the directed Hawkes support remains identifiable from the resulting count law.

Causal discovery from discrete-time event sequences. Temporal aggregation connects Hawkes processes to integer-valued autoregression (Kirchner, 2016), while TNPAR uses neural count autoregression for event-sequence causal discovery (Liu et al., 2024). SHP is the closest setting to ours because it models dependence within an observation interval, but its recursive construction and identifiability result assume an acyclic instantaneous graph under causal faithfulness and causal sufficiency (Qiao et al., 2023). C-SHP retains the discrete-time observation scheme while allowing the underlying Hawkes type graph to contain directed cycles, reciprocal edges, and self-loops.

Cyclic causal discovery. Feedback is also studied in general structural causal models. Constraint-based methods, for example, can characterize cyclic structures under suitable Markov and faithfulness conditions (Mooij & Claassen, 2020). Those results concern conditional-independence relations among static variables. Our variables are event types in a continuous-time branching process, whereas the observations are discrete-time counts, so our identifiability analysis instead starts from the count distribution induced by temporal aggregation.

Generating-function identifiability. High-order cumulants provide causal-order information in latent-variable models (Cai et al., 2023) and path information in Poisson branching structural causal models (Qiao et al., 2024). PGFL analysis also identifies local structures in the latter model (Xiang et al., 2024). Our result concerns a different distribution: the exact discrete-time count law induced by a continuous-time Hawkes process whose type graph may be cyclic. We use local PGFL derivatives and factorial cumulants as finite-sample evidence, while the global identifiability result relies on the complete count PGFL.

AI Assistance

We used large language models as general-purpose assistive tools for improving grammar, wording, and clarity, as well as for brainstorming intermediate derivation steps and suggesting occasional code snippets. The models did not independently determine the research problem, algorithm or model design, experimental protocol, data generation process, analysis, or conclusions. Our substantive methodological decisions, technical content, and core implementations were independently developed, examined, and verified. We reviewed and revised all AI-assisted text and code and take full responsibility for every part of the manuscript and its accompanying artifacts.

B Proofs of the Main Theoretical Results

The proofs follow the order of the main text. We first derive the continuous and discrete-window PGFL representations. We then prove the local K_2, K_3, D, and T signatures, followed by global support identifiability. The section ends with the discrete-time fitting criterion and its algorithmic details.

B.1 Proof of Theorem 1

Proof. Single-root cluster PGFL. Fix a type-v' root at time t. Let $\mathcal{C}_{v'}(t)$ denote the marked counting measure of the root and all of its descendants that occur by time T. For each $v \in V$, let $\mathcal{N}^{(1)}_{v|v'}(t)$ be the set of times of the direct type-v children generated by this root. The cluster decomposes recursively as

$$\mathcal{C}_{v'}(t) = \delta_{(t,v')} + \sum_{v\in V} \sum_{\tau_n \in \mathcal{N}^{(1)}_{v|v'}(t)} \mathcal{C}_v(\tau_n),$$

where $\delta_{(t,v')}$ represents the root and $\mathcal{C}_v(\tau_n)$ is the descendant cluster initiated by a type-v child born at time τ_n.

Conditional on the root, $\mathcal{N}^{(1)}_{v|v'}(t)$ is a nonhomogeneous Poisson process on $(t, T]$ with intensity

$$\alpha_{v'v} e^{-\beta(\tau-t)}, \qquad \tau \in (t, T].$$

Conditional on all first-generation child times and types, the descendant clusters initiated by distinct children are independent. The root therefore contributes $z_{v'}(t)$, while a type-v child born at τ_n contributes $F_v(\tau_n; \boldsymbol{z})$. Hence

$$F_{v'}(t; \boldsymbol{z}) = z_{v'}(t) \prod_{v \in V} \mathbb{E}\left[\prod_{\tau_n \in \mathcal{N}^{(1)}_{v|v'}(t)} F_v(\tau_n; \boldsymbol{z}) \right].$$

For a Poisson process with intensity $\lambda(\tau)$, its PGFL satisfies

$$\mathbb{E}\left[\prod_{\tau_n} f(\tau_n) \right] = \exp\left\{ \int \lambda(\tau)[f(\tau) - 1] \, \mathrm{d}\tau \right\}.$$

Applying this identity to the direct type-v children gives

$$\mathbb{E}\left[\prod_{\tau_n \in \mathcal{N}^{(1)}_{v|v'}(t)} F_v(\tau_n; \boldsymbol{z}) \right] = \exp\left\{ \int_t^T \alpha_{v'v} e^{-\beta(\tau - t)} [F_v(\tau; \boldsymbol{z}) - 1] \, \mathrm{d}\tau \right\}.$$

Multiplying these independent contributions over all $v \in V$ yields

$$F_{v'}(t; \boldsymbol{z}) = z_{v'}(t) \exp\left\{ \sum_{v \in V} \alpha_{v'v} \int_t^T e^{-\beta(\tau - t)} [F_v(\tau; \boldsymbol{z}) - 1] \, \mathrm{d}\tau \right\},$$

which proves Eq. (3).

Global PGFL. For the empty-history process, let $\mathcal{I}_{v'}$ be the set of type-v' immigrant times in $(0, T]$. It is a homogeneous Poisson process with rate $\mu_{v'}$. The complete Hawkes process is the superposition of the clusters initiated by these immigrants:

$$\mathcal{N} = \sum_{v' \in V} \sum_{t_n \in \mathcal{I}_{v'}} \mathcal{C}_{v'}(t_n).$$

Clusters initiated by distinct immigrants are independent. Conditional on all immigrant times, the PGFL contribution of an immigrant of type v' at time t_n is $F_{v'}(t_n; \boldsymbol{z})$. Therefore

$$G_T[\boldsymbol{z}] = \prod_{v' \in V} \mathbb{E}\left[\prod_{t_n \in \mathcal{I}_{v'}} F_{v'}(t_n; \boldsymbol{z}) \right].$$

Applying the Poisson PGFL identity to the type-v' immigrant process gives

$$\mathbb{E}\left[\prod_{t_n \in \mathcal{I}_{v'}} F_{v'}(t_n; \boldsymbol{z}) \right] = \exp\left\{ \mu_{v'} \int_0^T [F_{v'}(t; \boldsymbol{z}) - 1] \, \mathrm{d}t \right\}.$$

Finally, multiplying over $v' \in V$ gives

$$G_T[\boldsymbol{z}] = \exp\left\{ \sum_{v' \in V} \mu_{v'} \int_0^T [F_{v'}(t; \boldsymbol{z}) - 1] \, \mathrm{d}t \right\},$$

which proves Eq. (4). □

B.2 Proof of Theorem 2

Proof. We first connect the continuous-time PGFL to the observed window counts. Because $z_v(t) = z_{v,k}$ for $t \in I_k$, every realization satisfies

$$\prod_v \prod_{x \in N_v \cap (0,T]} z_v(x) = \prod_{k=1}^K \prod_v z_{v,k}^{Y_{v,k}}.$$

Taking expectations gives $G_T[\boldsymbol{z}] = G_K(\boldsymbol{Z})$.

We next differentiate the convolution state defined before Theorem 2. By Leibniz' rule,

$$\dot{U}^{(k)}_{v'}(s) = F_{v'}(t_k - s; \boldsymbol{Z}) - 1 - \beta U^{(k)}_{v'}(s).$$

At time $t_k - s \in I_k$, Eq. (3) becomes

$$F_{v'}(t_k - s; \boldsymbol{Z}) = z_{v',k} \exp\left\{ \sum_{v \in V} \alpha_{v'v} U^{(k)}_v(s) \right\}.$$

Substituting this expression gives the first ODE in Eq. (7).

The accumulated cluster contribution over the first s units of the reverse-time interval is

$$J_{v'}^{(k)}(s) = \int_0^s z_{v',k} \exp\left\{ \sum_v \alpha_{v'v} U_v^{(k)}(y) \right\} \mathrm{d}y.$$

The fundamental theorem of calculus gives the second ODE and $J_{v'}^{(k)}(0) = 0$.

The remaining boundary conditions follow directly from the integral definition of $U_{v'}^{(k)}$. For the final interval,

$$U_{v'}^{(K)}(0) = \int_T^T e^{-\beta(\tau - T)} [F_{v'}(\tau; \boldsymbol{Z}) - 1] \, \mathrm{d}\tau = 0.$$

For $k < K$, both $U_{v'}^{(k)}(0)$ and $U_{v'}^{(k+1)}(\Delta)$ equal

$$\int_{t_k}^T e^{-\beta(\tau - t_k)} [F_{v'}(\tau; \boldsymbol{Z}) - 1] \, \mathrm{d}\tau,$$

which proves the backward propagation rule.

At $s = \Delta$, changing variables from s to $t = t_k - s$ gives

$$J_{v'}^{(k)}(\Delta) = \int_{t_{k-1}}^{t_k} F_{v'}(t; \boldsymbol{Z}) \, \mathrm{d}t.$$

Splitting Eq. (4) over the K intervals therefore gives

$$\ell_K(\boldsymbol{Z}) = \sum_{k=1}^K \sum_{v'} \mu_{v'} [J_{v'}^{(k)}(\Delta) - \Delta],$$

which is Eq. (8).

It remains to prove the two single-window formulas. Fix I_k and solve the same ODE afresh on that interval, as specified in the theorem. Conditional on $X(t_{k-1})$, inherited excitation decays exponentially. Thus the time-dependent exogenous rate of type v' at elapsed time $r \in (0, \Delta]$ is

$$\mu_{v'} + e^{-\beta r} \big(A^\top X(t_{k-1}) \big)_{v'}.$$

Applying the cluster PGFL argument from Theorem 1 to these exogenous events gives

$$\begin{aligned} \ell_k(\boldsymbol{z} \mid X(t_{k-1})) = \sum_{v' \in V} \int_0^\Delta & \left[\mu_{v'} + e^{-\beta r} \big(A^\top X(t_{k-1}) \big)_{v'} \right] \\ & \times [F_{v'}(t_{k-1} + r; \boldsymbol{z}) - 1] \, \mathrm{d}r. \end{aligned} \tag{B.1}$$

For this single-window solve,

$$J_{v'}^{(k)}(\Delta) - \Delta = \int_0^\Delta [F_{v'}(t_{k-1} + r; \boldsymbol{z}) - 1] \, \mathrm{d}r$$

and

$$U_{v'}^{(k)}(\Delta) = \int_0^\Delta e^{-\beta r} [F_{v'}(t_{k-1} + r; \boldsymbol{z}) - 1] \, \mathrm{d}r.$$

Substitution into Eq. (B.1) yields

$$\ell_k(\boldsymbol{z} \mid X(t_{k-1})) = \sum_{v' \in V} \mu_{v'} [J_{v'}^{(k)}(\Delta) - \Delta] + X(t_{k-1})^\top A \boldsymbol{U}^{(k)}(\Delta),$$

which proves Eq. (9).

Finally, the tower property gives

$$\ell_k(\boldsymbol{z}) = \log \mathbb{E}[\exp\{\ell_k(\boldsymbol{z} \mid X(t_{k-1}))\}],$$

and substituting the conditional formula proves Eq. (10). For $\boldsymbol{Z} \in [0,1]^{d \times K}$, every cluster PGFL lies in $[0, 1]$, so the states are bounded. The ODE vector field is locally Lipschitz on this bounded region, which gives uniqueness. □

B.3 Proof of Theorem 3, Parts (ii) and (iii)

Fix I_k and condition on X_k, as in the Motivating Example. From Eq. (B.1),

$$K^{(2)}_{ab}(X_k) = \sum_{v\in V} \int_0^\Delta \nu_v(t; X_k)\, \partial_{z_a}\partial_{z_b} F_v(t_{k-1}+t; \boldsymbol{z})|_{\boldsymbol{z}=\mathbf{1}}\, dt. \tag{B.2}$$

Every term is nonnegative because a PGFL has nonnegative coefficients. In the isolated two-type model, if $\alpha_{ab} = \alpha_{ba} = 0$, the two types generate separate clusters and every mixed derivative is zero. If $\alpha_{ab} > 0$, a type-a root has positive probability of producing a type-b child before the window ends; since $\nu_a(t; X_k) \geq \mu_a > 0$, the corresponding term is positive on an interval. The case $\alpha_{ba} > 0$ is symmetric. Thus $K^{(2)}_{ab}(X_k)$ has the zero–positive behavior in Part (ii), while equality of mixed derivatives shows that it cannot determine direction.

For Part (iii), specialize further to the single-edge, no-self-loop model $a \to b$. The conditional log-PGFL is exactly Eq. (11). Direct differentiation, as in the Motivating Example, gives

$$K^{(3)}_{abb}(X_k) = \mu_a \alpha^2_{ab} \int_0^\Delta c_\beta(s)^2\, ds > 0, \qquad K^{(3)}_{baa}(X_k) = 0.$$

Indeed, z_b appears in the descendant exponential, whereas z_a appears only linearly. The entrance-state term is linear in z_b and independent of z_a, so it contributes to neither repeated-index derivative. Reversing the edge exchanges the two expressions; with neither edge, both vanish. This proves Part (iii).

B.4 Conditional Self-Loop and Directional Derivatives

We prove Parts (i) and (iv) of Theorem 3 in the same way as the Motivating Example: specialize the conditional cluster PGFL, differentiate at $\boldsymbol{z} = \mathbf{0}$, and isolate the coefficient whose presence is being tested. Fix I_k, condition on $X_k = X(t_{k-1})$, and measure time from the left endpoint of the window. Earlier events and baseline immigrants together produce type-i roots at rate

$$\nu_i(t; X_k) = \mu_i + (A^\top X_k)_i e^{-\beta t} > 0. \tag{B.3}$$

As in the main text, set

$$\tau(s) = \frac{1-e^{-\beta s}}{\beta}, \qquad E^0_i(s) = \exp\left\{-\tau(s)\sum_j \alpha_{ij}\right\}, \qquad v_i(s) = \int_0^s e^{-\beta(s-u)} E^0_i(u).$$

Here $v_i(0) = 0$ and $v_i(s) > 0$ for $s > 0$. At the PGFL origin, the one-window ODE satisfies $U^{(k)}_i(s; \mathbf{0}) = -\tau(s)$ and

$$\partial_{z_b} U^{(k)}_i(s; \boldsymbol{z})\Big|_{\boldsymbol{z}=\mathbf{0}} = \mathbf{1}_{\{i=b\}} v_b(s). \tag{B.4}$$

Self-loop derivative. In this proof, $F_i(s; \boldsymbol{z})$ abbreviates $F_i(t_k - s; \boldsymbol{Z}^{[k]}(\boldsymbol{z}))$ in reverse time. Two derivatives with respect to z_b retain two type-b events and no other observed type. For a type-b root, direct differentiation of the cluster PGFL gives

$$\partial^2_{z_b} F_b(s; \boldsymbol{z})\big|_{\boldsymbol{z}=\mathbf{0}} = 2\alpha_{bb} E^0_b(s) v_b(s).$$

Superposing roots according to Eq. (B.3) therefore yields

$$D_b(X_k) = 2\alpha_{bb} \int_0^\Delta \nu_b(t; X_k) E^0_b(\Delta - t) v_b(\Delta - t)\, dt. \tag{B.5}$$

The integral is strictly positive because ν_b, E^0_b, and v_b are positive away from an endpoint. Hence $D_b(X_k) = 0$ exactly when $\alpha_{bb} = 0$, and it is positive exactly when $\alpha_{bb} > 0$. This proves Part (i).

Directional derivative without a target self-loop. Now take $a \neq b$ and differentiate once in z_a and twice in z_b. When $\alpha_{bb} = 0$, the surviving clusters have a type-a root or a type-b root. Grouping these two contributions gives

$$T_{abb}(X_k) = \alpha_{ab}\left[\alpha_{ab}\mathcal{R}_{ab}(X_k) + 2\alpha_{ba}\mathcal{S}_{ab}(X_k)\right], \tag{B.6}$$

where

$$\mathcal{R}_{ab}(X_k) = \int_0^\Delta \nu_a(t; X_k) E^0_a(\Delta - t) v_b(\Delta - t)^2\, dt, \tag{B.7}$$

$$\mathcal{S}_{ab}(X_k) = \int_0^\Delta \nu_b(t; X_k) E^0_b(\Delta - t) \int_0^{\Delta - t} e^{-\beta(\Delta - t - s)} E^0_a(s) v_b(s)\, ds\, dt. \tag{B.8}$$

Both integrals are strictly positive. Thus every surviving term contains the tested coefficient α_{ab}, exactly as in the single-edge calculation of the Motivating Example. Consequently, $T_{abb}(X_k) = 0$ if and only if $\alpha_{ab} = 0$, including when the reverse edge $b \to a$ is present.

Directional derivative with a target self-loop. If $\alpha_{bb} > 0$, the second z_b derivative also acts on descendants created by the target self-loop. The same direct differentiation separates the desired edge contribution from the additional self-loop contribution:

$$T_{abb}(X_k) = \alpha_{ab}\mathcal{Q}_{ab}(X_k) + 2\alpha_{ba}\alpha_{bb}\mathcal{C}_{ab}(X_k). \tag{B.9}$$

The positive terms used in the main theorem are

$$\mathcal{Q}_{ab}(X_k) = \int_0^{\Delta} \nu_a(t; X_k)E_a^0(\Delta - t)\left[\alpha_{ab}v_b(\Delta - t)^2 + 2\alpha_{bb}\int_0^{\Delta-t} e^{-\beta(\Delta-t-s)}E_b^0(s)v_b(s)\,ds\right]dt$$
$$+ 2\alpha_{ba}\int_0^{\Delta} \nu_b(t; X_k)E_b^0(\Delta - t)\int_0^{\Delta-t} e^{-\beta(\Delta-t-s)}E_a^0(s)v_b(s)\,ds\,dt, \tag{B.10}$$

$$\mathcal{C}_{ab}(X_k) = \int_0^{\Delta} \nu_b(t; X_k)E_b^0(\Delta - t)\left[\int_0^{\Delta-t} e^{-\beta(\Delta-t-s)}E_b^0(s)v_a(s)\,ds + v_a(\Delta - t)v_b(\Delta - t)\right]dt. \tag{B.11}$$

All integrands are nonnegative and positive on an interval, so $\mathcal{Q}_{ab}(X_k), \mathcal{C}_{ab}(X_k) > 0$. Subtracting the second term in Eq. (B.9) leaves

$$T_{abb}^{\mathrm{corr}}(X_k) = \alpha_{ab}\mathcal{Q}_{ab}(X_k),$$

which has the required zero–positive behavior. This proves Part (iv).

Observable positive-weight scores. For completeness, we connect the origin derivatives to observable low-count events. Let

$$P_0(X_k) = \Pr(\boldsymbol{Y}_k = \mathbf{0} \mid X_k) > 0, \qquad r_i(X_k) = \frac{\Pr(\boldsymbol{Y}_k = \boldsymbol{e}_i \mid X_k)}{P_0(X_k)} = \partial_{z_i}\ell_k(\boldsymbol{z} \mid X_k)|_{\boldsymbol{z}=\mathbf{0}}\,. \tag{B.12}$$

Extracting the coefficients of the conditional PGFL and applying the derivative rule for its logarithm gives

$$\mathcal{S}_{bb} = 2\mathbf{1}_{\{\boldsymbol{Y}_k = 2\boldsymbol{e}_b\}} - r_b(X_k)^2\mathbf{1}_{\{\boldsymbol{Y}_k = \mathbf{0}\}},$$
$$\mathbb{E}[\mathcal{S}_{bb} \mid X_k] = P_0(X_k)D_b(X_k). \tag{B.13}$$

Similarly,

$$\mathcal{S}_{abb} = 2\mathbf{1}_{\{\boldsymbol{Y}_k = \boldsymbol{e}_a + 2\boldsymbol{e}_b\}} - 2r_b(X_k)\mathbf{1}_{\{\boldsymbol{Y}_k = \boldsymbol{e}_a + \boldsymbol{e}_b\}}$$
$$+ r_a(X_k)r_b(X_k)^2\mathbf{1}_{\{\boldsymbol{Y}_k = \mathbf{0}\}},$$
$$\mathbb{E}[\mathcal{S}_{abb} \mid X_k] = P_0(X_k)T_{abb}(X_k), \qquad \alpha_{bb} = 0. \tag{B.14}$$

For the second identity, the third derivative of the log-PGFL initially contains the additional term $-2r_a(X_k)\Pr(\boldsymbol{Y}_k = 2\boldsymbol{e}_b \mid X_k)$. When $\alpha_{bb} = 0$, the self-loop identity gives $2\Pr(\boldsymbol{Y}_k = 2\boldsymbol{e}_b \mid X_k) = r_b(X_k)^2P_0(X_k)$, which produces the displayed score.

The algorithm replaces the unobserved X_k by the count-based $\widetilde{X}$ in Eq. (32). This is a finite-sample approximation: the identities above are exact only when conditioning on the true entrance state. They are also different from derivatives of the marginal log-PGFL; in general,

$$\partial_{z_a}\partial_{z_b}^2\ell_k(\boldsymbol{z})\big|_{\boldsymbol{z}=\mathbf{0}} \neq \mathbb{E}[T_{abb}(X_k)], \tag{B.15}$$

because marginalization introduces additional cumulants of X_k.

B.5 Proof of Theorem 4

The proof follows the three claims in the main theorem. First we control the random entrance state of a finite-origin window. We then derive the negative-axis behavior of L_b and $H_{a|b}$. Finally, analytic continuation transfers equality of the observed count laws to these witnesses.

Fix the window origin t_{k-1} and retain the main-text notation $X_k = X(t_{k-1})$, ℓ_k, L_b, $H_{a|b}$, and $\boldsymbol{U}^{(k)}$. Set

$$\tau(s) = \frac{1 - e^{-\beta s}}{\beta}, \qquad E_i^0(s) = \exp\left\{-\tau(s)\sum_j \alpha_{ij}\right\}. \tag{B.16}$$

At $\boldsymbol{z} = q\boldsymbol{e}_b$, $q < 0$, the one-window ODE has the form

$$U_i^{(k)}(s; q\boldsymbol{e}_b) = -\tau(s) + \mathbf{1}_{\{i=b\}}B_q(s), \qquad B_q' = -\beta B_q + qE_b^0 e^{\alpha_{bb}B_q}, \quad B_q(0) = 0. \tag{B.17}$$

The scalar solution satisfies $B_q \leq 0$. The random-history factor in the marginal PGFL is

$$\exp\left\{-\tau(\Delta)\sum_i\Big(\sum_j \alpha_{ij}\Big)(X_k)_i + B_q(\Delta)(A^\top X_k)_b\right\}.$$

For the derivative $H_{a|b}$, history enters only through

$$c_j(q) = \frac{\mathbb{E}\Big[(A^\top X_k)_j e^{-\tau(\Delta)\sum_i(\sum_l \alpha_{il})(X_k)_i + B_q(\Delta)(A^\top X_k)_b}\Big]}{\mathbb{E}\big[e^{-\tau(\Delta)\sum_i(\sum_l \alpha_{il})(X_k)_i + B_q(\Delta)(A^\top X_k)_b}\big]}. \tag{B.18}$$

Lemma 1 (Finite-origin history control)**.** *The entrance state has a finite first moment and a sufficiently small positive exponential moment, and* $\Pr(X_k = \mathbf{0}) = e^{-t_{k-1}\sum_i \mu_i} > 0$*. If b has a parent, define*

$$\delta_b = e^{-\beta t_{k-1}} \min_{i:\alpha_{ib}>0} \alpha_{ib} > 0 \tag{B.19}$$

and

$$c_a^\infty = \frac{\mathbb{E}\big[(A^\top X_k)_a e^{-\tau(\Delta)\sum_i(\sum_l \alpha_{il})(X_k)_i}\mathbf{1}_{\{(A^\top X_k)_b=0\}}\big]}{\mathbb{E}\big[e^{-\tau(\Delta)\sum_i(\sum_l \alpha_{il})(X_k)_i}\mathbf{1}_{\{(A^\top X_k)_b=0\}}\big]}. \tag{B.20}$$

Then, for a constant C independent of q,

$$|c_a(q) - c_a^\infty| + |c_b(q)| \le C e^{\delta_b B_q(\Delta)}. \tag{B.21}$$

If b has no parents, $(A^\top X_k)_b = 0$, $c_b = 0$, and c_a is independent of q.

Proof. Put $\mu_\Sigma = \sum_i \mu_i$ and $\bar\alpha = \max_i \sum_j \alpha_{ij}$. The total intensity is at most $\mu_\Sigma + \bar\alpha \sum_i N_i(t-)$, so a linear birth–immigration process with these rates dominates the total count. When $\bar\alpha > 0$, its finite-time PGFL is

$$\left\{\frac{e^{-\bar\alpha t}}{1-(1-e^{-\bar\alpha t})z}\right\}^{\mu_\Sigma/\bar\alpha},$$

which is finite for some $z > 1$; when $\bar\alpha = 0$, the dominating count is Poisson. Since $\|X_k\|_1 \le \sum_i N_i(t_{k-1})$, this proves $\mathbb{E}e^{\eta\|X_k\|_1} < \infty$ for sufficiently small $\eta > 0$ and hence $\mathbb{E}\|X_k\|_1 < \infty$. No event before t_{k-1} occurs with probability $e^{-t_{k-1}\sum_i \mu_i}$, which gives the asserted atom at zero.

If $(A^\top X_k)_b > 0$, some parent type i occurred before the window. Its contribution is at least $e^{-\beta t_{k-1}}$, and hence $(A^\top X_k)_b \ge \delta_b$. Split the expectations in Eq. (B.18) according to whether $(A^\top X_k)_b$ is zero. Since $B_q(\Delta) \le 0$, the nonzero part is bounded by $e^{\delta_b B_q(\Delta)}$, while the atom $X_k = \mathbf{0}$ keeps each denominator bounded away from zero. Comparing the resulting ratios proves Eq. (B.21). The no-parent case is immediate. □

Lemma 2 (Negative-axis structural separation)**.** *For every b,*

$$L_b(q) = \begin{cases} -\mu_b|q| \displaystyle\int_0^\Delta E_b^0(s)\,ds + O(1), & \alpha_{bb} = 0, \\ O(\log|q|), & \alpha_{bb} > 0. \end{cases} \tag{B.22}$$

For $a \ne b$ and $\alpha_{bb} = 0$,

$$H_{a|b}(q) \longrightarrow \begin{cases} C_+ > 0, & \alpha_{ab} = 0,\ \alpha_{ba} = 0, \\ -\infty \text{ linearly in } |q|, & \alpha_{ab} = 0,\ \alpha_{ba} > 0, \\ 0, & \alpha_{ab} > 0,\ \alpha_{ba} = 0, \\ -C_- < 0, & \alpha_{ab} > 0,\ \alpha_{ba} > 0, \end{cases} \tag{B.23}$$

where $C_+, C_- > 0$. If $\alpha_{bb} > 0$, then

$$\begin{cases} \log|q|\,|H_{a|b}(q)| \to 0, & \alpha_{ab} > 0, \\ \liminf_{q\to-\infty} \log|q|\,|H_{a|b}(q)| > 0, & \alpha_{ab} = 0. \end{cases} \tag{B.24}$$

These conclusions include the random entrance state and the cancellation case treated below.

Proof. Step 1: a common derivative identity. At $q\mathbf{e}_b$, define only the two nonzero transverse sensitivities

$$P_q = \partial_{z_a} U_a^{(k)}, \qquad Q_q = \partial_{z_a} U_b^{(k)}.$$

Directly differentiating the ODE gives

$$\begin{aligned} P_q' &= -\beta P_q + f_q, & f_q &= E_a^0 e^{\alpha_{ab} B_q}, \\ Q_q' &= -\beta Q_q + h_q(\alpha_{ba} P_q + \alpha_{bb} Q_q), & h_q &= q E_b^0 e^{\alpha_{bb} B_q}, \end{aligned}$$

with zero initial values. Differentiating the marginal-window formula then gives the exact identity

$$H_{a|b}(q) = \mu_a \int_0^\Delta f_q(s)\,ds + \mu_b \int_0^\Delta h_q(s)[\alpha_{ba} P_q(s) + \alpha_{bb} Q_q(s)]\,ds + c_a(q) P_q(\Delta) + c_b(q) Q_q(\Delta). \tag{B.25}$$

Step 2: target without a self-loop. If $\alpha_{bb} = 0$, Eq. (B.17) gives $B_q = qv_b$ with $v_b(s) = \int_0^s e^{-\beta(s-u)} E_b^0(u)\, du > 0$, and $h_q = qE_b^0$. Consequently,

$$\ell_k(q\boldsymbol{e}_b \mid \mathbf{0}) = -\Delta \sum_i \mu_i + \mu_b q \int_0^\Delta E_b^0(s)\, ds.$$

The marginal history term lies between $\log \Pr(X_k = \mathbf{0})$ and zero, which proves the first case of Eq. (B.22).

If $\alpha_{ab} = 0$, both $f_q = E_a^0$ and P_q are positive and independent of q. When $\alpha_{ba} = 0$, $Q_q = 0$ and Eq. (B.25) tends to a positive constant. When $\alpha_{ba} > 0$, its second integral is strictly negative and linear in $|q|$. Lemma 1 makes the history term lower order, proving the first two cases of Eq. (B.23).

Now let $\alpha_{ab} > 0$. Since $v_b(s) = s + O(s^2)$ at zero, the endpoint Laplace estimate for the variation-of-constants formula gives

$$|q|P_q(s) \to \frac{e^{-\beta s}}{\alpha_{ab}}, \qquad q \int_0^\Delta E_b^0(s) P_q(s)\, ds \to -\frac{1}{\alpha_{ab}} \int_0^\Delta E_b^0(s) e^{-\beta s}\, ds. \tag{B.26}$$

The same representation gives a uniform bound on $|q|P_q$, so dominated convergence is applicable. Hence the first and history terms in Eq. (B.25) vanish. If $\alpha_{ba} = 0$, $Q_q = 0$ and $H_{a|b} \to 0$. If $\alpha_{ba} > 0$, the displayed limit and variation of constants give

$$Q_q(\Delta) \to -\frac{\alpha_{ba} e^{-\beta\Delta}}{\alpha_{ab}} \int_0^\Delta E_b^0(s)\, ds < 0.$$

Lemma 1 implies $c_a(q)P_q(\Delta) \to 0$ and $c_b(q)Q_q(\Delta) \to 0$, so $H_{a|b}$ has a finite strictly negative limit. This proves the last two cases of Eq. (B.23).

Step 3: the scalar scale created by a target self-loop. Suppose $\alpha_{bb} > 0$ and set $\omega_q = \alpha_{bb}|q|$. Define

$$W(s) = \exp\left\{ -\alpha_{bb} B_q(s) + \Big(\sum_j \alpha_{bj}\Big)\tau(s) \right\}.$$

Equation (B.17) becomes

$$W' = \omega_q + \Big(\sum_j \alpha_{bj}\Big) W - \beta W \log W, \qquad W(0) = 1. \tag{B.27}$$

Its unique equilibrium satisfies

$$\Lambda_q = W_0\left(\frac{\omega_q}{\beta} e^{-(\sum_j \alpha_{bj})/\beta} \right), \qquad W_q^\star = \frac{\omega_q}{\beta \Lambda_q}, \qquad \Lambda_q \sim \log|q|, \tag{B.28}$$

where W_0 is the positive Lambert branch. Scalar comparison gives $1 \le W(s) < W_q^\star$. With $\zeta = W/W_q^\star$ and $t = \beta\Lambda_q s$,

$$\frac{d\zeta}{dt} = 1 - \zeta - \frac{\zeta \log \zeta}{\Lambda_q} \ge 1 - \zeta.$$

Thus $W/W_q^\star \to 1$ uniformly on every $[\varepsilon, \Delta]$, and

$$-B_q(\Delta) = \Theta(\Lambda_q) = \Theta(\log|q|). \tag{B.29}$$

Lemma 1 consequently gives, for some $\chi > 0$,

$$|c_a(q) - c_a^\infty| + |c_b(q)| = O(|q|^{-\chi}) = o((\log|q|)^{-1}). \tag{B.30}$$

We also need one integral consequence of the same comparison. For every $\vartheta > 0$, splitting the scaled time interval at $t = 1$ and using $\zeta(t) \ge c[(W_q^\star)^{-1} + t]$ before that point and $\zeta(t) \ge 1 - e^{-1}$ afterward gives

$$\Lambda_q \int_0^\Delta W(s)^{-\vartheta}\, ds \le C \begin{cases} \Lambda_q (W_q^\star)^{-\vartheta}, & 0 < \vartheta < 1, \\ (\log W_q^\star + \Lambda_q)/W_q^\star, & \vartheta = 1, \\ (W_q^\star)^{-1} + \Lambda_q (W_q^\star)^{-\vartheta}, & \vartheta > 1. \end{cases}$$

Every term on the right tends to zero by Eq. (B.28); hence

$$\Lambda_q \int_0^\Delta W(s)^{-\vartheta}\, ds \longrightarrow 0. \tag{B.31}$$

Because $h_q = -|q|/W$, integration of the B_q equation now shows that the empty-entrance part of L_b is $O(\log|q|)$. Its history term remains between $\log \Pr(X_k = \mathbf{0})$ and zero, proving the second case of Eq. (B.22).

Step 4: cross-edge separation when the target has a self-loop. Put $S_q = \alpha_{ba} P_q + \alpha_{bb} Q_q$. Then

$$S_q' = (-\beta + \alpha_{bb} h_q) S_q + \alpha_{ba} f_q, \qquad S_q(0) = 0,$$

so $S_q \geq 0$. Integrating this equation in Eq. (B.25) gives

$$\begin{aligned} H_{a|b} = &\left(\mu_a - \frac{\mu_b \alpha_{ba}}{\alpha_{bb}}\right) \int_0^\Delta f_q(s)\, ds + \left(c_a - \frac{\alpha_{ba}}{\alpha_{bb}} c_b\right) P_q(\Delta) \\ &+ \frac{\mu_b + c_b}{\alpha_{bb}} S_q(\Delta) + \frac{\beta \mu_b}{\alpha_{bb}} \int_0^\Delta S_q(s)\, ds. \end{aligned} \tag{B.32}$$

If $\alpha_{ab} > 0$, let $\vartheta = \alpha_{ab}/\alpha_{bb}$. Then $f_q \leq C W^{-\vartheta}$, so Eq. (B.31) implies $\Lambda_q \int f_q \to 0$. Variation of constants gives $0 \leq P_q(\Delta) \leq \int f_q$, $0 \leq S_q(\Delta) \leq \alpha_{ba} \int f_q$, and $0 \leq \beta \int S_q \leq \alpha_{ba} \int f_q$. Equation (B.32) and boundedness of the c_j therefore imply $\Lambda_q |H_{a|b}| \to 0$.

If $\alpha_{ab} = 0$, $f_q = E_a^0$ and $P_q(\Delta)$ is a fixed positive constant. Moreover,

$$0 \leq \sup_{[0,\Delta]} S_q \leq \frac{\alpha_{ba} \max E_a^0}{\beta(1 + \Lambda_q)}. \tag{B.33}$$

Using Eq. (B.30) in Eq. (B.32) gives

$$H_{a|b}(q) = C_0 + \frac{\mu_b}{\alpha_{bb}} \left[S_q(\Delta) + \beta \int_0^\Delta S_q(s)\, ds \right] + o(\Lambda_q^{-1}), \tag{B.34}$$

where

$$C_0 = \left(\mu_a - \frac{\mu_b \alpha_{ba}}{\alpha_{bb}}\right) \int_0^\Delta E_a^0(s)\, ds + c_a^\infty \int_0^\Delta e^{-\beta(\Delta - s)} E_a^0(s)\, ds.$$

If $C_0 \neq 0$, the missing-edge conclusion is immediate. If $C_0 = 0$, then $\alpha_{ba} > 0$; otherwise strict positivity of μ_a would give $C_0 > 0$. To obtain a lower bound in this cancellation case, take $I = [\Delta/2, 3\Delta/4]$ and $m_a = \min_{[\Delta/4, 3\Delta/4]} E_a^0 > 0$. Interior convergence of $W/W_q^\star$ implies $\beta + \omega_q/W \leq 3\beta\Lambda_q$ on $[\Delta/4, 3\Delta/4]$ for large $|q|$. Restricting the variation-of-constants formula for $S_q(s)$ to $[s - (3\beta\Lambda_q)^{-1}, s]$ gives

$$S_q(s) \geq \frac{\alpha_{ba} m_a (1 - e^{-1})}{3\beta\Lambda_q}, \qquad s \in I.$$

Together with Eq. (B.33), this gives

$$\int_0^\Delta S_q(s)\, ds = \Theta(\Lambda_q^{-1}). \tag{B.35}$$

The remainder in Eq. (B.34) is $o(\Lambda_q^{-1})$ and cannot cancel this positive term. Thus $\liminf \Lambda_q |H_{a|b}| > 0$. Since $\Lambda_q \sim \log |q|$, both cases of Eq. (B.24) follow. □

Lemma 3 (Analytic transfer from the observed window law)**.** *For every b and $a \neq b$, the ODE/Laplace-defined L_b and $H_{a|b}$ are real analytic on $(-\infty, 0)$. Equality of two complete window PGFLs on $[0,1]^d$ implies equality of these witnesses for every $q < 0$.*

Proof. Lemma 1 gives $\mathbb{E} e^{\eta \|X_k\|_1} < \infty$ for some $\eta > 0$. Hence the Laplace transform of X_k is analytic in a complex neighborhood of every real nonpositive argument. Its value is positive on the real axis, so a local analytic logarithm exists there. The ODE in Eq. (B.17) has a bounded finite-window solution near every fixed $q_0 < 0$, and standard analytic dependence of ODE solutions applies to q and the transverse coordinate z_a. Composing the ODE terminal state with the preceding Laplace logarithm proves local analyticity of L_b and $H_{a|b}$ throughout $(-\infty, 0)$.

Near $\boldsymbol{z} = \mathbf{0}$, the marginal PGFL is strictly positive; explicitly,

$$e^{\ell_k(\mathbf{0})} = e^{-\Delta \sum_i \mu_i} \mathbb{E}\left[e^{-\tau(\Delta) \sum_i (\sum_j \alpha_{ij})(X_k)_i} \right] > 0. \tag{B.36}$$

Equality of the complete labelled count laws therefore gives the same normalized local log-PGFL and the same transverse derivatives near the origin. The two models' witnesses agree on an interval $(-\varepsilon, 0)$. The real-analytic identity theorem extends this equality to all $q < 0$. This continues only the model-defined ODE/Laplace witnesses; it does not claim that the original PGFL power series converges for every negative q. □

Proof of Theorem 4. Suppose the two models have the same labelled window law. By Lemma 3, their L_b and $H_{a|b}$ coincide on the entire negative axis.

For each b, Eq. (B.22) distinguishes linear negative growth from logarithmic growth, and hence identifies whether α_{bb} is zero. If this common target-loop status is zero, the first two cases of Eq. (B.23) are exactly $\alpha_{ab} = 0$, while the last two are exactly $\alpha_{ab} > 0$. If the target has a self-loop, Eq. (B.24) distinguishes the same two support states by their rates of decay. Applying these classifications to every diagonal entry and ordered pair proves $\mathrm{supp}(A) = \mathrm{supp}(\widetilde{A})$.

The two models may have different positive decay rates. Each decay rate changes constants in its own auxiliary ODE but not the separating $|q|$ and $\log|q|$ scales used above. □

B.6 Derivation of the Discrete-Time Likelihood Criterion

The exact within-bin event positions are unavailable. The working criterion therefore averages the exponential kernel over independent uniform positions inside each bin. Put

$$x_\Delta = \beta\Delta, \qquad a(x_\Delta) = \frac{(e^{x_\Delta}-1)(1-e^{-x_\Delta})}{x_\Delta^2}, \qquad c_0(x_\Delta) = \frac{1}{x_\Delta}\left(1 - \frac{1-e^{-x_\Delta}}{x_\Delta}\right).$$

For a source in bin I_{k-m} and a target in I_k, $m \geq 1$, let $u, w \sim \mathrm{Unif}(0,\Delta)$ be their within-bin positions. Independence gives

$$\mathbb{E}[e^{-\beta(m\Delta+w-u)}] = e^{-\beta m\Delta}\mathbb{E}[e^{\beta u}]\mathbb{E}[e^{-\beta w}] = a(x_\Delta)e^{-\beta m\Delta}.$$

For two events in the same bin, temporal order must also be averaged:

$$\mathbb{E}[e^{-\beta(w-u)}\mathbf{1}_{\{u<w\}}] = \frac{1}{\Delta^2}\int_0^\Delta \int_0^w e^{-\beta(w-u)}\,du\,dw = c_0(x_\Delta).$$

The target itself is excluded when source and target types coincide. Hence

$$s^{(v)}_{v',k} = a(x_\Delta)\sum_{\ell<k} e^{-\beta(t_{k-1}-t_{\ell-1})}Y_{v',\ell} + c_0(x_\Delta)(Y_{v',k} - \mathbf{1}_{\{v'=v\}}), \tag{B.37}$$

and $\bar{\lambda}_{v,k} = \mu_v + \sum_{v'} \alpha_{v'v}s^{(v)}_{v',k}$.

For the compensator, a source at $t_{k-1}+u$ contributes expected kernel mass

$$\frac{1}{\beta}\left[1 - e^{-\beta(T-t_{k-1})}\frac{e^{x_\Delta}-1}{x_\Delta}\right].$$

Combining this mass with $\sum_{v,k\in\mathcal{I}_v} Y_{v,k}\log\bar{\lambda}_{v,k}$ and the baseline compensator $T\sum_v \mu_v$ gives $\mathcal{L}$ in Eq. (29).

This criterion also uses $\mathbb{E}_{\mathrm{U}}[\log\lambda] \approx \log\mathbb{E}_{\mathrm{U}}[\lambda]$ at bins containing target events. It is therefore a discrete-time working criterion, not the exact likelihood of the binned counts and not a likelihood bound under the true conditional law of event positions. As $x_\Delta \to 0$, $a(x_\Delta) \to 1$ and $c_0(x_\Delta) \to 1/2$, as expected.

C Algorithmic Details

This appendix follows the three structural-learning layers in Section 4; the final fixed-graph refit updates parameters only. The bin width Δ, shared decay β, and calibration parameters are fixed method settings and are omitted from the compact input line of Algorithm 1. The symbols below denote tuning parameters rather than reported experimental values. Every gate is a finite-sample decision rule, whereas Theorem 4 concerns the complete population count law. The implementation stores $A_{uv} = \alpha_{uv}$ for parent u and child v.

C.1 Discrete-Time Regressors, MM Fitting, and the Backbone

Count construction. Because the input retains counts but not positions within a bin, the discrete-time criterion uses the kernel in Eq. (28). Each recorded time t_e is first replaced by its integer bin $b_e = \lfloor t_e/\Delta \rfloor$. Let $b_0 = \min_e b_e$ and $b_1 = \max_e b_e$. All records are placed on the common grid $k_e = b_e - b_0 + 1$, with implicit zero bins and common horizon $T = (b_1 - b_0 + 1)\Delta$. The timestamp is discarded after this mapping, before any fit or structural statistic is computed. The range from the earliest to latest occupied bin is a working observation range; it does not reveal unrecorded zero-count periods outside that range. Using the same T for every record assumes that all records share this observation span. If their true horizons differ, this convention may change exposure terms and is an additional finite-sample approximation, not a consequence of the empty-history identifiability theorem. Independently shifted records must first be aligned to a meaningful common discrete clock. Put $x_\Delta = \beta\Delta$ and use $a(x_\Delta)$ and $c_0(x_\Delta)$ from Eq. (28). If t_{rk} is bin k's left endpoint and T_r is record r's observation horizon, define

$$\begin{aligned} R_{u,rk} &= \sum_{\ell<k} e^{-\beta(t_{rk}-t_{r\ell})}Y_{u,r\ell}, \\ s^{(v)}_{u,rk} &= a(x_\Delta)R_{u,rk} + c_0(x_\Delta)\left(Y_{u,rk} - \mathbf{1}_{\{u=v\}}\right), \\ c_u &= \sum_{r,k}\frac{Y_{u,rk}}{\beta}\left[1 - e^{-\beta(T_r-t_{rk})}\frac{e^{x_\Delta}-1}{x_\Delta}\right]. \end{aligned} \tag{C.1}$$

The regressor is used only in bins containing target events, so the self-count correction excludes the selected event. In the current common-grid implementation $T_r = T$ for every record and $T_\Sigma = \sum_r T_r$. These are the quantities used in Eq. (29); Appendix B.6 derives their within-bin averaging factors.

Fixed-graph fitting. At the current parameters, responsibilities are

$$p_{0,v,rk} = \frac{\mu_v}{\bar{\lambda}_{v,rk}}, \qquad p_{u,v,rk} = \frac{\alpha_{uv} s^{(v)}_{u,rk}}{\bar{\lambda}_{v,rk}}. \tag{C.2}$$

Jensen's inequality gives the responsibility-weighted lower bound. Maximizing it with the linear compensator yields Eq. (30). Sources with zero exposure are excluded from division, and coefficients absent from the candidate graph remain zero. Fits are cached by child and parent set, so only changed columns are recomputed.

Backbone search. The Layer 1 graph has zero free off-diagonal entries at initialization: $A^{(0)}_{ij} = 0$ for $i \neq j$. Every diagonal is temporarily active as an auxiliary self-loop; this is the initial graph used in Algorithm 1. For a sampled window, write $Y = (Y_v)_{v \in V}$ for its count vector. Write $\bar{Y}_i := \widehat{\mathbb{E}}[Y_i]$ for the empirical mean count of type i. The empirical second factorial cumulant is $\widehat{K}_{2,ij} = \widehat{\mathbb{E}}[Y_i Y_j - \mathbf{1}_{\{i=j\}} Y_i] - \bar{Y}_i \bar{Y}_j$. These moments pool sampled windows rather than condition on a fixed entrance state as in Theorem 3. The same notation is used in the empirical K_3 formulas below. A bootstrap-normalized absolute score ranks pairs. A burn-in tied to the decay time is clipped to the available horizon. The small-dimension branch exposes every unordered pair; otherwise the union of per-node top-ranked pairs is expanded in batches. Exposure does not alter the graph: every newly exposed pair remains absent until its absent, forward, reverse, and reciprocal states have been refitted and compared. Infeasible or unstable candidates are rejected, and numerical ties preserve the current state. Exposure stops at its budget or after the configured number of batches without accepted new pairs. Unexposed pairs need not be evaluated, so the large-graph procedure is not exhaustive support search.

C.2 Origin Sampling and Diagonal Selection

Entrance-state approximation. The count-based entrance-state approximation $\widetilde{X}$ is defined in Eq. (32). Unrolling that recurrence gives $\widetilde{X}_{u,r,k} = \frac{1-e^{-\beta\Delta}}{\beta\Delta} \sum_{j=1}^{k-1} e^{-\beta\Delta(k-1-j)} Y_{u,rj}$ for each record r. The uniform within-bin position is a working approximation, so this observable approximation is generally not the exact entrance state $X(t_{k-1})$. Sampling retains nonzero windows and nearby zero windows, and subsamples distant zero windows with inverse-inclusion weights. For the population identity, use the generic window count vector Y defined above and set $I_y = \mathbf{1}_{\{Y=y\}}$, $P_0(X_k) = \Pr(\boldsymbol{Y}_k = \mathbf{0} \mid X_k)$, and $r_b(X_k) = \Pr(\boldsymbol{Y}_k = \boldsymbol{e}_b \mid X_k)/P_0(X_k)$. The ideal conditional score

$$\mathcal{S}_{bb} = 2I_{2\boldsymbol{e}_b} - r_b(X_k)^2 I_{\mathbf{0}} \tag{C.3}$$

satisfies, for the population r_b,

$$\mathbb{E}[\mathcal{S}_{bb} \mid X_k] = P_0(X_k) D_b(X_k), \qquad P_0(X_k) > 0. \tag{C.4}$$

The implementation substitutes $\widetilde{X}$ for the unobserved X_k, fits the singleton ratios by nonnegative ridge regression on held-out records, and averages the approximate score with the sampling weights. Thus Eq. (C.4) motivates the estimator but is not asserted after replacing X_k by $\widetilde{X}$. Record-level bootstrap refits the ratio regressions; $\widehat{D}^{\mathrm{low}}_b$ denotes the resulting lower confidence bound, not a fixed-state derivative estimate.

Diagonal decision. With prescribed thresholds $0 \leq \tau_{\mathrm{weak}} \leq \tau_{\mathrm{strong}}$, the free-diagonal rule is

$$\mathrm{keep}(b \to b) = \mathbf{1}\{(\widehat{D}^{\mathrm{low}}_b > 0 \ \wedge \ \widehat{\alpha}_{bb}/\beta \geq \tau_{\mathrm{weak}}) \ \vee \ \widehat{\alpha}_{bb}/\beta \geq \tau_{\mathrm{strong}}\}. \tag{C.5}$$

The chosen diagonal is fixed before subsequent fitting. There is no diagonal-K_3 acceptance gate in this version. A disabled self-loop branch is an estimator configuration, not a theoretical restriction.

C.3 Target-Specific Reciprocal Review

A target self-loop changes the interpretation of the local directional derivative, so exposed reciprocal pairs touching a retained loop receive a separate bounded review before the ordinary finite-q search. Each one-way alternative is refitted and must satisfy the stability screen and a prescribed BIC-loss allowance relative to the reciprocal fit. The evidence is computed separately for each ordered target.

For a looped target b, local positive-q regression estimates a conditional affine approximation at selected $q \in (0,1)$. Put $\psi(\widetilde{X}) = (1, \widetilde{X}^\top)^\top$, and let w denote the inverse-inclusion weight from origin sampling. Here Y_{-b} is the count vector with coordinate b removed, and $Y_{-\{a,b\}}$ removes both indicated coordinates. The regression uses the weighted matrix $\sum w \mathbf{1}_{\{Y_{-b}=0\}} q^{Y_b} \psi\psi^\top$ and right-hand side $\sum w \mathbf{1}_{\{Y_a=1, Y_{-\{a,b\}}=0\}} q^{Y_b} \psi$, with ridge regularization. Its intercept is compared with the model's empty-entrance witness $H^{(0)}_{a|b}(q) := \partial_{z_a} \ell_k(\boldsymbol{z} \mid \mathbf{0})|_{\boldsymbol{z}=q\boldsymbol{e}_b}$. Residuals across the evaluation points are scaled by their bootstrap covariance to define the local loss.

For a loop-free target, the empirical score is

$$\mathcal{S}_{abb} = 2I_{\boldsymbol{e}_a+2\boldsymbol{e}_b} - 2r_b(X_k)I_{\boldsymbol{e}_a+\boldsymbol{e}_b} + r_a(X_k)r_b(X_k)^2 I_{\mathbf{0}}. \tag{C.6}$$

For population ratio functions based on the true entrance state and $\alpha_{bb} = 0$, its conditional mean is $P_0(X_k)T_{abb}(X_k)$. The implementation instead fits these functions on $\widetilde{X}$ and compares the approximate score's weighted mean with an empty-entrance model derivative. For a prescribed finite-difference step $\delta_{\mathrm{FD}} > 0$, the latter is evaluated by

$$T^{(0)}_{abb} \simeq \frac{2H^{(0)}_{a|b}(0) - 5H^{(0)}_{a|b}(\delta_{\mathrm{FD}}) + 4H^{(0)}_{a|b}(2\delta_{\mathrm{FD}}) - H^{(0)}_{a|b}(3\delta_{\mathrm{FD}})}{\delta^2_{\mathrm{FD}}}. \tag{C.7}$$

These quantities generally have different estimands: the empirical score uses the count-based entrance-state approximation, whereas the model quantity fixes an empty true entrance. Accordingly, this comparison is a theory-guided finite-sample heuristic, not an unbiased moment equation or a new consequence of the conditional-T theorem. It does not estimate the self-loop correction term in Part (iv) of Theorem 3.

Available finite discrepancies are summed across the two ordered directions. A one-way candidate is accepted only if the local loss gain exceeds its margin and the BIC allowance passes; ties are resolved using the score gap. Bounded sweeps permit further reviews. No new adjacency or separate conditional-T completion step is performed.

C.4 Finite Negative-Axis Evidence and Reliability

The complete-law theorem separates edge configurations along the negative axis. The finite-sample procedure evaluates one prescribed coordinate $q = q_H < 0$ and applies reliability screening. For each ordered pair, pooled complete windows give

$$\widehat{H}_{a|b}(q) = \frac{\sum \mathbf{1}_{\{Y_a=1, Y_{-\{a,b\}}=0\}} q^{Y_b}}{\sum \mathbf{1}_{\{Y_{-b}=0\}} q^{Y_b}}. \tag{C.8}$$

Zero windows remain in the denominator. For numerical stability, positive and negative terms are accumulated separately on the log scale. A zero denominator makes the ratio undefined; a zero numerator with a nonzero denominator gives a valid zero.

Whole records are resampled. Let B_{boot} be the number of bootstrap replicates and B^{valid}_{ab} the number with a defined ratio; write $\kappa_{ab} = B^{\mathrm{valid}}_{ab}/B_{\mathrm{boot}}$. Undefined ratios are excluded, never replaced with zero. Relative standard error is computed from rescaled signed-log replicate ratios, avoiding overflow. A zero pooled ratio cannot pass the relative-error test. An ordered coordinate passes when the independent-record count exceeds $n_{\min}$, the pooled denominator is positive, $\kappa_{ab} \geq \pi_{\min}$, and the relative error is below τ_H. An unordered pair enters the search only when both directions pass and neither endpoint has a retained loop.

Theoretical $H^{\mathrm{th}}_{G,a|b}(q_H; t_H)$ uses the transient PGFL calculation at the prescribed negative coordinate q_H and entrance time t_H following the empty-history origin. Empirical windows are pooled over origins; this fixed-time model comparison is a working regularizer rather than an asserted identity with that mixture. The procedure does not evaluate $q \to -\infty$ and therefore does not implement the complete-law identifiability proof.

Let $\mathcal{V}_H$ be the set of reliable ordered coordinates. When $\mathcal{V}_H = \varnothing$, this layer abstains and retains the current graph. Otherwise, with regularization weight $\lambda_H \geq 0$ and numerical floor $\varepsilon_H > 0$, the finite-q search uses

$$S_H(G) = \frac{\mathcal{L}_G}{N_{\mathrm{ev}}} - \lambda_H L_H(G), \qquad L_H(G) = \frac{1}{|\mathcal{V}_H|} \sum_{(a,b)\in\mathcal{V}_H} \left[\frac{H^{\mathrm{th}}_{G,a|b}(q_H; t_H) - \widehat{H}_{a|b}(q_H)}{\max\{|\widehat{H}_{a|b}(q_H)|, \varepsilon_H\}} \right]^2. \tag{C.9}$$

Eligible pairs are temporarily reset to the reciprocal state before a bounded greedy toggle search. Each toggle changes one directed edge and refits the affected column. Only strict stable improvements of Eq. (C.9) are accepted. The temporary reset itself need not improve that score. The BIC weight is zero in this layer; the explicit BIC guards in other layers remain separate.

C.5 Record-Averaged Marginal K_3 Direction Proposals and Their Guard

After the support review, marginal repeated-index third factorial cumulants supply an observable direction score. Let $\mathcal{I}_r$ be the usable window set of record r after applying the common horizon, burn-in, and window selection rules, and write $\boldsymbol{Y}_{r,k} = (Y_{v,rk})_{v\in V}$ for its count vector in window k. For any function f of the window count vector, define

$$\widehat{\mathbb{E}}_r[f(Y)] = \frac{1}{|\mathcal{I}_r|} \sum_{k\in\mathcal{I}_r} f(\boldsymbol{Y}_{r,k}), \qquad \bar{Y}_{r,i} = \widehat{\mathbb{E}}_r[Y_i]. \tag{C.10}$$

For distinct i, j, the record-level repeated-index factorial cumulant is

$$\begin{aligned}\widehat{K}^{(r)}_{3,jji} = \widehat{\mathbb{E}}_r[Y_j(Y_j-1)Y_i] - \widehat{\mathbb{E}}_r[Y_j(Y_j-1)]\bar{Y}_{r,i} \\ - 2\widehat{\mathbb{E}}_r[Y_jY_i]\bar{Y}_{r,j} + 2\bar{Y}^2_{r,j}\bar{Y}_{r,i}\end{aligned} \tag{C.11}$$

and the statistic used by the algorithm is the equal-record average

$$\widehat{K}^{\mathrm{pool}}_{3,jji} = \frac{1}{R}\sum_{r=1}^{R}\widehat{K}^{(r)}_{3,jji}. \tag{C.12}$$

Thus a longer record does not receive extra weight through its number of usable windows. When record lengths differ, the implementation first restricts them to the configured common usable horizon and applies the same burn-in and window-selection rule before Eq. (C.10). Whole-record bootstrap resampling preserves dependence among windows within a record and provides regularized standard errors.

The statistic in Eq. (C.12) is a low-dimensional projection of the observable marginal count law. The implemented $\widehat{K}_3$ is the factorial cumulant of the marginal window-count distribution, computed record-wise and then averaged across records. It uses the same repeated-index cumulant coordinate as the conditional quantity in Theorem 3, but it is not, in general, the average of the fixed-state conditional cumulants. The conditional calculation motivates the repeated-index coordinates, while Theorem 4 establishes that the complete marginal law retains the directed support. For orientation, define

$$\Delta_3 = z_{ij} - z_{ji}, \qquad z_{ij} = \frac{\widehat{K}^{\mathrm{pool}}_{3,jji}}{\widehat{\mathrm{se}}(\widehat{K}^{\mathrm{pool}}_{3,jji})}, \qquad \Delta_{\mathrm{BIC}} = S_{\mathrm{BIC}}(G_{i\to j}) - S_{\mathrm{BIC}}(G_{j\to i}). \tag{C.13}$$

A positive Δ_3 proposes $i \to j$, whereas a negative value proposes $j \to i$. This proposal is not accepted on its own. Here $G_{i\to j}$ and $G_{j\to i}$ are equal-complexity refitted one-way alternatives, so their BIC ordering agrees with their fitted-$\mathcal{L}$ ordering.

For a fitted one-way candidate $G_{i\to j}$, define its normalized finite-q theoretical penalty by

$$P(i\to j) = \frac{|H^{\mathrm{th}}_{G_{i\to j},i|j}(q_H;t_H)| + \max\{H^{\mathrm{th}}_{G_{i\to j},j|i}(q_H;t_H),0\}}{\max\{|H^{\mathrm{th}}_{G_{i\to j},i|j}(q_H;t_H)|, |H^{\mathrm{th}}_{G_{i\to j},j|i}(q_H;t_H)|, \varepsilon_H\}}. \tag{C.14}$$

The resulting directional contrast is

$$\Delta_H = \frac{P(j\to i) - P(i\to j)}{P(j\to i) + P(i\to j) + \varepsilon_H}. \tag{C.15}$$

Across-pair agreement with Δ_{BIC}, the number of reliable comparisons, and their separation select the configured fusion mode. In H-primary mode, Δ_H is used above its adaptive reliability floor and otherwise the procedure falls back to Δ_3. Conservative mode uses $\Delta_3 + w_H\Delta_H$, where w_H is calibrated from sign-change breakpoints. Dead-zone ties abstain. Thus the selected proposal is either Δ_H (falling back to Δ_3 below its reliability floor) in H-primary mode, or $\Delta_3 + w_H\Delta_H$ in conservative mode. These data-adaptive choices are finite-sample calibrations rather than independent population tests.

Proposals are assessed against a common reference graph. Let $\tau_{\mathrm{eff}} > 0$ be the weak-excitation threshold and $\gamma_{\mathrm{rec}} \in \mathbb{R}$ the allowed reciprocal advantage measured in BIC-penalty units. Reciprocal protection is triggered when

$$S_{\mathrm{BIC}}(G_{i\leftrightarrow j}) - \max\{S_{\mathrm{BIC}}(G_{i\to j}), S_{\mathrm{BIC}}(G_{j\to i})\} > \gamma_{\mathrm{rec}}\frac{\omega_{\mathrm{BIC}}\log N_{\mathrm{ev}}}{2N_{\mathrm{ev}}}. \tag{C.16}$$

A direction proposal is accepted only when the finite-q H comparison is reliable and the equal-complexity refitted score does not trigger reciprocal protection. The calibrated guard additionally considers whether the weaker fitted excitation $\min(\widehat{\alpha}_{ij}, \widehat{\alpha}_{ji})/\beta$ falls below τ_{eff} or whether Eq. (C.16) is false. Pairs touching retained loops are excluded, and an unstable combined refit restores the reference graph. Only the repeated-index coordinates above are computed; no additional post-orientation moment-refinement layer is used.

C.6 Tuning, Termination, and Computational Scope

Tuning parameters comprise the candidate budget, MM tolerance and iteration budget, ω_{BIC}, window-sampling and ridge controls, bootstrap replicate count and confidence level, $\tau_{\mathrm{weak}}, \tau_{\mathrm{strong}}$, the local positive-$q$ coordinates and BIC allowance, $q_H, t_H, \lambda_H, \varepsilon_H, n_{\min}, \pi_{\min}, \tau_H$, the mode-selection margins, and $\tau_{\mathrm{eff}}, \gamma_{\mathrm{rec}}$. They control fitting, reliability, and search breadth. No experiment-specific values are specified here.

Candidate fits must satisfy the implementation's spectral-radius check $\rho(A)/\beta < 1$ with numerical tolerance. A final stable refit is retained; if it fails the screen, the preceding accepted fit is kept. Explicit budgets bound all sweeps and candidate expansions, ensuring finite execution for terminating numerical routines. Scores differ

between layers and reciprocal seeding can decrease them, so there is no claim of a single globally increasing objective or global optimality.

If C candidate fits each take at most I MM iterations and M source–target regressor entries are evaluated per iteration, the fitting work is of order CIM before caching. The repeated-index direction stage evaluates at most quadratically many ordered coordinates and does not construct a full third-order tensor. This does not make the complete algorithm quadratic: candidate-graph fitting and its MM iterations retain their own costs. Parent-set fit caches and reusable moment summaries reduce repeated computation. Bootstrap and finite-q ODE evaluations can dominate runtime. Parallel execution is an optional implementation optimization, not a required or platform-independent capability.