

# Causal Discovery from Interval-Based Event Sequences

Lénaïg Cornanguer<sup>1</sup>, Joscha Cüppers<sup>1</sup>, Jilles Vreeken<sup>1</sup>

<sup>1</sup>CISPA Helmholtz Center for Information Security, Saarbrücken, Germany  
{lenaig.cornanguer,joscha.cueppers,jv}@cispa.de

## Abstract

In this paper we address the problem of discovering causal relationships from observational event sequence data. Existing methods typically assume that events are instantaneous point events, however in many real-world settings, events have duration. For example, in healthcare, a patient’s symptoms may persist over a time interval and influence clinical actions while ongoing. To address this, we introduce a causal model for interval-based event sequences that captures rich causal structures, including interactions between events and causal mechanisms that depend on whether other events are ongoing. We prove that our model is identifiable in the limit and present a practical causal discovery algorithm, NIAGARA, grounded in the algorithmic Markov condition. To select among candidate models, we employ a minimum description length (MDL) criterion, enabling robust inference even with limited data. We validate our approach on synthetic and real data and demonstrate its utility on a real-world medical case study, where it uncovers meaningful causal relationships from noisy, interval-based event data.

**Code** — <http://eda.rg.cispa.io/niagara/>

**Dataset** — <https://physionet.org/content/mimiciii/1.4/>

**Dataset** — <https://competition.huaweicloud.com/information/1000041487/circumstance>

## 1 Introduction

Uncovering causal relationships from observational data is a core challenge across many domains such as healthcare, economics, and cybersecurity. While most existing methods focus on continuous data, the problem of causal discovery in discrete events remains relatively underexplored. Yet, many real-world processes involve events occurring over time with causal dependencies.

Consider a clinical setting where patients are admitted to the emergency room and treated based on their evolving symptoms. We aim to understand which symptoms cause which treatments. Most causal discovery methods for event sequences rely on *Granger causality* (Granger 1969), which defines causality as a predictive relationship, a cause is something whose past improves the prediction of another variable. While useful for forecasting, this notion is not causal in the

interventional sense. Intervening on a “cause” identified by Granger methods may not affect the outcome. For instance, a correlated symptom might be selected as the cause of an antibiotic prescription, even though treating that symptom does not eliminate the underlying infection, and hence does not change the actual treatment decision.

More recent approaches (Qiao et al. 2023; Zhang et al. 2020; Cüppers et al. 2024) adopt *Pearl’s Structural Causal Model (SCM)* framework (Pearl 2009), enabling counterfactual reasoning and intervention. Some, such as CASCADE (Cüppers et al. 2024), do not support interactions among multiple parents and might infer that *admission* sometimes causes *antibiotic treatment*, but not always, missing that the treatment is only administered if an *infection* is ongoing at the time of admission. Other methods, such as SHP (Qiao et al. 2023), can model parent interactions but are still limited to *instantaneous* point events. Yet in many domains, it is not the start of an event that matters, but its persistence over time. Looking only at the start of a symptom may not be sufficient to explain repeated antibiotic administration.

This is precisely the gap we address. We introduce a causal model for *interval-based events*, where events have durations, may overlap, and interact to cause effects. Our model captures rich causal mechanisms, for example, a treatment triggered by an admission *only if* an infection is ongoing, or a medication causing periodic monitoring *while* it is being administered.

Our approach goes beyond recovering a causal graph: it identifies a full *Structural Causal Model*, including explicit causal mechanisms. Additionally, we recover which specific event occurrences caused each observed effect.

To discover this structure from data, we propose NIAGARA, a causal discovery algorithm grounded in the *Algorithmic Markov Condition* (AMC) (Janzing and Schölkopf 2010), which posits that the true causal model is the one that most compactly explains the data. As Kolmogorov complexity is not computable, we approximate it via the *Minimum Description Length* principle (Grünwald 2007), which has proven effective, for causal discovery in continuous and time-series data (Mian, Marx, and Vreeken 2021; Mameche et al. 2025).

Our main contributions are as follows:

- (a) We introduce a causal model for event sequences that supports event durations, multi-parent interactions, and conditioning-based causal mechanisms;
- (b) We prove identifiability of our model in the limit;

- (c) We define an MDL-based scoring function tailored to interval-based events;
- (d) We instantiate our model into NIAGARA, a practical causal discovery algorithm with theoretical guarantees.

## 2 Preliminaries

In this section we introduce the information-theoretic foundations underpinning our causal discovery approach.

### 2.1 Information-Theoretic Causal Discovery

Our approach builds on Pearl’s notion of causality (Pearl 2009), which assumes the existence of an underlying causal structure represented as a directed acyclic graph (DAG) over events. An edge from event  $e_i$  to  $e_j$  indicates that  $e_i$  is a direct cause (or parent) of  $e_j$  (child), and interventions on  $e_i$  would affect the distribution of occurrences of  $e_j$ .

To discover a causal graph from data, we rely on the *Algorithmic Markov Condition (AMC)* (Janzing and Schölkopf 2010), which postulates that the true causal model corresponds to the factorization of the joint distribution that yields the minimal Kolmogorov complexity (Li and Vitányi 1993). The Kolmogorov complexity  $K(x)$  of a finite binary string  $x$  is the length of shortest algorithmic description of  $x$ . For distribution  $P$  this is the length of the program that approximates  $P$  arbitrary well.

Applied to causal discovery over multivariate data  $X$ , the AMC implies that for the true causal graph  $\mathcal{G}$ ,

$$K(P(X)) \stackrel{\pm}{=} \sum_i K(P(X_i \mid pa(i))),$$

where  $X_i$  is a variable,  $pa(i)$  denotes its parents nodes in  $\mathcal{G}$ , and  $\stackrel{\pm}{=}$  denotes equality up to an additive constant.

As Kolmogorov complexity is not computable, we approximate it using the statistically well-founded *Minimum Description Length (MDL)* principle (Grünwald 2007).

### 2.2 Two-Part Minimum Description Length

The MDL principle states that the best explanation for data is the one that compresses it most. Given a model class  $\mathcal{M}$ , the best model  $\mathcal{M} \in \mathcal{M}$  minimizes the total code length

$$L(\mathcal{M}) + L(\mathcal{D} \mid \mathcal{M}),$$

where  $L(\mathcal{M})$  is the length of the model description (in bits), and  $L(\mathcal{D} \mid \mathcal{M})$  is the length of the data given  $\mathcal{M}$ .

Next, we introduce our causal model class  $\mathcal{M}$ .

## 3 Theory

We first describe our problem setting and then formalize our causal model for sequences of events with duration, illustrated with a toy example inspired by a medical setting.

### 3.1 Notations

We consider the problem of discovering causal relationships among discrete interval-based event sequences.

Let  $\Sigma = \{e_1, \dots, e_n\}$  be a set of events. We define an event occurrence as a tuple  $x = (t_s, t_e)$  where  $t_s, t_e \in [0, \mathcal{T}] \subset \mathbb{R}$  denote, respectively, the start and end time.  $E_i$  is the set of

occurrences of event  $e_i$ , and  $S_{i,s}$  denote the sequences of start times in  $E_i$ . An event  $e_i$  is said to be *ongoing* at time  $t$  if there exists an occurrence  $x \in E_i$  such that  $t_s \leq t < t_e$ . We define the function  $on : \Sigma \rightarrow \mathcal{I}$  that returns the union of time intervals  $I_i$  during which  $e_i$  is ongoing.

An example of an event are tachycardia episodes. Tachycardia occurrences denote episodes of high heart rate, with start and end times, between which tachycardia is ongoing.

### 3.2 Causal Model

We model each event  $e_i$  as being generated by a set of independent *data generating processes*  $\Gamma_i$ , each responsible for a subset of its occurrences. The union of all processes in  $\Gamma_i$  fully explains the observed occurrences of  $e_i$ .

Every event has at least a background process, modeled as a homogeneous Poisson process (PP), which captures spontaneous or noise-driven occurrences in the absence of causal influence. In our example, tachycardia might occur without apparent or measured cause, e.g., stress. This corresponds to its background generating process.

Causal effects arise when the occurrences or ongoing state of other events influence the generation of  $e_i$ . We distinguish between two types of causal influences,

- **Triggering:** An occurrence of another event  $e_j$  triggers an occurrence of  $e_i$  after some delay;
- **Conditioning:** The ongoing state of one or more parent events influence the generation of  $e_i$  either
  - (a) By modulating the intensity of a Poisson process generating occurrences of  $e_i$ , or
  - (b) By conditioning a triggering process, i.e., occurrences of another event trigger  $e_i$  only when the condition holds.

All these influences can act in combination. We formalize this via the notion of a causal mechanism.

**Definition 1 (Causal Mechanism).** A causal mechanism for event  $e_i$  is a tuple  $\kappa = (\tau_\kappa, F_\kappa)$  where:

- $\tau_\kappa \in \Sigma \cup \{\emptyset\}$  is an optional triggering event. If  $\tau_\kappa = \emptyset$ , the mechanism is non-triggering.
- $F_\kappa$  is a propositional logic formula over predicates of the form  $on(e_j)$  for  $e_j \in \Sigma$ , using conjunctions and negations. It defines the time intervals during which the mechanism is active.

Each causal mechanism induces a data generating process  $\gamma \in \Gamma_i$  that models the generation of occurrences of  $e_i$ .

**Definition 2 (Data Generating Process).** A data generating process for event  $e_i$  is a tuple  $\gamma = (T_\gamma, \theta_\gamma, \kappa_\gamma)$ , where:

- $T_\gamma \in \{PP, Triggered\}$  specifies the process type,
- $\theta_\gamma$  are the parameters of the process (e.g., rate or delay distribution),
- $\kappa_\gamma$  is the causal mechanism controlling this process (empty if it is the background process).

In our example, in addition to the background generating process, *tachycardia* may be triggered by the start of *antibiotic administration*, when another ongoing treatment interacts adversely with the antibiotics. This corresponds to

a generating mechanism  $\gamma$  with a causal mechanism  $\kappa_\gamma$  that includes both a triggering parent ( $\tau_\kappa$ ) and a conditioning parent ( $F_\kappa$ ). Second, we can imagine that ongoing antibiotic administration leads to regular monitoring of specific vital signs. This corresponds to a generating mechanism  $\gamma$  with a causal mechanism  $\kappa_\gamma$  that has a conditioning parent only, i.e.,  $\tau_\kappa = \emptyset$  and  $F_\kappa = \text{on}(\text{antibiotics})$ .

We show in Fig. 1 the four kind of data generating processes that can be induced by different causal mechanisms. A generating process without causal mechanism is a back-

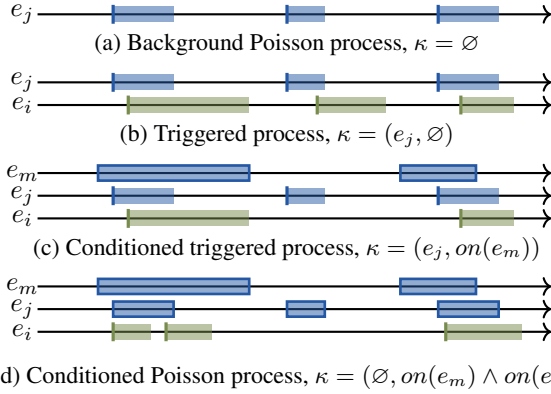


Figure 1: Data generating processes with different causal mechanisms  $\kappa = (\tau_\kappa, F_\kappa)$ . Each box represents an event occurrence from its start to its end. The cause events (blue) affect the generation of the effect event occurrences (green), either via triggering by their start (bold start), or via conditioning (bold box).

ground Poisson process (a), there is no cause. In (b), we have a triggered process without conditioning, the start of a parent occurrence causes (triggers) the start of a child occurrence. Then come the mechanisms that could not be captured without considering the duration of events. In (c), the causal mechanism is based on the interaction of a triggering parent ( $e_j$ ) and of a conditioning parent  $e_m$ , the triggering only occurs when the conditioning parent is ongoing. In (d), child occurrences are caused via a Poisson process that is active only when both  $e_m$  and  $e_j$  are ongoing.

The set of causal mechanisms  $K$  over  $\Sigma$  induces a causal graph  $\mathcal{G}$ , which does not encompass information about the causal mechanism nature or potential parent interaction.

**Definition 3 (Causal Graph).** A causal graph  $\mathcal{G} = (V, E)$  is a directed acyclic graph (DAG) where  $V$  is a finite set of nodes, each associated to an event  $e_i$ , and  $E$  is a finite set of edges where  $e_j \rightarrow e_i$  whenever  $e_j$  is involved in a causal mechanism  $\kappa \in K_i$ .

### 3.3 Data Modeling

Let  $e_i$  be generated through a set of data generating processes  $\Gamma_i$ , and  $E_{i,\gamma} \subseteq E_i$  be the subset of occurrences generated through a given process  $\gamma \in \Gamma_i$ . We now describe how we model each process  $\gamma$ , depending on  $\kappa_\gamma$ . We will see that, in all cases, we can associate  $E_{i,\gamma}$  to a sequence of delays  $\Delta_\gamma$  for which we define a probability density function  $\phi(d|\theta)$ .

This will allow a unified definition of the log-likelihood of the data given the model.

**Background Poisson Process** We recall that we assume that all events are generated by a background homogeneous Poisson process (Last and Penrose 2018)  $\gamma$  where  $T_\gamma = \text{Poisson}$  and  $\kappa_\gamma = \emptyset$  (e.g., Fig. 1 (a)). A Poisson process has a rate, or *intensity* parameter  $\lambda$  that determines the number of occurrences generated per time unit. In the case of events with causal parents, this background process corresponds to noise with a rate expected to be very low, i.e., we expect only a few of its occurrences to be generated by it. In a Poisson process, the delays between the start of two occurrences are independent and exponentially distributed. We can therefore model the sequence of delay between occurrences in  $E_{i,\gamma}$ .

The sequence of delays between the start of consecutive occurrences in  $E_{i,\gamma}$  is given by

$$\Delta_\gamma = \{t_{s,k} - t_{s,k-1}\}_{k=2}^{|E_{i,\gamma}|}.$$

Following the Poisson process assumption, each delay  $d \in \Delta_\gamma$  is drawn from an exponential distribution with rate parameter  $\lambda_\gamma$ . Accordingly, the probability density function of the delays is

$$\phi(d|\theta) = p_{\text{exp}}(d; \lambda_\gamma),$$

where  $\theta_\gamma = \{\lambda_\gamma\}$ .

In our running example, for tachycardia episodes generated by the background process, the delay sequence consists of the delays between the start of consecutive episodes.

**Conditioned Poisson Process** Events occurrences can also be generated by a Poisson process that is conditioned on one or more other events, in such case,  $T_\gamma = \text{Poisson}$ , and  $\kappa_\gamma = (\emptyset, F_\kappa)$  (e.g., Fig. 1 (d)). These processes are only “active” over specific intervals defined by the ongoing state of those conditioning events, more specifically over intervals  $I_\gamma$  defined by the evaluation of  $F_\kappa$ . In this particular case of *non-homogeneous* Poisson process (Last and Penrose 2018), the function rate is defined as

$$\lambda_\gamma(t) = \begin{cases} c_\gamma & \text{if } t \in I_\gamma, \\ 0 & \text{else.} \end{cases}$$

The sequence of delays between the start of consecutive occurrences in  $E_{i,\gamma}$  is given by

$$\Delta_\gamma = \left\{ \int_{t_{s,k-1}}^{t_{s,k}} \mathbb{1}_{I_\gamma}(t) dt \right\}_{k=2}^{|E_{i,\gamma}|},$$

where  $\mathbb{1}_{I_\gamma}$  is the indicator function that returns 1 if  $t \in I_\gamma$ , the integral simply corresponds to the delay between two occurrences disregarding the periods out of the intervals. The probability density function of the delays is

$$\phi(d|\theta) = p_{\text{exp}}(d; c_\gamma),$$

where  $\theta_\gamma = \{c_\gamma\}$ .

In our example with antibiotics causing monitoring, it means that we disregard the periods without antibiotic administration to compute the delay between two controls.

**(Conditioned) Triggered Process** A triggered process  $\gamma$  is a process where  $T_\gamma = \text{Triggered}$ , and  $\kappa_\gamma = (e_j \in \Sigma, F_\kappa)$  (e.g., Fig. 1 (b-c)). Unlike in a Poisson process, there is a one-to-one matching between the occurrences of the parent and of the child, the start of a parent occurrence triggers the start of a child one. If  $F_\kappa \neq \emptyset$ , it is conditioned by the ongoing state of other events and therefore only active over the interval defined by the formula, otherwise, it is active over  $I_\gamma = [0, \mathcal{T}]$ .

We assume that the start of an occurrence  $x = (t_s, t_e)$  of  $e_j$ , subject to  $t_s \in I_\gamma$ , triggers an occurrence of  $e_i$  with probability  $\alpha_\gamma$ , and after a delay  $d$  that follows an exponential distribution of intensity  $\lambda_\gamma$ . We assume that causes precede their effects, but the data capture process might make them appear simultaneous, meaning that triggering delays of  $d = 0$  may happen. The triggering probability  $\alpha_\gamma$  is here to account of noise in the process. A parent occurrence that did not trigger a child occurrence is associated with a delay of  $d = \infty$ . The global probability density function for all delays is

$$\phi(d|\theta_\gamma) = \begin{cases} 1 - \alpha_\gamma & \text{if } d = \infty, \\ p_{\text{exp}}(d; \lambda_\gamma) \cdot \alpha_\gamma & \text{else,} \end{cases}$$

where  $\theta = \{\alpha_\gamma, \lambda_\gamma\}$ .

In our example with triggered tachycardia, the delay sequence consists of the delays between the start of antibiotic administration and the start of triggered episodes.

### 3.4 Identifiability

Our causal model is identifiable in the limit, i.e., each unique set of causal mechanisms  $K$  over  $\Sigma$  will lead to a unique interarrival time distributions given infinite data.

**Theorem 1 (Identifiability).** *Assume the following conditions hold:*

- (i) *Triggered processes have strictly positive average delay (i.e., causes precede their effects);*
- (ii) *The causal mechanisms are stationary over time;*
- (iii) *The causal Markov condition, faithfulness, and sufficiency hold.*

*Then, the true causal model over interval-based events is identifiable in the limit from observational data.*

A complete proof of Theorem 1 is provided in Appendix A.1, leveraging the unique characteristics of the distributions induced by each possible data generating mechanism. In the next section we use the Minimum Description Length principle to instantiate our model.

## 4 MDL-based Model Selection

Following identifiability, asymptotically, the negative log-likelihood is sufficient to select the true causal model given observational data. However, in real scenarios, we only have access to a finite amount of data for causal discovery. We therefore base our approach on the algorithmic Markov condition, and approximate the Kolmogorov complexity using a two-part MDL score that we describe below. It balances the model's fit to the data—measured by the negative log-likelihood—against the model complexity, i.e., the model parameters and the combinations of causal parents.

**Model Cost** As the causal graph does not encompass information about the interactions between parents, we do not encode the graph but rather the description of all the causal mechanisms  $\kappa \in K$ , which implicitly describes the graph. For a given event  $e_i$ , we encode each of the causal mechanisms  $\kappa \in K_i$  associated to one of its generating processes  $\gamma \in \Gamma_i$  by identifying the triggering event  $\tau_\kappa$  (if any), the subset of  $n_F$  parents involved in  $F_\kappa$  (if any), and whether they are negated or not in the logic formula, leading to the cost function

$$L(K_i) = \sum_{\kappa \in K_i} \mathbb{1}_{\tau_\kappa \neq \emptyset} \log_2 |\Sigma| + \log_2 \binom{|\Sigma|}{n_F} + n_F.$$

The total model cost over all events  $L(K)$  is the sum of all event model costs.

**Data Cost** The cost of the data given the causal model corresponds to the negative log-likelihood of the event occurrences. As described in the previous section, we more specifically look at delays extracted according to the data generating mechanism considered. The cost of encoding the occurrences of an event  $e_i$  corresponds to

$$L(E_i|\Gamma_i) = \sum_{\gamma \in \Gamma_i} \sum_{d \in \Delta_{i,\gamma}} -\log_2(\phi_\gamma(d|\theta_\gamma)).$$

Again, the total data cost  $L(\mathcal{D}|\Gamma)$  is the sum of the data cost of all events, i.e.,  $L(\mathcal{D}|\Gamma) = \sum_{i=1}^n L(E_i|\Gamma_i)$ .

**Parameter Cost** Finally, we need to encode the parameters  $\Theta$  of the different data generating processes  $\Gamma$ ,

$$L(\Theta) = \sum_{\gamma \in \Gamma} \sum_{v \in \theta_\gamma} L_{\mathbb{R}}(v),$$

where  $L_{\mathbb{R}}(v) = L_{\mathbb{N}}(d) + L_{\mathbb{N}}(\lceil v \cdot 10^d \rceil) + 1$  is an encoding for real-valued parameter  $v$  up to precision  $p$ , with  $L_{\mathbb{N}}$  the MDL-optimal encoding for integer (Rissanen 1983).

**Overall MDL Cost** The global MDL cost of a causal model  $\mathcal{M}$  corresponds to the sum of the cost of the causal model, of its parameters, and of the data encoded by it,

$$L(\mathcal{M}, \mathcal{D}) = L(\Gamma) + L(\Theta) + L(\mathcal{D}|\Gamma).$$

The best causal model  $\mathcal{M}^* \in \mathcal{M}$  is the one achieving the smallest MDL cost,

$$\mathcal{M}^* = \arg \min_{\mathcal{M} \in \mathcal{M}} L(\mathcal{M}, \mathcal{D}).$$

## 5 Algorithm

To implement our approach into a practical causal discovery method, we must address three main challenges. First, because we consider causal mechanisms with interacting parents, we cannot identify the causal parents one by one. For example, if we look at the link between tachycardia and antibiotics over all patients, we will not see the causal relation as most of them are not under the interacting treatment. Second, as data is limited in real-world, we need to mitigate the uncertainty induced when selecting a data generating mechanism, even given exhaustive enumeration. Finally, we need to

find the correct partition of occurrences  $E_i$  for a given set of generating mechanisms  $\Gamma_i$ , i.e., which generating mechanism is responsible for the generation of each occurrence.

We first present our MDL-based strategy for causal mechanism selection given exhaustive enumeration. We then describe occurrence partitioning procedures that aim to minimize the overall negative log-likelihood. Finally, we optionally propose to estimate the causal superset of an event without having to fit all models to reduce the search space.

## 5.1 Causal Mechanism Selection

To select a model, we would ideally like to test all possible models and choose the one that minimizes our MDL objective. With infinite data and optimal matching, the minimizing model corresponds to the true causal graph, according to our identifiability guarantee. Unfortunately, testing all possible models is infeasible, even with finite data. We hence adopt a greedy approach, we select the causal mechanisms in decreasing order of certainty, all events considered at once, and use the acyclic constraint of the graph as a safety net. Given an event  $e_i$  and a set generating mechanisms  $\Gamma_{candidate}$  induced by causal mechanisms  $K_{candidate}$ , the certainty is measured by the gain in integrating the causal mechanisms  $L(E_i|\{\gamma_{background}\}) - L(E_i|\Gamma_{candidate} \cup \{\gamma_{background}\})$ .

## 5.2 Occurrence Partitioning

Let an event  $e_i$  be generated by a set of processes  $\Gamma_i$ . We first present how to distinguish occurrences generated by a triggered process  $E_{i,trig}$  from the occurrences generated by the background Poisson process  $E_{i,noise}$ . We then present how to disentangle multiple Poisson processes.

**Background Process and Triggered Processes** We consider the case of an event generated by a background Poisson process and one or more triggered (conditioned or not) processes. Here, the difficulty arises from the dual nature of noise, some parent occurrences did not trigger a child occurrence, and some child occurrences are due to its background noise process. We need to infer which parent occurrence (if not noise) triggered each child occurrence.

Given a set of generating mechanisms  $\Gamma_i$ , we partition occurrences of  $e_i$  across the background Poisson process  $\gamma_b$  and one or more triggered processes  $\gamma_p \in \Gamma_{trig}$  by minimizing the global negative log-likelihood via an iterative reassignment procedure, shown in Algorithm 1.

**Multiple Poisson Processes** We disentangle multiple Poisson processes by leveraging their estimated intensities over non-overlapping active intervals. First, we estimate the noise intensity over intervals where no conditioned process is active, and assign corresponding occurrences to it. Then, we estimate the intensity of each conditioned process using intervals where it is active alone. Assuming noise has lower intensity, we assign corresponding occurrences to the conditioned process. Occurrences in overlapping intervals are assigned proportionally and spread evenly, based on the inferred intensities (as detailed in Appendix A.3).

---

### Algorithm 1 OCCURRENCEMATCHING

---

**Require:** Child occurrences  $E_i$ , triggering parent occurrences  $E_{pa}$ , active intervals  $I_{trig}$   
**Ensure:** Matches  $m \subset \{(\{E_{pa} \cup \infty\} \times \{E_i \cup \infty\})^*\}$

- 1: Filter out the triggering parent occurrences  $E_{pa}$  with conditioning intervals  $I_{trig}$  (if any)
- 2: Initialize  $m$  by matching each  $c \in E_i$  to nearest past  $p \in E_{pa}$
- 3: Compute initial cost
- 4:  $L(E_i|\Gamma_i) = L(E_{i,b}|\gamma_b) + \sum_{\gamma_p \in \Gamma_{i,trig}} L(E_{i,p}|\gamma_p)$
- 5: **while** the cost can be improved **do**
- 6:   Identify  $(p, c) \in m$  with highest delay cost
- 7:   Remove  $c$  from  $E_{i,p}$  and assign  $c$  to background
- 8:   Recompute  $L'(E_i|\Gamma_i)$
- 9: **end while**

---

## 5.3 Causal Superset Estimation

To handle large search space, we propose an optional step of causal superset estimation. To prune the search space, we use simple conditional frequency-based heuristics to detect potential candidate mechanisms. It requires two parameters, the length  $w$  of the time window to consider, and the maximum of noise assumed  $\epsilon_{max}$ , which control the maximal proportion  $\epsilon_{max}$  of child occurrences generated by the background process, and the maximal triggering rate  $\alpha = 1 - \epsilon_{max}$ .

For each pair of events  $e_i$  (child) and  $e_j$  (parent), we compute the following statistics,

- $p_w(x_i|x_j)$ , the proportion of child occurrences preceded by a parent occurrence within  $w$  time units,
- $p_w(x_j|x_i)$ , the proportion of parent occurrences followed by a child occurrence within  $w$  time units (triggering probability  $\hat{\alpha}$ ),
- $p(x_i|on(e_j))$ , the proportion of child occurrences while the parent is ongoing,
- $p_w(x_i|on(e_j))$ , the proportion of child occurrences preceded by the parent being ongoing within  $w$  time units.

High value of both  $p_w(x_i|x_j)$  and  $p_w(x_j|x_i)$  indicate strong alignment between  $e_i$  and  $e_j$ , suggesting a triggered process. A high  $p(x_i|on(e_j))$ , in the other hand, suggests a conditioned Poisson process. By combining these statistics for multiple candidate parents, we can also infer mechanisms with parent interaction. For instance, a high  $p_w(x_i|x_j)$  (even if  $p_w(x_j|x_i)$  is low) and high  $p_w(x_i|on(e_j))$  may reflect a conditioned triggering mechanism.

A statistic is considered high if  $p > (1 - \epsilon_{max})$ . The parameters  $w$  and  $\epsilon_{max}$  control the strictness of the selection, larger values will be more permissive and result in more selected mechanisms. Full selection rules covering all combinations with up to two parents are provided in the Appendix A.3.

## 5.4 Computational Complexity

For each event, we evaluate a bounded number of candidate causal mechanisms involving up to a finite number of parents. The most expensive operation—iterative occurrence matching for triggered processes—has a worst-case complexity of

$O(n_{ch}^2)$ , where  $n_{ch}$  is the number of child occurrences. Overall, our method has a worst-case complexity of  $O(|\Sigma| \cdot n_{ch}^2)$ , which is polynomial in the number of occurrences. A full derivation, along with the complexity for the unbounded case (which is exponential in the number of events), is provided in the Appendix A.3.

## 5.5 Guarantees

Given the correct partition of occurrences across data generating mechanisms  $\Gamma$ , our method is guaranteed to recover the true causal model under exhaustive enumeration and in the limit. This follows directly from the model identifiability results and the assumption of the causal Markov condition.

When using the causal superset estimation procedure, we guarantee that the inferred set of candidate mechanisms is a superset of the true ones, assuming a maximal noise tolerance of  $\epsilon_{max} = 1$  and a time window of length  $w = \mathcal{T}$ .

However, the occurrence partitioning procedures themselves are not guaranteed to recover the ground truth. For triggered processes, the algorithm minimizes the negative log-likelihood, but it may converge to a local minimum. Even in the simpler case of a single triggered process, matching causes to effects is ambiguous when earlier causes produce effects that occur after those of later causes. For Poisson processes, accurate partitioning depends on low noise and sufficient non-overlapping active intervals for each process. In less favorable conditions, the approximation may degrade.

That said, our global algorithm is deterministic and guaranteed to converge to a local minimum, as the two iterative procedures—triggered process partitioning and causal mechanism selection—monotonically decrease the total cost until termination.

## 6 Related Work

Causal discovery from observational data has been extensively studied, though primarily in the context of continuous variables (Spirtes and Zhang 2016). For discrete tabular data, identifiability is harder to achieve; approaches typically rely on variable conditional distribution and model selection principles like BIC (Cai et al. 2018) or MDL (Budhathoki and Vreeken 2018a). However, these methods do not naturally extend to event sequences, where variables are binary (event/no event). Time can provide additional structure required for causal discovery.

In the event sequence setting, most existing work adopts the notion of *Granger causality* (Granger 1969), where an event is said to cause another if its past helps to predict it. Early methods include MDL-based CUTE (Budhathoki and Vreeken 2018b), while recent approaches use autoregressive models (Liu et al. 2024) or neural point processes with attribution, as in CAUSE (Zhang et al. 2020), to capture more complex relationships. Hawkes processes are also widely used, often with sparsity-inducing regularization (Xu, Farajtabar, and Zha 2016; Salehi et al. 2019; Jalaldoust, Hlaváčková-Schindler, and Plant 2022; Cai et al. 2022).

Beyond prediction, few methods address causal discovery in the sense of Pearl, with identifiability guarantees. SHP (Qiao et al. 2023) extends Hawkes processes to leverage

both lagged and instantaneous effects. Closer to our work, CASCADE (Cüppers et al. 2024) uses the Algorithmic Markov Condition to infer causality based on triggering mechanisms. Our approach differs by a wider range of causal mechanisms, by allowing parent interactions, and assuming additive noise.

Finally, interval-based events have been considered recently (Ho et al. 2023), where several scoring functions were proposed to assess pairwise relationships. However, in the absence of formal identifiability guarantees and rigorous evaluation, these metrics remain associative rather than truly causal.

## 7 Experiments

We implemented NIAGARA with support for up to two interacting parents per event—available online—and conducted an experimental study to answer the following questions:

- Q.1 Do the identifiability and algorithmic guarantees hold empirically?
- Q.2 How robust is NIAGARA to finite data, large event sets, and noise?
- Q.3 How does NIAGARA compare to state-of-the-art causal discovery methods on real-world data?

We first evaluate on controlled synthetic data and then present results on a real-world dataset and a medical case study.

### 7.1 Synthetic Data

We first describe how we generated the synthetic data, then present and discuss the main results.

**Methodology** We generated synthetic data according to our causal model under varying parameters: number of events (default  $|\Sigma| = 10$ ), observation period (default  $\mathcal{T} = 500,000$ ), and noise level (default  $\epsilon_{max} = 0.05$ ).

We report the performance of our approach both with and without the causal superset estimation procedure, referred to as NIAGARA and NIAGARA-heuristic, respectively. We compare against the most relevant existing methods, CAUSE and SHP, which can capture complex interactions between causal parents, and CASCADE, which is based on the AMC and considers triggering mechanisms.

We evaluate the discovered causal model by comparing the predicted graphs with the true graph using with the structural Hamming distance (SHD) which is a standard graph distance, and the structural intervention distance (SID) (Peters and Bühlmann 2015) which reflects differences in implied causal effects and is defined only for acyclic graphs (thus excluding CAUSE)), and the F1 score which is the harmonic mean of precision and recall over the graph edges.

**Identifiability** Figure 2a shows that NIAGARA recovers the true causal graph as the observation period  $\mathcal{T}$  increases, confirming identifiability in the large-data limit. Yet, the score remains close to perfect on limited data.

**Scalability** Figure 2b demonstrates that performance remains stable as the number of events—and thus the number of possible mechanisms—increases. The causal superset heuristic substantially reduces runtime (from a median runtime of 8346 s to 793 s at 40 events) without degrading performance.

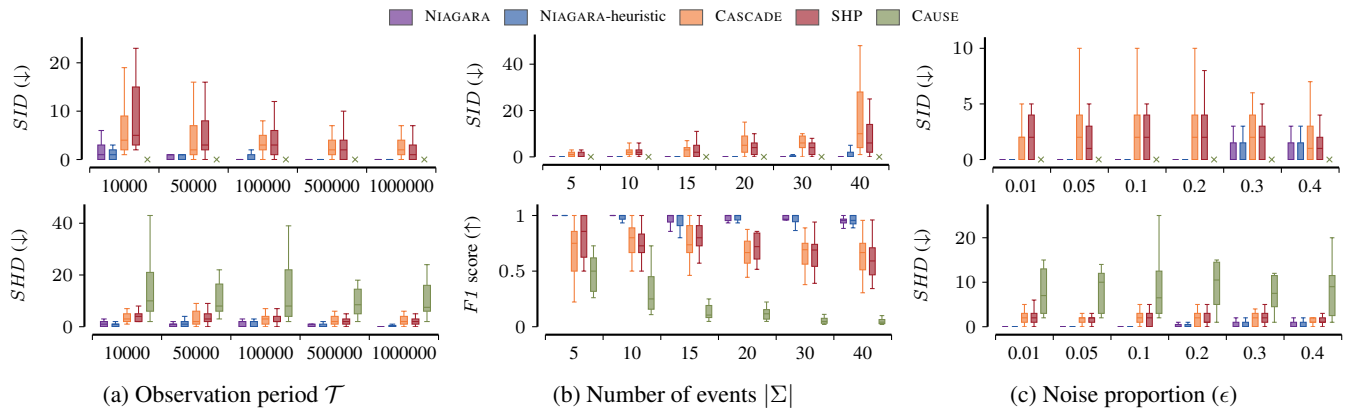


Figure 2: Results on the synthetic data. On average, 375 occurrences of source events are generated by default, which in turn cause other events. For  $\mathcal{T} = 10,000$ , it goes down to 7 occurrences, and for  $\mathcal{T} = 1,000,000$ , it goes up to 750 occurrences.

**Noise** To simulate increasing noise, we raised the proportion  $\epsilon$  of child occurrences generated by the background process, while proportionally decreasing the triggering rate ( $\alpha = 1 - \epsilon$ ). We observed that the performance of NIAGARA only starts decreasing from a noise proportion of  $\epsilon = 0.3$ , demonstrating a good robustness to noise.

## 7.2 Real-World Data

We applied NIAGARA to the three public datasets of the PCIC 2021 causal discovery challenge. The data consists of alarms over time intervals across telecommunication devices. We achieved F1 scores of 0.57 for 18V\_55N\_Wireless, 0.35 for 24V\_439N\_Microwave, and 0.36 for 25V\_474N\_Microwave, which is competitive with methods specifically designed for topological data (Cai et al. 2022), and significantly superior to most of the other scores reported for this dataset (details in Appendix A.4). This shows that our approach is competitive for real-world causal discovery.

## 7.3 Medical Case Study

MIMIC-III (Johnson et al. 2016) is a large publicly available database containing healthcare data from patients admitted to the intensive care units of the Beth Israel Deaconess Medical Center in Boston. It includes information such as medical procedures and vital sign measurements, typically recorded at a coarse temporal resolution (e.g., hourly for most events).

**Methodology** For this case study, we focused on patients diagnosed with sepsis to form a clinically homogeneous cohort, yielding 1,184 patients. To define our event set, we extracted all entries related to service transfers, fluid administration, and medical procedures, which all have a start and an end time, resulting in 477,491 occurrences over 532 distinct events. Since test results and clinical notes were not included, the dataset is subject to significant confounding.

**Findings** Sepsis treatment primarily involves antibiotics and intravenous fluids to support blood pressure and organ function (Evans et al. 2021). Among the most confident causal links, we found that antibiotic administrations often co-occurred with dextrose (used as a solvent), and

were conditioned on the absence of other ongoing antibiotics (e.g.,  $Dextrose5\% \xrightarrow{\kappa} PenicillinG \xrightarrow{\kappa} Vancomycin$  with  $\tau_{\kappa} = Dextrose5\%$  and  $F_{\kappa} = \neg PenicillinG$ ). These mechanisms suggest that antibiotics are typically given one at a time, especially when their microbial spectra overlap.

Ongoing invasive ventilation appears in numerous causal mechanisms leading to conditioned Poisson processes. It causes periodic filter replacements, administration of pulmonary-specific nutritional formulas to feed the patient via tube, and the use of muscle relaxants to facilitate mechanical ventilation (e.g.,  $Invasive\ ventilation \xrightarrow{\kappa} Filter\ change$  with  $\tau_{\kappa} = \emptyset$  and  $F_{\kappa} = Invasive\ ventilation$ ). Interestingly, it also appeared to cause notifications to organ donation services, which may reflect institutional protocols for identifying potential organ donors in patients with a poor prognosis.

The discovered relationships align well with everyday understanding and provide strong, intuitive support for NIAGARA validity. Moreover, other causal discovery methods would not have been able to discover them, as they cannot express such parent relations and/or do not consider the time periods over which events are ongoing.

## 8 Conclusion

In this paper, we addressed the problem of causal discovery from interval-based event sequences. We introduced a causal model for events with durations that supports rich causal structures, including interactions between multiple parent events, and proved that it is identifiable in the limit. To recover causal structure from purely observational data, we proposed an MDL-based score and introduced NIAGARA, a practical causal discovery algorithm. Experiments on synthetic data show that NIAGARA reliably recovers the ground truth model and outperforms state-of-the-art methods. On real-world medical data, it discovered causal mechanisms that align well with established domain knowledge. Like all causal discovery approaches, ours relies on assumptions—most notably, causal sufficiency. A promising direction for future work is to investigate how this assumption can be relaxed to identify confounding in event sequences.

## References

- Budhathoki, K.; and Vreeken, J. 2018a. Accurate causal inference on discrete data. In *2018 IEEE International Conference on Data Mining (ICDM)*, 881–886. IEEE.
- Budhathoki, K.; and Vreeken, J. 2018b. Causal inference on event sequences. In *Proceedings of the 2018 SIAM International Conference on Data Mining*, 55–63. SIAM.
- Cai, R.; Qiao, J.; Zhang, K.; Zhang, Z.; and Hao, Z. 2018. Causal Discovery from Discrete Data using Hidden Compact Representation. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Cai, R.; Wu, S.; Qiao, J.; Hao, Z.; Zhang, K.; and Zhang, X. 2022. THPs: Topological Hawkes processes for learning causal structure on event sequences. *IEEE Transactions on Neural Networks and Learning Systems*.
- Chickering, D. M. 2002. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov): 507–554.
- Cüppers, J.; Xu, S.; Musa, A.; and Vreeken, J. 2024. Causal discovery from event sequences by local cause-effect attribution. *Advances in Neural Information Processing Systems*, 37: 24216–24241.
- Evans, L.; Rhodes, A.; Alhazzani, W.; Antonelli, M.; Coopersmith, C. M.; French, C.; Machado, F. R.; McIntyre, L.; Ostermann, M.; Prescott, H. C.; et al. 2021. Surviving sepsis campaign: international guidelines for management of sepsis and septic shock 2021. *Critical care medicine*, 49(11): e1063–e1143.
- Granger, C. W. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, 424–438.
- Grünwald, P. 2007. *The Minimum Description Length Principle*. MIT Press.
- Ho, N.; Le, T. C.; Huynh, N. T.; Ngo, L.; et al. 2023. Causal Associations between Temporal Events. In *2023 IEEE International Conference on Big Data (BigData)*, 1135–1142. IEEE.
- Jalaldoust, A.; Hlaváčková-Schindler, K.; and Plant, C. 2022. Causal discovery in Hawkes processes by minimum description length. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 6978–6987.
- Janzing, D.; and Schölkopf, B. 2010. Causal inference using the algorithmic Markov condition. *IEEE Transactions on Information Theory*, 56(10): 5168–5194.
- Johnson, A. E.; Pollard, T. J.; Shen, L.; Lehman, L.-w. H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Anthony Celi, L.; and Mark, R. G. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3(1): 1–9.
- Last, G.; and Penrose, M. 2018. *Lectures on the Poisson process*, volume 7. Cambridge University Press.
- Li, M.; and Vitányi, P. 1993. *An Introduction to Kolmogorov Complexity and its Applications*. Springer.
- Liu, Y.; Cai, R.; Chen, W.; Qiao, J.; Yan, Y.; Li, Z.; Zhang, K.; and Hao, Z. 2024. TNPART: topological neural poisson autoregressive model for learning granger causal structure from event sequences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 20491–20499.
- Mameche, S.; Cornanguer, L.; Ninad, U.; and Vreeken, J. 2025. SPACETIME: Causal Discovery from Non-Stationary Time Series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 19405–19413.
- Mian, O. A.; Marx, A.; and Vreeken, J. 2021. Discovering fully oriented causal networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 8975–8982.
- Pearl, J. 2009. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2nd edition.
- Peters, J.; and Bühlmann, P. 2015. Structural intervention distance for evaluating causal graphs. *Neural computation*, 27(3): 771–799.
- Qiao, J.; Cai, R.; Wu, S.; Xiang, Y.; Zhang, K.; and Hao, Z. 2023. Structural Hawkes Processes for Learning Causal Structure from Discrete-Time Event Sequences. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*, 5702–5710. ijcai.org.
- Rissanen, J. 1983. A Universal Prior for Integers and Estimation by Minimum Description Length. 11(2): 416–431.
- Salehi, F.; Trouleau, W.; Grossglauser, M.; and Thiran, P. 2019. Learning hawkes processes from a handful of events. *Advances in neural information processing systems*, 32.
- Spirtes, P.; and Zhang, K. 2016. Causal discovery and inference: concepts and recent methodological advances. In *Applied informatics*, volume 3, 3. Springer.
- Xu, H.; Farajtabar, M.; and Zha, H. 2016. Learning granger causality for hawkes processes. In *International conference on machine learning*, 1717–1726. PMLR.
- Zhang, W.; Panum, T.; Jha, S.; Chalasani, P.; and Page, D. 2020. Cause: Learning granger causality from event sequences using attribution methods. In *International Conference on Machine Learning*, 11235–11245. PMLR.