

TNPAR: Topological Neural Poisson Auto-Regressive Model for Learning Granger Causal Structure from Event Sequences

Yuequn Liu¹, Ruichu Cai^{1,2*}, Wei Chen¹, Jie Qiao¹, Yuguang Yan¹, Zijian Li^{1,3}, Keli Zhang⁴, Zhifeng Hao⁵

¹School of Computer Science, Guangdong University of Technology, Guangzhou, China

²Peng Cheng Laboratory, Shenzhen, China

³Machine Learning Department, Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, United Arab Emirates

⁴Huawei Noah's Ark Lab, Huawei, Paris, France

⁵College of Science, Shantou University, Shantou, China

{weanyq, cairuichu, chenweidelight, qiaojie.chn}@gmail.com, ygyan@gdut.edu.cn, leizigin@gmail.com, zhangkeli1@huawei.com, haozhifeng@stu.edu.cn

Abstract

Learning Granger causality from event sequences is a challenging but essential task across various applications. Most existing methods rely on the assumption that event sequences are independent and identically distributed (i.i.d.). However, this i.i.d. assumption is often violated due to the inherent dependencies among the event sequences. Fortunately, in practice, we find these dependencies can be modeled by a topological network, suggesting a potential solution to the non-i.i.d. problem by introducing the prior topological network into Granger causal discovery. This observation prompts us to tackle two ensuing challenges: 1) how to model the event sequences while incorporating both the prior topological network and the latent Granger causal structure, and 2) how to learn the Granger causal structure. To this end, we devise a unified topological neural Poisson auto-regressive model with two processes. In the generation process, we employ a variant of the neural Poisson process to model the event sequences, considering influences from both the topological network and the Granger causal structure. In the inference process, we formulate an amortized inference algorithm to infer the latent Granger causal structure. We encapsulate these two processes within a unified likelihood function, providing an end-to-end framework for this task. Experiments on simulated and real-world data demonstrate the effectiveness of our approach.

Introduction

Causal discovery from multi-type event sequences is important in many applications. In the realm of intelligent operation and maintenance, identifying the causal structure behind alarm sequences can expedite the location of root causes (Vuković and Thalmann 2022). In social analysis, recovering the causal structure behind users' behavior sequences can inform the development of effective advertising strategies (Chen et al. 2020).

Lots of methods have been proposed for causal discovery from multi-type event sequences. They fall into two main categories: 1) constraint-based methods, which utilize independence-based tests or measures to estimate the




Figure 1: An example of the topological event sequences generated by a mobile network in an operation and maintenance scenario. In this context, $\{n_1, n_2, n_3\}$ represent the network elements, and $\{v_1, v_2, v_3\}$ denote the event types of the alarms. The term $P_{v_1}(t_1)$ stands for the distribution of v_1 alarm at t_1 timestamp, and so on. Black arrows represent the correct causal edges, whereas the red arrow represents an incorrect edge. (a) Generation process of topological event sequences. The event distribution varies with time and is influenced by both a topological network (connected by the solid lines) and a Granger causal structure (depicted in part (b)). (b) Ground truth of the Granger causal structure. (c) Granger causal structure learned under i.i.d. assumption.

causal structure, and 2) point process-based methods, which employ point processes to model the generation process of event sequences. While the former, including methods such as PCMCI (PC with Momentary Conditional Independence test) (Runge 2020) and transfer entropy-based methods (Chen et al. 2020; Mijatovic et al. 2021), is founded on strict assumptions of the causal mechanism assumptions, the latter, like Hawkes process-based methods (Xu, Farajtabar, and Zha 2016; Zhou, Zha, and Song 2013) and neural point process models (Shchur et al. 2021), is based on the concept of Granger causality (Granger 1969). Given the weaker assumption of Granger causality, identifying the Granger causal structure just by assessing if a sequence is predictive of another, point process-based methods often find preference in real-world applications (Shojaie and Fox 2022) and are therefore the focal point of this work.

*Corresponding author.

However, these existing point process-based methods are significantly dependent on the i.i.d. assumption to model the generation process of event sequences. In reality, this assumption often falls short due to the complex dependencies that permeate sequences. Consequently, the performance of these methods often leaves much to be desired in real-world applications. For instance, as illustrated in Fig. 1, non-i.i.d. alarm sequences in an operation and maintenance scenario highlight this issue. These event sequences are produced by network elements interconnected by a topological network, indicating the event sequences of different elements are dependent on each other, thereby violating the i.i.d. assumption. In this illustrated scenario, the event type v_2 in n_2 is only caused by the event types in topologically connected elements (specifically, v_1 in n_2). Existing methods would incorrectly identify the edge $v_3 \rightarrow v_2$ because they erroneously treat event sequences on different nodes as independent. By considering the prior topological network, we can correctly learn that v_2 is actually caused by v_1 and discern the true causal structure. Therefore, it is vital to exploit the topological network to model the dependencies among non-i.i.d. event sequences and accurately recover the causal structure from them.

This gives rise to two ensuing challenges: 1) how to explicitly model the generation process of topological event sequences, incorporating both the prior topological network and the latent causal structure, and 2) how to learn the latent Granger causal structure from topological event sequences. To tackle these challenges, Cai et al. (Cai et al. 2022) propose the topological Hawkes process (THP), which extends the time-domain convolution property of the Hawkes process to the time-graph domain. Nevertheless, THP’s applicability in real-world scenarios remains limited due to the complex generation process and distribution of real-world scenarios, which are often too intricate to be accurately represented by the Hawkes process. Hence, it is crucial to devise a distribution-free generation model for non-i.i.d. event sequences that can incorporate both the topological network and the latent Granger causal structure.

In response to these challenges, we develop a unified Topological Neural Poisson Auto-Regressive (TNPARG) model, comprising a generation process and an inference process. In the generation process, we introduce a variant of the Poisson Auto-Regressive model to depict the generation process of topological event sequences. By extending the Auto-Regressive model with Poisson Process and incorporating both the topological network and the Granger causal structure, TNPARG effectively overcomes the non-i.i.d. problem. By representing the distribution using a series of Poisson processes with varying parameters, TNPARG displays enhanced flexibility compared to existing methods that employ a single fixed distribution for the entire sequence. In the inference process, we design an amortized inference algorithm (Zhang et al. 2018; Gershman and Goodman 2014) to effectively learn the Granger causal structure. By integrating these generation and inference processes, we establish a unified likelihood function for TNPARG model, which can be optimized within an end-to-end paradigm.

In summary, our main contributions are as follows:

1. We develop a topological neural Poisson auto-regressive method that accurately learns causal structure from topological event sequences.
2. By incorporating the prior topological network into the generation model, we offer a distribution-free solution to the non-i.i.d. challenge in causal discovery.
3. By treating the causal structure as a latent variable, we devise an amortized inference method to effectively identify the causal structure among the event types.

Related work

Temporal Point Process

Temporal point processes are stochastic processes utilized to model event sequences. They can be bifurcated into two main categories: statistical point processes and neural point processes. On the one hand, statistical point processes emphasize designing appropriate intensity functions, the parameters of which often carry specific physical interpretations. Notable examples of statistical point processes encompass the Poisson process (Cox 1955), Hawkes process (Hawkes 1971), reactive point process (Ertekin, Rudin, and McCormick 2015), and self-correcting process (Isham and Westcott 1979), among others. On the other hand, neural point processes (Shchur et al. 2021) exploit the potent learning capabilities of neural networks to implement the intensity functions, often outperforming statistical point process methods in prediction performance.

Granger Causality for Event Sequences

There are many types of methods for Granger causal discovery from event sequences. For instance, the Hawkes-process-based methods assume that the past events stimulate the occurrence of related events in the future if and only if the former Granger caused the latter. Typical Hawkes process-based methods (Zhou, Zha, and Song 2013; Xu, Farajtabar, and Zha 2016) focus on designing appropriate intensity functions and regularized techniques. Recently, the THP algorithm (Cai et al. 2022) extends the Hawkes process to address the non-i.i.d. issue. Other methods like Graphical Event Models (GEMs) (Bhattacharjya, Subramanian, and Gao 2018) formalize Granger causality using process independence. In contrast, neural point process-based methods (Xiao et al. 2019; Zhang et al. 2020) do not depend on specific assumptions of generation functions like the Hawkes process or GEMs. Instead, they employ neural networks to model event sequences with Granger causality. In another vein, our work is also related to the Granger causal discovery from time series. For instance, Brillinger (Brillinger 1994) aggregates event sequences into time series, which enables the analysis of event sequences using auto-regressive models. The Granger Causality alignment (GCA) model (Li et al. 2023) combines the auto-regressive model with the generation model to learn the Granger causality. Amortized Causal Discovery (ACD) (Löwe et al. 2022) trains an amortized model to infer causal graphs from time series. In addition, we provide a comparison of our work with related methods in Appendix A.

Granger Causal Discovery Using Neural Poisson Auto-Regressive

Problem Definition

Let an undirected graph $\mathcal{G}_N(\mathbf{N}, \mathbf{E}_N)$ represent the topological network among nodes \mathbf{N} , and a directed graph $\mathcal{G}_V(\mathbf{V}, \mathbf{E}_V)$ represent the Granger causal graph among event types \mathbf{V} . Here, \mathbf{E}_N signifies the physically connected edges between nodes, and \mathbf{E}_V denotes the causal edges between event types. In this context, an event could trigger its subsequent events in its own node and its topologically connected nodes. A series of *topological event sequences* of length m , symbolized as $\mathbf{X} = \{(v_i, n_i, t_i) | i \in \{1, \dots, m\}\}$, are generated by a Granger causal structure $\mathcal{G}_V(\mathbf{V}, \mathbf{E}_V)$ within a topological network $\mathcal{G}_N(\mathbf{N}, \mathbf{E}_N)$. Each item in the series consists of $v_i \in \mathbf{V}$ representing the event type, $n_i \in \mathbf{N}$ representing the topological node, and $t_i \in [0, T]$ representing the occurrence timestamp.

We can consider the topological event sequences \mathbf{X} as a set of counting processes $\{C_{v_i, n_j}(t) | v_i \in \mathbf{V}, n_j \in \mathbf{N}, t \in [0, T]\}$, where $C_{v_i, n_j}(t)$ denotes the count of event type v_i that has occurred up to t at n_j . In this work, we divide the continuous interval $[0, T]$ into $\lceil T/\Delta \rceil$ small intervals, where $\Delta \in \mathbb{R}^+$ can be determined according to real-world applications. Then, the occurrence numbers of the counting processes can be denoted as $\mathbf{O}^{V, N} = \{O_t^{v_i, n_j} | t \in \{1, \dots, \lceil T/\Delta \rceil\}, i \in \{1, \dots, |\mathbf{V}|\}, j \in \{1, \dots, |\mathbf{N}|\}\}$, where $O_t^{v_i, n_j} = C_{v_i, n_j}(t \times \Delta) - C_{v_i, n_j}((t-1) \times \Delta)$ represents the occurrence number of v_i at n_j within $((t-1) \times \Delta, t \times \Delta]$. Within each interval, we assume the counting process follows a Poisson process (Stoyan et al. 2013) with parameter $\lambda_t^{v_i, n_j}$. Thus, the probability of $O_t^{v_i, n_j}$ can be expressed as follows:

$$P(O_t^{v_i, n_j} = o) = \frac{(\lambda_t^{v_i, n_j} \Delta)^o}{o!} e^{-\lambda_t^{v_i, n_j} \Delta}; o = 0, 1, \dots \quad (1)$$

Consequently, we can formulate the problem of causal discovery from topological event sequences as:

Definition 1 (Granger causal discovery from topological event sequences). *Given the occurrence numbers of topological event sequences, $\mathbf{O}^{V, N} = \{O_t^{v_i, n_j} | t \in \{1, \dots, \lceil T/\Delta \rceil\}, i \in \{1, \dots, |\mathbf{V}|\}, j \in \{1, \dots, |\mathbf{N}|\}\}$, and the prior topological network $\mathcal{G}_N(\mathbf{N}, \mathbf{E}_N)$, the goal of Granger causal discovery from topological event sequences is to infer the Granger causal structure $\mathcal{G}_V(\mathbf{V}, \mathbf{E}_V)$ among the event types.*

In the following, we propose the TNPARG model to solve the above problem, comprising a generation process and an inference process.

Generation Process of the Event Sequences via Topological Neural Poisson Auto-Regressive Model

TNPARG's generation process is depicted using solid lines in Fig. 2. In this process, $O_t^{v_i, n_j}$ is determined by a combination of historical data $\{\mathbf{O}_{t-1}^{V, N}, \mathbf{O}_{t-2}^{V, N}, \dots, \mathbf{O}_{t-\Omega}^{V, N}\}$, the causal matrices $\mathbf{A}_{0:K}$, and the topological matrices $\mathbf{B}_{0:K}$. Here, $\mathbf{A}_{0:K}$ represents a set of matrices $\{\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_K\}$, where \mathbf{A}_k is a $|\mathbf{V}| \times |\mathbf{V}|$ binary matrix that indicates the




Figure 2: An illustration of the generation and inference processes for TNPARG. In this figure, $\mathbf{A}_{0:K}$ represents the causal matrices of \mathcal{G}_V , and $\mathbf{B}_{0:K}$ represents the topological matrices of \mathcal{G}_N . Solid lines signify the generation process for the data $O_t^{v, n}$, and dashed lines correspond to the inference process for the causal matrices $\mathbf{A}_{0:K}$. Observed variables are denoted by the solid circles, with the latent variables represented by the dashed circles.

Granger causality between event types at a geodesic distance (Bouttier, Di Francesco, and Guitter 2003) of k within \mathcal{G}_N . Let $A_k^{i, j}$ denote the element in row i and column j of \mathbf{A}_k . Then, if $A_k^{i, j} = 0$, it implies that, at a geodesic distance of k , event type v_i does not have a Granger causality with v_j . Otherwise, $A_k^{i, j}$ will be 1. Similarly, $\mathbf{B}_{0:K}$ corresponds to another set of matrices $\{\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_K\}$, each of which \mathbf{B}_k is a $|\mathbf{N}| \times |\mathbf{N}|$ binary matrix that indicates the physical connections of nodes at a geodesic distance of k . Let $B_k^{i, j}$ denote the element in row i and column j of \mathbf{B}_k . If the geodesic distance between node n_i and n_j is k in \mathcal{G}_N , then $B_k^{i, j} = 1$. Otherwise, $B_k^{i, j}$ will be set to 0.

In this work, we devise a model, combining the neural auto-regressive model and the Poisson process, to describe the above generation process. To initiate, we introduce a traditional neural auto-regressive model with Ω time lags:

$$O_t^{v_i} = g(\mathbf{O}_{t-1:t-\Omega}^V) + \epsilon_t^{v_i}, \quad (2)$$

where $\mathbf{O}_{t-1:t-\Omega}^V = \{\mathbf{O}_{t-1}^V, \mathbf{O}_{t-2}^V, \dots, \mathbf{O}_{t-\Omega}^V\}$ represents a set of historical data, $g(\cdot)$ is a nonlinear function implemented by a neural network that generates data in the t -th time interval based on the historical data, and $\epsilon_t^{v_i}$ denotes the noise term for event type v_i in the t -th time interval. In order to incorporate the topological relationships of nodes into our model, we extend $\mathbf{O}_{t-1:t-\Omega}^V$ to include the topologically connected historical data, using the function $f_{n_j}^{\mathcal{G}_N}(\mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{B}_{0:K})$. This function applies a filter to set the non-topologically connected data for node n_j to 0 in $\mathbf{O}_{t-1:t-\Omega}^{V, N}$ based on $\mathbf{B}_{0:K}$, which effectively removes the non-topologically connected information in the generation process of $O_t^{v_i}$. Building upon this, our auto-regressive model, incorporating topological information, can be expressed as follows:

$$O_t^{v_i, n_j} = g(f_{n_j}^{\mathcal{G}_N}(\mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{B}_{0:K})) + \epsilon_t^{v_i, n_j}. \quad (3)$$

Furthermore, to incorporate the Granger causal structure into the above model, we substitute $\mathbf{O}_{t-1:t-\Omega}^{V, N}$ in Eq. (3) with a function $f_{v_i}^{\mathcal{G}_V}(\mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{A}_{0:K})$. This function acts as

a filter to set the historical data that do not have a Granger causality with v_i to 0 in $\mathbf{O}_{t-1:t-\Omega}^{V,N}$ based on $\mathbf{A}_{0:K}$. By incorporating both the topological network and the Granger causal structure, we propose the following model for generating topological event sequences:

$$O_t^{v_i, n_j} = g(f_{n_j}^{\mathcal{G}_N}(f_{v_i}^{\mathcal{G}_V}(\mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{A}_{0:K}), \mathbf{B}_{0:K})) + \epsilon_t^{v_i, n_j}. \quad (4)$$

Inference Process of the Granger Causal Structure via Amortized Inference

The inference process of TNPAR is depicted using dashed lines in Fig. 2. As shown in the diagram, the causal matrices $\mathbf{A}_{0:K}$ are inferred based on a combination of the current occurrence number $O_t^{v_i, n_j}$, historical occurrence numbers $\mathbf{O}_{t-1:t-\Omega}^{V,N}$ and the topological network $\mathbf{B}_{0:K}$. Note that the causal relationship is invariant across the varying samples. This nature inspires us to implement the inference process with amortized inference (Zhang et al. 2018; Gershman and Goodman 2014), in which a shared function is employed to predict the varying posterior distribution of each sample. Consequently, we consider the generation process function, defined in Eq. (4), as the shared function and designate the causal matrices as the function’s parameter to implement the amortized inference process.

In more detail, to amortize the inference of the shared function, we optimize an inference model to predict the distribution of causal matrices. Given the data $O_t^{v_i, n_j}$ and $\mathbf{O}_{t-1:t-\Omega}^{V,N}$, as well as the topological network $\mathbf{B}_{0:K}$, we use $q_\phi(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K})$ to approximate the true distribution $p_\theta(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K})$. Here, ϕ and θ represent the variational and generation parameters, respectively. The value of θ and ϕ can be learned by maximizing the ensuing likelihood function for each $O_t^{v_i, n_j}$:

$$\begin{aligned} & \log P(O_t^{v_i, n_j} | \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K}) \\ &= D_{KL}(q_\phi^A || p_\theta^A) + \mathcal{L}(\phi, \theta; O_t^{v_i, n_j} | \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K}), \end{aligned} \quad (5)$$

where $q_\phi^A = q_\phi(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K})$, $p_\theta^A = p_\theta(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K})$. In Eq. (5), the right-hand side’s (RHS) first term represents the Kullback-Leibler (KL) divergence of the approximate posterior from the true posterior, and the RHS’s second term is the variational evidence lower bound (ELBO) on the log-likelihood of $O_t^{v_i, n_j}$. By defining the ELBO as \mathcal{L}_e and considering the KL divergence is non-negative, we can deduce from the Eq. (5) that:

$$\begin{aligned} & \log P(O_t^{v_i, n_j} | \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K}) \geq \mathcal{L}_e \\ &= \mathbb{E}_{q_\phi^A}[-\log q_\phi^A + \log p_\theta(O_t^{v_i, n_j}, \mathbf{A}_{0:K} | \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K})]. \end{aligned} \quad (6)$$

The ELBO can be further expressed as:

$$\begin{aligned} \mathcal{L}_e &= \mathbb{E}_{q_\phi^A}[\log p_\theta(O_t^{v_i, n_j} | \mathbf{A}_{0:K}, \mathbf{O}_{t-1:t-\Omega}^{V,N}, \mathbf{B}_{0:K})] \\ &\quad - D_{KL}(q_\phi^A || p_\theta(\mathbf{A}_{0:K})). \end{aligned} \quad (7)$$

Building upon this equation, the goal of the inference task is to optimize \mathcal{L}_e with respect to both ϕ and θ , thereby obtaining the posterior distribution of latent Granger causal structure given the data and the topological network.

Theoretical Analysis

In this subsection, we will theoretically analyze the identifiability of the proposed model. To do so, according to the definition in (Tank et al. 2018), we begin with the definition of the Granger non-causality of multi-type event sequences as follows:

Definition 2 (Granger non-causality of multi-type event sequences). *Given data \mathbf{O}^V during time interval $t \in \{1, \dots, \lceil T/\Delta \rceil\}$, which is generated according to Eq. (2). For the original sample values of historical variables $\mathbf{O}_{t-1:t-\Omega}^{v_1}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_{|V|}}$, let $\hat{\mathbf{O}}_{t-1:t-\Omega}^{v_i}$ represent the different sample values of the historical variables $\mathbf{O}_{t-1:t-\Omega}^{v_i}$ (i.e., $\mathbf{O}_{t-1:t-\Omega}^{v_i} \neq \hat{\mathbf{O}}_{t-1:t-\Omega}^{v_i}$). We can determine the Granger non-causality of event type v_i with respect to event type v_j , if the following condition holds for all t :*

$$\begin{aligned} & g(\mathbf{O}_{t-1:t-\Omega}^{v_1}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_i}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_{|V|}}) \\ &= g(\mathbf{O}_{t-1:t-\Omega}^{v_1}, \dots, \hat{\mathbf{O}}_{t-1:t-\Omega}^{v_i}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_{|V|}}). \end{aligned}$$

That is, $O_t^{v_j}$ is invariant to $\mathbf{O}_{t-1:t-\Omega}^{v_i}$ with $g(\cdot)$.

By extending the above definition to the topological event sequences, we define the notion of Granger non-causality of topological event sequences as follows:

Definition 3 (Granger non-causality of topological event sequences). *Assume that the occurrence numbers of event sequences are generated by Eq. (4). For the original sample values of historical variables $\mathbf{O}_{t-1:t-\Omega}^{v_1, n_1}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_{|V|}, n_{|V|}}$, let $\hat{\mathbf{O}}_{t-1:t-\Omega}^{v_i, n_k}$ represent the different sample values of the historical variables $\mathbf{O}_{t-1:t-\Omega}^{v_i, n_k}$ (i.e., $\hat{\mathbf{O}}_{t-1:t-\Omega}^{v_i, n_k} \neq \mathbf{O}_{t-1:t-\Omega}^{v_i, n_k}$). We can determine the Granger non-causality of event type v_i with respect to event type v_j under topological network \mathcal{G}_N , if the following condition holds for all t :*

$$\begin{aligned} & g(\mathbf{O}_{t-1:t-\Omega}^{v_1, n_1}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_i, n_1}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_i, n_{|N|}}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_{|V|}, n_{|N|}}) \\ &= g(\mathbf{O}_{t-1:t-\Omega}^{v_1, n_1}, \dots, \hat{\mathbf{O}}_{t-1:t-\Omega}^{v_i, n_1}, \dots, \hat{\mathbf{O}}_{t-1:t-\Omega}^{v_i, n_{|N|}}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_{|V|}, n_{|N|}}). \end{aligned}$$

That is, $O_t^{v_j, n_k}$ is invariant to $\mathbf{O}_{t-1:t-\Omega}^{v_i, n_1}, \dots, \mathbf{O}_{t-1:t-\Omega}^{v_i, n_{|N|}}$ with $g(\cdot)$ across all nodes in topological network \mathcal{G}_N .

Definition 3 implies that if two event types are Granger non-causality among a topological network, then they are Granger non-causality across all nodes. Based on this definition, we can derive a proposition about the identification of the model as follows.

Proposition 1. *Given the Granger causal structure $\mathcal{G}_V(\mathbf{V}, \mathbf{E}_V)$, the topological network $\mathcal{G}_N(\mathbf{N}, \mathbf{E}_N)$ and the max geodesic distance K , along with the assumption that the data generation process adheres to Eq. (4), we can deduce that $v_i \rightarrow v_j \notin \mathbf{E}_V$, if and only if $A_k^{v_i, v_j} = 0$ for every $k \in \{0, \dots, K\}$.*

The proof of Proposition 1 is provided in Appendix B. This proposition inspires a methodology for TNPAR to identify Granger causality, which involves evaluating whether the elements of $\mathbf{A}_{0:K}$ are zero.

Algorithm and Implementation

Inference Process of TNPAR

In this subsection, we employ an encoder to implement the above inference process of TNPAR. Here, the $O_t^{v_i, n_j}$ and $\mathbf{O}_{t-1:t-\Omega}^{V, N}$ are used jointly as inputs to the encoder, with $\mathbf{B}_{0:K}$ serving as a mask. The output is the variational posterior of the Granger causal structure, represented as $q_\phi(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{B}_{0:K})$.

To elaborate, each element within $\mathbf{O}_{t-1:t-\Omega}^{V, N}$ is initially filtered in accordance with both $\mathbf{B}_{0:K}$ and n_j . The resulting values are then combined with $O_t^{v_i, n_j}$ and converted into a vector, where each item denotes the representation of the occurrence numbers at a geodesic distance of k for n_j . Following this, the encoder applies a multilayer perceptron (MLP) (Hastie et al. 2009) to the input, which facilitates the propagation of input information across multiple layers and nonlinear activation functions. Ultimately, the encoder generates the posterior distribution of Granger causal structure $q_\phi(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{B}_{0:K})$.

In a manner similar to variational inference for categorical latent variables, we assume that the posterior distribution of each causal edge in the Granger causal structure follows a Bernoulli distribution. More specifically, the encoder produces a latent vector $\mathbf{Z} \in [0, 1]^{|V| \times |V| \times K}$, where $z_{i,j,k}$, the $(i \times j \times k)$ -th element of \mathbf{Z} , represents the parameter of Bernoulli distribution for corresponding edge. Consequently, $q_\phi(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{B}_{0:K})$ can be expressed as:

$$q_\phi(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{B}_{0:K}) = \prod_{i=1}^{|V|} \prod_{j=1}^{|V|} \prod_{k=0}^K q_\phi(A_k^{v_i, v_j} | z_{i,j,k}) \quad (8)$$

$$\text{with } q_\phi(A_k^{v_i, v_j} = 1 | z_{i,j,k}) = \sigma_\beta(z_{i,j,k}).$$

Here, $\sigma_\beta(x) = 1/(1 + \exp(-\beta x))$ denotes the sigmoid function with an inverse temperature parameter β .

Furthermore, for backpropagation through the discrete distribution $q_\phi(\mathbf{A}_{0:K} | O_t^{v_i, n_j}, \mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{B}_{0:K})$, we employ the Gumbel-Softmax trick (Jang, Gu, and Poole 2017) to generate a differentiable sample as:

$$\hat{A}_k^{v_i, v_j} = \frac{\exp(\log(q_\phi(A_k^{v_i, v_j} = 1 | z_{i,j,k}) + \varepsilon_1)/\tau)}{\sum_{e=0}^K \exp(\log(q_\phi(A_k^{v_i, v_j} = e | z_{i,j,k}) + \varepsilon_e)/\tau)}, \quad (9)$$

where each ε_* is an i.i.d. sample drawn from a standard Gumbel distribution, while τ is the temperature parameter controlling the randomness of $\hat{A}_k^{v_i, v_j}$.

Generation Process of TNPAR

Next, we propose a decoder to implement the generation process of $\mathbf{O}_t^{V, N}$. In this part, $\mathbf{O}_{t-1:t-\Omega}^{V, N}$ is filtered according to both $\hat{\mathbf{A}}_{0:K}$ and $\mathbf{B}_{0:K}$, as specified in Eq. (4). This filtered result is then used as input to the decoder. Notably, for neural network training, we treat $\hat{\mathbf{A}}_{0:K}$ as a weight matrix during training, and as a mask during testing. This can be accomplished using soft/hard Gumbel-Softmax sampling. Similar to the encoder, an MLP is used in the decoder.

Aligning with numerous works on counting processes that assume the distribution of occurrence numbers $O_t^{v_i, n_j}$ in each time interval is stationary (Stoyan et al. 2013; Babu and Feigelson 1996), we model each $p_\theta(O_t^{v_i, n_j} | \mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{A}_{0:K}, \mathbf{B}_{0:K})$ using a Poisson process as denoted in Eq. (1). In this context, the conditional distribution of $O_t^{v_i, n_j}$ can be formalized as:

$$p_\theta(O_t^{v_i, n_j} | \mathbf{O}_{t-1:t-\Omega}^{V, N}, \mathbf{A}_{0:K}, \mathbf{B}_{0:K}) = \frac{(\lambda_t^{v_i, n_j} \Delta)^{O_t^{v_i, n_j}}}{O_t^{v_i, n_j}!} e^{-\lambda_t^{v_i, n_j} \Delta}, \quad (10)$$

where $\lambda_t^{v_i, n_j}$, provided by the output of the decoder, represents the intensity parameter of the Poisson process.

Optimization of TNPAR

Based on the above analysis, we adopt the ELBO \mathcal{L}_e , defined in Eq. (7), as the objective function to estimate ϕ and θ and infer the Granger causal structure. To adapt to real-world scenarios, we further introduce acyclic constraint and sparsity constraint into the objective function.

In many real-world applications, the causal structure \mathcal{G}_V is acyclic. For instance, based on experts' prior knowledge, the causal structure in the aforementioned operation and maintenance scenario of a mobile network is a DAG (Directed Acyclic Graph). With this in mind, we introduce an acyclic constraint proposed by (Yu et al. 2019), denoted as:

$$h(\mathbf{G}) \equiv \text{tr}((\mathbf{I} + \frac{1}{|\mathbf{V}|} \mathbf{G})^{|\mathbf{V}|}) - |\mathbf{V}| = 0, \quad (11)$$

where \mathbf{G} is a matrix in which $G_{i,j} = 0$ signifies that v_i has no Granger causality with v_j . Conversely, $G_{i,j} > 0$. The function $\text{tr}(\cdot)$ calculates the trace of an input matrix. Then, a causal structure \mathbf{G} is acyclic if and only if $h(\mathbf{G}) = 0$; Otherwise, $h(\mathbf{G}) > 0$. In this study, we set $G_{i,j} = \sum_{k=0}^K A_k^{v_i, v_j}$ and populate the diagonal of \mathbf{G} with 0, which implies that self-excitation is permitted in an acyclic graph. Subsequently, we design an acyclic regularized term \mathcal{L}_c as:

$$\mathcal{L}_c = \mathbb{E}_{q_\phi}[\text{tr}((\mathbf{I} + \frac{1}{|\mathbf{V}|} \mathbf{G})^{|\mathbf{V}|}) - |\mathbf{V}|] = 0. \quad (12)$$

Because the causal structure tends to be sparse in most real-world applications, we employ a ℓ_1 -norm sparsity constraint as follows:

$$\mathcal{L}_s = \mathbb{E}_{q_\phi}[\sum_{i=1}^{|\mathbf{V}|} \sum_{j=1}^{|\mathbf{V}|} \sum_{k=0}^K A_k^{v_i, v_j}] \leq \kappa, \quad (13)$$

where κ is a small positive constant. Therefore, the training procedure boils down to the following optimization:

$$\max \mathcal{L}_e \quad \text{s.t. } \mathcal{L}_c = 0, \mathcal{L}_s \leq \kappa.$$

By leveraging the Lagrangian multiplier method, we define the total loss function \mathcal{L}_{total} as:

$$\mathcal{L}_{total} = -\mathcal{L}_e + \lambda_c \mathcal{L}_c + \lambda_s \mathcal{L}_s, \quad (14)$$

where λ_c and λ_s denote the regularized hyperparameters.

In summary, our proposed model is trained using the following objective:

$$(\phi^*, \theta^*) = \arg \min_{\phi, \theta} \mathcal{L}_{total}. \quad (15)$$

For model training, we can adopt Adam (Kingma and Ba 2015) stochastic optimization.

Experiment

In this section, we apply the proposed TNPARG method and the baselines to both simulated data and metropolitan cellular network alarm data. For all methods, we employ five different random seeds and present the results in graphical format, with error bars included. Our evaluation metrics for the experiments include Precision, Recall, F1 score (Powers 2011), Structural Hamming Distance (SHD), and Structural Intervention Distance (SID) (Peters and Bühlmann 2015). Specifically, Precision refers to the fraction of predicted edges that exist among the true edges. Recall is the fraction of true edges that have been successfully predicted. F1 score is the weighted harmonic mean of both Precision and Recall, and is calculated as $F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$. SHD represents the number of edge insertions, deletions, or flips needed to transform a graph into another graph. SID is a measure that quantifies the closeness between two DAGs based on their corresponding causal inference statements.

Dataset

Simulated Data The simulated data is generated as follows: a) a Granger causal structure \mathcal{G}_V and a topological network \mathcal{G}_N are randomly generated; b) root event records are generated via the Poisson process using a base intensity parameter μ in the Hawkes process. These root event records are spontaneously generated by the system; c) based on the root event records, propagated event records are discretely generated according to both the time interval Δ and the excitation intensity α . Here, α represents the event excitation intensity in the Hawkes process. Given that the event sequences can be quite sparse in real-world data, we offer a time interval parameter Δ to the generation process, dividing the time domain $[0, T]$ to small intervals with indexes as $\{1, \dots, \lceil T/\Delta \rceil\}$. Then, the event records can be summarized within the same timestamp. Note that the $\Delta \geq 0$. If $\Delta = 0$, it implies the use of the original event sequences.

All simulated data are generated by varying the parameters of the generation process one at a time while maintaining default parameters. We set these default parameters according to the data generation process used in the PCIC 2021 competition¹, inclusive of the additional parameter Δ , as follows: #nodes = 40, #event type = 20, #sample size = 20000, $\mu = [0.00003, 0.00005]$, $\alpha = [0.02, 0.03]$, $\Delta = 2$.

Metropolitan Cellular Network Alarm Data This real-world dataset is collected in a business scenario by a multinational communications company and is available in the PCIC 2021 competition. The data comprises a series of alarm records generated within a week according to both a topological network \mathcal{G}_N and a causal structure \mathcal{G}_V . Specifically, the topological network \mathcal{G}_N includes 3087 network elements. The alarm records encompass 18 alarm event types, totaling 228,030 alarm event records. It is noteworthy that, due to the characteristics of the equipment and as confirmed by experts, the collected timestamps of alarm events exhibit certain time intervals.

¹<https://competition.huaweicloud.com/information/1000041487/dataset>

Baselines and Model Variants

The following causal discovery methods are chosen as baseline comparisons in our experiments: PCMCI (Runge 2020), ACD (Löwe et al. 2022), MLE_SGL (Xu, Farajtabar, and Zha 2016), ADM4 (Zhou, Zha, and Song 2013), and THP (Cai et al. 2022). In order to assess the components of our proposed method, we also introduce three variants of our method, namely TNPARG_NT, TNPARG_NC, and TNPARG_MG. Specifically, TNPARG_NT does not incorporate topological information and treats each event sequence independently. TNPARG_MG also does not introduce the topological network and takes the event sequences in all nodes as one single sequence. TNPARG_NC is a variant that excludes both acyclicity and sparsity regularization from the loss function. Detailed information about the baselines and variants can be found in Appendix C.

Results on Simulated Data

The F1 scores of different methods on the simulated data are depicted in Fig. 3. Due to space limitations, results pertaining to Precision, Recall, SHD, and SID on the simulated data can be found in Appendix D.

As shown in Fig. 3, when compared with the results of PCMCI, ACD, ADM4, and MLE_SGL (which do not consider the topological network behind the event sequences), both TNPARG and THP achieve superior performance across all cases. These results suggest that incorporating the topological network assists in learning Granger causality. Moreover, TNPARG and its variants significantly outstrip other baselines under most cases, demonstrating that the neural point process approach is powerful for modeling the generation process and the amortized inference technique is effective for Granger causal discovery.

Traversing the value of parameters α , μ , and Δ , the F1 scores of the TNPARG are generally better than those of other methods. These results signify that our method is relatively insensitive to the parameters of the generation process. Under settings with varying sample sizes, or changing numbers of event types or nodes, TNPARG achieves the highest F1 scores, illustrating the robustness of TNPARG. Even in cases with small sample sizes or high dimensional event types, TNPARG still outperforms other methods, corroborating the effectiveness of TNPARG.

Among the baseline methods, THP delivers the best results as it considers the topological network. Both ADM4 and MLE_SGL outperform PCMCI and ACD, which demonstrates the efficacy of the point process methods in causal discovery for event sequences.

Results on Metropolitan Cellular Network Alarm Data

The experimental results of all algorithms on real-world data are presented in Table 1. From these results, it can be seen that the models incorporated the topological network (including TNPARG, TNPARG_NC, THP) outperform that those rely on the i.i.d. assumption. Specifically, the F1 score of TNPARG is higher than that of THP, demonstrating

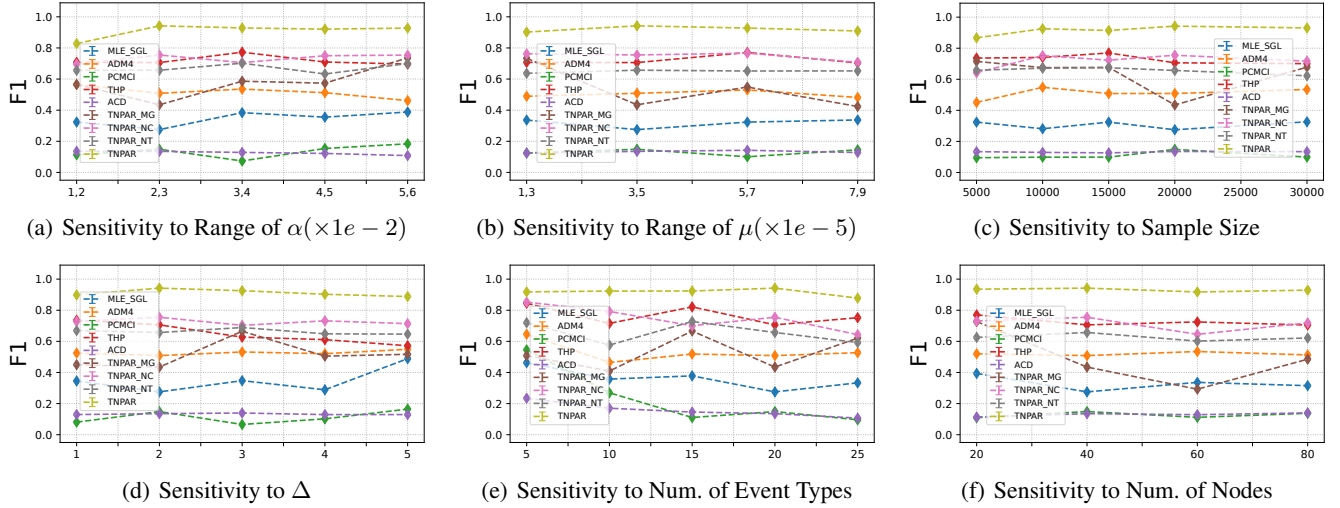


Figure 3: Results on the simulated data

Algorithm	Prec.	Rec.	F1	SID	SHD
ACD	0.25	0.32	0.27	230.0	127.0
MLE_SGL	0.39	0.12	0.18	263.0	100.0
ADM4	0.49	0.34	0.40	292.0	95.0
PCMCI	0.47	0.31	0.37	306.0	102.0
THP	0.58	0.35	0.44	266.0	85.0
TNPAR_NT	0.31	0.48	0.38	257.0	150.4
TNPAR_NC	0.41	0.60	0.49	271.3	123.5
TNPAR_MG	0.35	0.31	0.33	272.5	115.2
TNPAR	0.46	0.65	0.54	241.6	105.5

Table 1: Result on real-world data

that TNPAR can achieve the best trade-off between Precision and Recall without needing a prior distribution assumption. Although the Precision of TNPAR is marginally lower than that of THP, PCMCI, and ADM4, TNPAR exhibits the highest Recall. These results indicate that TNPAR is relatively insensitive to the weak Granger causal strength between event types, which accounts for the slightly higher SHD of TNPAR compared to other baselines. Moreover, TNPAR attains the best SID relative to other methods except ACD. These results underscore TNPAR’s superior capability for causal inference (Peters and Bühlmann 2015).

Ablation Study

Ablation experiments were conducted on both simulated and real-world data to assess the effectiveness of incorporating topological information and regularization in TNPAR. The experimental results given in Fig. 3 reveal a significant improvement in the performance of TNPAR compared to TNPAR_NT, TNPAR_MG, and TNPAR_NC. These results underscore the importance of incorporating topological information, as well as acyclicity and sparsity regularization. Additionally, in most instances, TNPAR_NT, TNPAR_MG, and TNPAR_NC perform comparably to THP, despite the

fact that the THP fully leverages topological information and sparsity constraint. These results demonstrate the robustness of TNPAR.

Case Study

In the real-world data experiment, our method successfully infers the Granger causality across the nodes of the topological network, aligning with experts’ knowledge. From the results, we observe that the count of causal edges diminishes as the geodesic distance, denoted as k , increases. This finding aligns with our intuition that some events are influenced solely by their node. For instance, the causal relationship $BD_STATUS \rightarrow TU_AIS$ is only valid for $k = 0$. Here, BD_STATUS represents a device disconnecting from a node, and TU_AIS signifies that some tasks running on this node have been disrupted. Another notable discovery is that some edges exclusively transit across nodes, suggesting that these event types are triggered only by their neighbors. For example, consider $ETH_LINK_DOWN \rightarrow MW_LOF$, where ETH_LINK_DOWN denotes an error on the Ethernet interface in a downstream node, and MW_LOF indicates disconnection of an upstream node.

Conclusion

In this paper, we studied the Granger causal discovery problem on the topological event sequences. By leveraging the prior topological network and regarding the causal structure as a latent variable, we successfully addressed the challenge of unified modeling the topological network and latent causal structure, which enable us to propose the amortized causal discovery method. The experimental results on both simulation and real-world data demonstrate the effectiveness of our proposed method. TNPAR offers a viable solution for causal discovery in real-world settings with non-i.i.d. data. To further enhance its effectiveness, our future work will focus on improving sample efficiency, integrating prior knowledge, and so on.

Acknowledgments

This research was supported in part by National Key R&D Program of China (2021ZD0111501), National Science Fund for Excellent Young Scholars (62122022), Natural Science Foundation of China (61876043, 61976052, 62206061, 62206064), the major key project of PCL (PCL2021A12), Guangzhou Basic and Applied Basic Research Foundation (2023A04J1700).

References

- Babu, G. J.; and Feigelson, E. D. 1996. Spatial point processes in astronomy. *Journal of Statistical Planning and Inference*, 50(3): 311–326.
- Bhattacharjya, D.; Subramanian, D.; and Gao, T. 2018. Proximal graphical event models. *Advances in Neural Information Processing Systems*, 31.
- Bouttier, J.; Di Francesco, P.; and Guitter, E. 2003. Geodesic distance in planar graphs. *Nuclear physics B*, 663(3): 535–567.
- Brillinger, D. R. 1994. Time series, point processes, and hybrids. *Canadian Journal of Statistics*, 22(2): 177–206.
- 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*.
- Chen, W.; Cai, R.; Hao, Z.; Yuan, C.; and Xie, F. 2020. Mining hidden non-redundant causal relationships in online social networks. *Neural Computing and Applications*, 32(11): 6913–6923.
- Cox, D. R. 1955. Some statistical methods connected with series of events. *Journal of the Royal Statistical Society: Series B (Methodological)*, 17(2): 129–157.
- Ertekin, Ş.; Rudin, C.; and McCormick, T. H. 2015. Reactive point processes: A new approach to predicting power failures in underground electrical systems. *The Annals of Applied Statistics*, 9(1): 122–144.
- Gershman, S.; and Goodman, N. 2014. Amortized inference in probabilistic reasoning. In *Annual Meeting of the Cognitive Science Society*, volume 36.
- Granger, C. W. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 424–438.
- Hastie, T.; Tibshirani, R.; Friedman, J. H.; and Friedman, J. H. 2009. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer.
- Hawkes, A. G. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1): 83–90.
- Isham, V.; and Westcott, M. 1979. A self-correcting point process. *Stochastic processes and their applications*, 8(3): 335–347.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Li, Z.; Cai, R.; Fu, T. Z.; Hao, Z.; and Zhang, K. 2023. Transferable time-series forecasting under causal conditional shift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Löwe, S.; Madras, D.; Zemel, R.; and Welling, M. 2022. Amortized causal discovery: Learning to infer causal graphs from time-series data. In *Conference on Causal Learning and Reasoning*, 509–525. PMLR.
- Mijatovic, G.; Antonacci, Y.; Loncar-Turukalo, T.; Minati, L.; and Faes, L. 2021. An information-theoretic framework to measure the dynamic interaction between neural spike trains. *IEEE Transactions on Biomedical Engineering*, 68(12): 3471–3481.
- Peters, J.; and Bühlmann, P. 2015. Structural intervention distance for evaluating causal graphs. *Neural Computation*, 27(3): 771–799.
- Powers, D. 2011. Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness & Correlation. *Journal of Machine Learning Technologies*, 2(1): 37–63.
- Runge, J. 2020. Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series datasets. In *Conference on Uncertainty in Artificial Intelligence*, 1388–1397.
- Shchur, O.; Türkmen, A. C.; Januschowski, T.; and Günnemann, S. 2021. Neural Temporal Point Processes: A Review. In *International Joint Conference on Artificial Intelligence*, 4585–4593.
- Shojaie, A.; and Fox, E. B. 2022. Granger causality: A review and recent advances. *Annual Review of Statistics and Its Application*, 9: 289–319.
- Stoyan, D.; Kendall, W. S.; Chiu, S. N.; and Mecke, J. 2013. *Stochastic geometry and its applications*. John Wiley & Sons.
- Tank, A.; Covert, I.; Foti, N.; Shojaie, A.; and Fox, E. 2018. Neural granger causality for nonlinear time series. *Stat*, 1050: 16.
- Vuković, M.; and Thalmann, S. 2022. Causal Discovery in Manufacturing: A Structured Literature Review. *Journal of Manufacturing and Materials Processing*, 6(1): 10.
- Xiao, S.; Yan, J.; Farajtabar, M.; Song, L.; Yang, X.; and Zha, H. 2019. Learning time series associated event sequences with recurrent point process networks. *IEEE Transactions on Neural Networks and Learning Systems*, 30(10): 3124–3136.
- Xu, H.; Farajtabar, M.; and Zha, H. 2016. Learning granger causality for hawkes processes. In *International Conference on Machine Learning*, 1717–1726.
- Yu, Y.; Chen, J.; Gao, T.; and Yu, M. 2019. DAG-GNN: DAG structure learning with graph neural networks. In *International Conference on Machine Learning*, 7154–7163.
- Zhang, C.; Bütepage, J.; Kjellström, H.; and Mandt, S. 2018. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8): 2008–2026.

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.

Zhou, K.; Zha, H.; and Song, L. 2013. Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes. In *Artificial Intelligence and Statistics*, 641–649.