
Event Outlier Detection in Continuous Time

Siqi Liu^{1,2} Milos Hauskrecht¹

Abstract

Continuous-time event sequences represent discrete events occurring in continuous time. Such sequences arise frequently in real-life. Usually we expect the sequences to follow some regular pattern over time. However, sometimes these patterns may be interrupted by unexpected absence or occurrences of events. Identification of these unexpected cases can be very important as they may point to abnormal situations that need human attention. In this work, we study and develop methods for detecting outliers in continuous-time event sequences, including unexpected absence and unexpected occurrences of events. Since the patterns that event sequences tend to follow may change in different contexts, we develop outlier detection methods based on point processes that can take context information into account. Our methods are based on Bayesian decision theory and hypothesis testing with theoretical guarantees. To test the performance of the methods, we conduct experiments on both synthetic data and real-world clinical data and show the effectiveness of the proposed methods.

1. Introduction

Continuous-time event sequences are defined by occurrences of various types of events in time. Event sequences may represent many real-world processes and observations including, e.g., arrival of packets to servers in network systems, or administration of drugs to patients. Continuous-time event sequences are typically modeled as point processes (Daley & Vere-Jones, 2003). While many point-process models have been developed in the literature in recent years, most works build and test models on event prediction tasks. The focus of this work is on outlier detection in event sequences that aims to identify unusual occurrences

¹Department of Computer Science, University of Pittsburgh, Pittsburgh, PA, USA ²Borealis AI, Vancouver, BC, Canada. Correspondence to: Siqi Liu <siqiliu@cs.pitt.edu>, Milos Hauskrecht <milos@pitt.edu>.

or absence of events in event sequences in real-time. Outlier detection is the basis of many critical real-world applications, such as fraud detection (Fawcett & Provost, 1997), network intrusion surveillance (Garcia-Teodoro et al., 2009), disease outbreak detection (Wong et al., 2003), and medical error detection (Hauskrecht et al., 2013; 2016).

Two types of outliers may arise in continuous-time event sequences. First, given the history and the recent absence of the events, the event may be overdue. We refer to these as *omission outliers*. Second, given the history, the event that has just arrived is unexpected: it either arrived too early or was not expected to occur at all given the context from the history. We refer to these as *commission outliers*. Figure 1 shows some examples of these two types of outliers.




Figure 1. Examples of omission (top) and commission (bottom) outliers. Each stem and rectangle / circle represent one event on the horizontal time line. Different colors and shapes represent different types of events. Red dashed boxes mark the outliers. Top: A blue-circle event is expected but has not been observed. Bottom: The second to last blue-circle event happened too early. The last blue-circle event was not expected given the occurrence pattern.

Both types of outliers are often related to problems of practical importance. Take for example, a person suffering from a disease and taking specific medications on a regular schedule to treat the disease. Given the history and the current time, we may infer that the person has not taken the medication yet, and the medication is overdue (omission). The detection of the overdue medication can be then used to generate a reminder alert. For commission, consider a patient who takes a medication too early compared to the normal

schedule. The detection of this event or its prevention is extremely important and may prevent adverse situations like high concentration of the drug and its possible toxic effects. Similar situations may happen when one receives a medication that is unrelated to his/her condition (i.e., context). The occurrence of this event may indicate a medical error, and again its timely detection is extremely important.

Despite the importance of these types of outlier detection problems, to the best of our knowledge, neither of them have been studied in the literature. The key contributions of this paper are the formulation of these outlier detection problems and the development of their algorithmic solutions. Specifically, we develop semi-supervised outlier detection methods (Chandola et al., 2009), where data generated by the normal point process are available, based on Bayesian decision theory and hypothesis testing that can leverage a variety of point-process models. We show the effectiveness of the methods on both synthetic and real-world clinical event sequences. Since event occurrences often depend on other related events, we adapt an existing model, the continuous-time LSTM (Mei & Eisner, 2017), to account for these events and their history as context. However, we stress that the proposed outlier detection methods can be combined with any type of point-process model.

2. Related Work

Point Processes Point processes are probabilistic models for discrete points in continuous domains (Daley & Vere-Jones, 2003). They have been widely used to model continuous-time event sequences. For a temporal point process, the conditional intensity function (CIF) characterizes the process. Some researchers have proposed to use Gaussian processes to model the CIF (Adams et al., 2009; Rao & Teh, 2011; Lloyd et al., 2015; 2016; Ding et al., 2018; Liu & Hauskrecht, 2019), while the others have developed flexible models based on Hawkes processes (Hawkes, 1971; Zhou et al., 2013; Wang et al., 2016; Xu et al., 2016; Lee et al., 2016; Apostolopoulou et al., 2019), where the CIF depends on the history explicitly. Recently, neural-network-based point-process models (Du et al., 2016; Mei & Eisner, 2017; Xiao et al., 2017; Li et al., 2018; Jia & Benson, 2020) have shown promising results in different settings. Despite the large amount of works on point processes for event sequences, none of them have studied the problem of defining and detecting outliers in continuous time.

Point Processes for Noisy Data Recently, researchers have started developing methods to deal with noisy data, such as incomplete or missing data (Xu et al., 2017; Shelton et al., 2018; Mei et al., 2019) and desynchronized data (Trouleau et al., 2019). They assume the data has been corrupted by some source (e.g., censoring or noise), and

the goal is to recover the original data or to learn a model nonetheless. Although the omission outliers we consider is related to missing data, our goal is to detect these outliers as accurately as possible, i.e., to distinguish them from normal data, instead of data recovery. Moreover, our method is assumed to be executed in an online manner, so we must decide whether there is an outlier based *only* on the history. We do not have access to the data in the future as in the missing data setting. Finally, detecting commission outliers is not related to any of these works.

Outlier Detection In general, outlier (or anomaly) detection (Chandola et al., 2009; Aggarwal, 2013) aims to identify data instances that are unusual when compared to other instances in data. When outlier detection is applied on a subset of dimensions given the rest as context, it is referred to as contextual (or conditional) outlier detection (Hauskrecht et al., 2007; Song et al., 2007; Chandola et al., 2009). Our goal in this work is to detect unusual absence and occurrences of events in event sequences that depend on other types of events defining the context. Unlike outlier detection in traditional time series (Gupta et al., 2013), where time is treated as merely the index of the observed values, for continuous-time event sequences, the time of each event is essentially a random variable defined by the generative process, and the goal is to detect unusual behaviors in time given the context.

Semi-Supervised Detection Semi-supervised outlier (or anomaly) detection is a term coined in (Chandola et al., 2009). It means that normal data (without any outliers or other corruptions) are available for training the model to detect outliers on test data later. It is a widely-adopted approach in both outlier detection (Fiore et al., 2013; Akcay et al., 2019) and other related tasks, such as intrusion detection (Xu & Shelton, 2010) and missing data imputation (Mei et al., 2019), although in some cases, the assumption is not clearly stated. Our work also adopts this approach, but we note that event outlier detection in continuous time has not been studied in any of the previous works.

3. Method

3.1. Problem Formulation

First, we formally define the problem of contextual outlier detection in continuous-time event sequences. An event sequence can be formulated as $S_x = \{(t_i : t_i \in \mathcal{T})_{i=1}^{N_x}\}$, i.e., a sequence of timestamps t_i of the events, where t_i is the time of the i -th event, N_x is the total number of the events in the sequence, and $\mathcal{T} \subseteq \mathbb{R}$ is the domain of time. We call S_x the *target sequence* and the events the *target events*, because they are the targets in which we aim to detect outliers. Meanwhile, we may observe contextual information

along with S_x . We assume the contextual information can be either represented as or converted to discrete events. We denote these events as $S_C = \{(t_i, u_i) : t_i \in \mathcal{T}, u_i \in \mathcal{C}\}_{i=1}^{N_C}$, where t_i and u_i are the time and type (mark) of the i -th event, N_C is the total number of the events, and $\mathcal{C} \subseteq \mathbb{Z}$ is the finite set of distinct marks for different types of events. We call S_C the *context sequence* and the events the *context events*.

We stress that S_x and S_C share the same time domain \mathcal{T} and, therefore, we can combine them into a single sequence $S_{\mathcal{M}} = \{(t_i, u_i) : (t_i, u_i) \in S_C \text{ or } t_i \in S_x, u_i = x\}$, where a new type $x \notin \mathcal{C}$ is assigned to all the events in S_x . For detecting outliers, we only rely on information in the *past* from both S_x and S_C . We denote the combined history of S_x and S_C up to time t as $\mathcal{H}_t^{\mathcal{M}} = \{(t_i, u_i) : (t_i, u_i) \in S_{\mathcal{M}}, t_i < t\}$. Meanwhile, $\mathcal{H}_t^x = \{t_i : t_i \in S_x, t_i < t\}$ is the history of the target sequence S_x only without any contextual information. Now, we are ready to define two types of outlier detection problems we wish to solve.

Commission Outlier Detection Given an observed target event at time t_n and the combined history $\mathcal{H}_{t_n}^{\mathcal{M}}$ up to time t_n , the goal is to assign a label $y_c(t_n) \in \{0, 1\}$ to t_n indicating whether it is a commission outlier. Notice that $y_c(t)$ is only defined if t is the time of a target event. In this work, instead of hard labels, we calculate a commission outlier score $s_c(t_n)$ for t_n to indicate how likely it is to be a commission outlier.

Omission Outlier Detection Define a blank interval $B \subseteq \mathcal{T}$ as an interval in which there is no event of type x (target event). Given a blank interval $B = (t_b, t_e)$ and the combined history $\mathcal{H}_{t_b}^{\mathcal{M}}$ up to time t_b , the goal is to assign a label $y_o(B) \in \{0, 1\}$ to B indicating whether there are any omission outliers in B . Notice that $y_o(B)$ is only defined when B is an interval with no events from S_x . In this work, instead of hard labels, we calculate an omission outlier score $s_o(B)$ for B to indicate how likely it is to contain any omission outliers.

3.2. Probabilistic Models

We develop algorithms for detecting outliers in continuous-time event sequences based on probabilistic models, specifically (temporal) point processes. Point processes are probabilistic models for discrete points in continuous domains. For an event sequence, the points are the events, and the domain is the time \mathcal{T} . In this case, the models are also called *temporal* point processes. A temporal point process can be defined as a counting process $N(\cdot)$ on \mathcal{T} , where $N(\tau)$ is the number of points in the interval $\tau \subseteq \mathcal{T}$.

For a temporal point process, the conditional intensity function (CIF), $\lambda(t)$, characterizes the probability of observing

an event in an infinitesimal time interval $[t, t + dt) \subseteq \mathcal{T}$ given the history up to time t , i.e. $\lambda(t)dt = p(N([t, t + dt)) = 1 | \mathcal{H}_t)$. For our problem, we only model the *target* events using a point process, while the history $\mathcal{H}_t = \mathcal{H}_t^{\mathcal{M}}$ contains both the target events and the context events. Because we always condition on \mathcal{H}_t , we omit \mathcal{H}_t in notations for the rest of the paper.

For a sequence of target events $S_x = \{t_i : t_i \in \mathcal{T}\}_{i=1}^{N_x}$ generated from the point process with CIF $\lambda(t)$, the probability density is

$$p(S_x) = \prod_{i=1}^{N_x} \lambda(t_i) \exp\left(-\int_{\mathcal{T}} \lambda(s)ds\right). \quad (1)$$

In this work, we focus on semi-supervised outlier detection (Chandola et al., 2009), meaning that we can obtain data generated from the normal point process (without outliers), from which we may train a model for detecting outliers on unseen data later. It has been widely adopted in previous works (see Section 2).

Assumption 3.1. *Data generated by the normal point process are available.*

We note that by definition, the majority of the data are normal rather than outliers, so obtaining normal data is usually much easier than finding and labeling outliers. Meanwhile, robust learning of the model in presence of outliers is an important but orthogonal direction for future work, since we can combine robust learning with our proposed outlier detection methods to eliminate Assumption 3.1.

In general, the type of model to use should be able to represent the dependencies between the target events and the context events. For a point-process model, it means that the CIF, instead of only depending on the target events, $\lambda(t) = f(\mathcal{H}_t^x)$, should also depend on the context events, $\lambda(t) = f(\mathcal{H}_t^{\mathcal{M}})$, where $f(\cdot)$ denotes the mapping represented by the model.

In this work, we use a model adapted from the continuous-time LSTM (Mei & Eisner, 2017) due to its good performance in previous studies (Mei & Eisner, 2017; Jia & Benson, 2020). See Appendix A for details of the model. However, we stress that if there is a better model, we can simply plug it in to get better performance.

3.3. Detecting Commission Outliers

Without knowing the true process that generates commission outliers, we make the following assumption when developing our method and later generalize it by relaxing the assumption of constant rate.

Assumption 3.2 (Independent and Constant Commission). *Commission outliers are generated by a process with a constant CIF independently from the normal point process.*

Suppose we are given a target event t_n and the history up to time t_n . Define a random variable Z_n , such that $Z_n = 1$ if t_n is a commission outlier, and $Z_n = 0$ otherwise. We are interested in calculating $p(Z_n = 1|t_n)$.

By Assumption 3.2 (Independent), the generative processes for the normal points and outliers can be viewed together as a marked point process. For each event t_n , there is a hidden mark Z_n indicating whether it is an outlier. The overall CIF is $\lambda_g(t) = \lambda_1(t) + \lambda_0(t)$, where $\lambda_0(t)$ is the CIF for the normal point process, and $\lambda_1(t)$ for the outlier process. Suppose we are at time t_c . The density of having the next event at t_n marginally and jointly with mark $Z_n = 1$ are respectively

$$p(t_n) = \lambda_g(t_n) \exp\left(-\int_{t_c}^{t_n} \lambda_g(t) dt\right), \quad (2)$$

$$p(Z_n = 1, t_n) = \lambda_1(t_n) \exp\left(-\int_{t_c}^{t_n} \lambda_g(t) dt\right). \quad (3)$$

Then we can derive the conditional distribution of the hidden mark

$$p(Z_n = 1|t_n) = \frac{\lambda_1(t_n)}{\lambda_g(t_n)} = 1 - \frac{\lambda_0(t_n)}{\lambda_g(t_n)}. \quad (4)$$

Therefore, the Bayes decision rule is

$$Z_n^* = \arg \max_{z \in \{0,1\}} p(Z_n = z|t_n) = \mathbb{I}[\lambda_1(t_n) > \lambda_0(t_n)] \quad (5)$$

where $\mathbb{I}[x] = 1$ if x is true, and 0 otherwise. However, this rule cannot be directly applied, because $\lambda_1(t_n)$ is unknown. From Assumption 3.2 (Constant), λ_1 is a constant, so the decision rule becomes $Z_n^* = \mathbb{I}[\lambda_0(t_n) < \theta_c]$, where θ_c is a threshold. This justifies ranking by

$$s_c(t_n) = -\lambda_0(t_n) \quad (6)$$

across all $n = 1, \dots, N_x$, so we use $-\lambda_0(t_n)$ as the commission outlier score: the higher the score, the more likely t_n is to be a commission outlier.

Nonconstant Rate We generalize the idea by relaxing the assumption of $\lambda_1(t)$ being a constant. The key is to treat $\lambda_1(t_n)$ as a random variable and then marginalize it out as

$$\begin{aligned} p(Z_n = 1|t_n) &= \mathbb{E}_{\lambda_1} [p(Z_n = 1|t_n, \lambda_1(t_n))] \\ &= 1 + \mathbb{E}_{\lambda_1} \left[\frac{s_c(t_n)}{\lambda_1(t_n) - s_c(t_n)} \right]. \end{aligned} \quad (7)$$

Then we can prove that $p(Z_n = 1|t_n)$ is still an increasing function of $s_c(t_n)$, which justifies ranking by s_c (see Appendix B for details).

3.4. Detecting Omission Outliers

Similar to detecting commission outliers, we make the following assumption for omission outliers.

Assumption 3.3 (Independent and Constant Omission). *Omission outliers are generated by independently removing the normal points with a constant probability.*

To derive the method, we first define some notations. Let $\lambda_0(t)$ denote the CIF of the normal point process, p_1 the probability of removing each normal point, and B a blank interval without any target points which we wish to decide whether contains omission outliers. For any interval $\tau \subseteq \mathcal{T}$, let $N(\tau)$ be the number of points observed, and $N_0(\tau)$ be the number of points generated by the normal point process with CIF $\lambda_0(t)$, so $N(\cdot)$ is the result of combining $N_0(\cdot)$ with random removal, and we can observe $N(\cdot)$ but not $N_0(\cdot)$. Furthermore, we define an auxiliary random variable K_B that counts the number of points removed in B .

For a blank interval B , we observe $N(B) = 0$, but $K_B = k$ can take different values $k = 0, 1, \dots$. The joint probability for each k is

$$\begin{aligned} p(K_B = k, N(B) = 0) &= p(K_B = k, N_0(B) = k) \\ &= p_1^k F_k(B) \end{aligned} \quad (8)$$

where $F_k(B)$ denotes the probability that k points are generated by the normal point process $N_0(\cdot)$ in B for $k = 0, 1, \dots$. It depends on $\lambda_0(t)$. Then we can calculate the posterior probability of $K_B = 0$

$$p(K_B = 0|N(B) = 0) = \frac{F_0(B)}{\sum_{k=0}^{\infty} p_1^k F_k(B)} \quad (9)$$

Define a random variable Z_B to indicate whether there are any omission outliers in the blank interval B . $Z_B = 0$ is equivalent to $K_B = 0$; $Z_B = 1$ is equivalent to $K_B > 0$.

$$\begin{aligned} p(Z_B = 1|N(B) = 0) &= 1 - p(Z_B = 0|N(B) = 0) \\ &= 1 - p(K_B = 0|N(B) = 0). \end{aligned} \quad (10)$$

Then the Bayes decision rule is

$$\begin{aligned} Z_B^* &= \arg \max_{z \in \{0,1\}} p(Z_B = z|N(B) = 0) \\ &= \mathbb{I}[p(K_B = 0|N(B) = 0) < 0.5]. \end{aligned} \quad (11)$$

Without further assumptions, $p(K_B = 0|N(B) = 0)$ (Eq. 9) cannot be evaluated in closed form, but we can get a lower bound

$$\begin{aligned} p(K_B = 0|N(B) = 0) &\geq F_0(B) \\ &= \exp\left(-\int_B \lambda_0(s) ds\right), \end{aligned} \quad (12)$$

because

$$\sum_{k=0}^{\infty} p_1^k F_k(B) \leq \sum_{k=0}^{\infty} F_k(B) = 1.$$

Then the posterior probability of B containing any omission outliers can also be bounded

$$p(Z_B = 1 | N(B) = 0) \leq 1 - \exp\left(-\int_B \lambda_0(s) ds\right). \quad (13)$$

Therefore, we propose to use

$$s_o(B) = \int_B \lambda_0(s) ds \quad (14)$$

as the omission outlier score. When we rank the blank intervals by $s_o(B)$, we essentially rank them by an upper bound of $p(Z_B = 1 | N(B) = 0)$.

Poisson Process There is a notable special case where we can get a closed-form $p(K_B = 0 | N(B) = 0)$. When the normal point process $N_0(\cdot)$ is an *inhomogeneous Poisson process*, we have

$$F_k(B) = \frac{\left(\int_B \lambda_0(s) ds\right)^k}{k!} \exp\left(-\int_B \lambda_0(s) ds\right) \quad (15)$$

for $k = 0, 1, \dots$. The posterior becomes

$$\begin{aligned} p(K_B = 0 | N(B) = 0) &= \frac{F_0(B)}{\sum_{k=0}^{\infty} p_1^k F_k(B)} \\ &= \exp\left(-p_1 \int_B \lambda_0(s) ds\right). \end{aligned} \quad (16)$$

Therefore, the posterior probability of B containing any omission outliers is

$$p(Z_B = 1 | N(B) = 0) = 1 - \exp\left(-p_1 \int_B \lambda_0(s) ds\right). \quad (17)$$

This justifies scoring the interval B by $s_o(B) = \int_B \lambda_0(s) ds$, because if we rank the intervals by their scores, the result will be the same as ranking by their posterior probabilities of containing omission outliers, $p(Z_B = 1 | N(B) = 0)$.

Nonconstant Rate Similar to commission outliers, we can generalize the idea by relaxing the assumption that p_1 is a constant and treating it as a random variable, in which case $p(Z_B = 1 | N(B) = 0)$ is still an increasing function of $s_o(B)$ (see Appendix C for details).

Inter-Event Time Without making any assumptions on the normal point process, we can alternatively justify $s_o(B) = \int_B \lambda_0(s) ds$ with hypothesis testing, if B is an *inter-event time* interval, i.e., a time interval between consecutive events t_{n-1} and t_n . Let T_n be the random variable

for the inter-event time corresponding to B . The null and alternative hypotheses are

$$H_0 : B \text{ is normal}; \quad H_1 : B \text{ contains omission outliers.}$$

Assuming the null hypothesis is true, i.e., B is generated by the normal point process with CIF $\lambda_0(t)$, the probability that the inter-event time is at least as long as $|B|$ is

$$p(T_n > |B|) = \exp\left(-\int_B \lambda_0(s) ds\right), \quad (18)$$

which is the p-value. A lower p-value means that the observation is more extreme, given that the null hypothesis is true, which means it is more likely to contain omission outliers. This justifies scoring by $s_o(B) = \int_B \lambda_0(s) ds$, where a higher score means that B is more likely to contain omission outliers.

3.5. Bounds on FDR and FPR

In this section, we prove some bounds on the performance of the proposed outlier scoring methods. We recall the definitions of false discovery rate (FDR) and false positive rate (FPR). Let y denote the true label (1=outlier, 0=normal) of an object (a target event or a blank interval) and \hat{y} denote the predicted label. Then FDR and FPR are defined as

$$FDR = p(y = 0 | \hat{y} = 1), \quad FPR = p(\hat{y} = 1 | y = 0).$$

Given the above definitions, we can prove the following theorems (see Appendix D for proofs).

Theorem 3.1. *If we use the commission outlier score $s_c(t_n) = -\lambda_0(t_n)$, where t_n is the time of a target event, with a threshold $\theta_c \leq 0$, such that the decision rule is $\hat{y}_c(t_n) = \mathbb{I}[s_c(t_n) > \theta_c]$, and let λ_1 denote the CIF of the independent process generating commission outliers, then the FDR is bounded above by $\frac{-\theta_c}{\lambda_1 - \theta_c}$.*

Theorem 3.2. *If we use the omission outlier score $s_o(B) = \int_B \lambda_0(s) ds$ for an inter-event time interval B , with a threshold $\theta_o \geq 0$, such that the decision rule is $\hat{y}_o(B) = \mathbb{I}[s_o(B) > \theta_o]$, then the FPR is bounded above by $\exp(-\theta_o)$.*

Theorem 3.3. *If we use the omission outlier score $s_o(B) = \int_B \lambda_0(s) ds$ for a blank interval B , with a threshold $\theta_o \geq 0$, such that the decision rule is $\hat{y}_o(B) = \mathbb{I}[s_o(B) > \theta_o]$, and assume that the normal point process is an inhomogeneous Poisson process and the probability of omission is p_1 , then the FDR is bounded above by $\exp(-p_1 \theta_o)$.*

4. Experiments

To test the proposed methods, we perform experiments on both synthetic and real-world event sequences. We compare the following methods in the experiments. **RND**: A baseline

that generates outlier scores by sampling from a uniform distribution. **LEN**: A baseline that detects outliers based on the empirical distribution of the inter-event time lengths (see Section 4.1). **PPOD** (Point-Process Outlier Detection): Our method based on a point-process model but only using the history of the target events as the context. **CPPOD** (Contextual PPOD): Our method based on a point-process model using the history of both the target events and the context events as the context.

A model adapted from the continuous-time LSTM (Mei & Eisner, 2017) is used for PPOD and CPPOD (see Appendix A). We choose the number of hidden units in the model from $\{64, 128, 256, 512, 1024\}$ by maximizing the likelihood on the internal validation set that consists of 20 percent of the training set. We stress that for training and validation we do not use any labeled outlier data. Our implementation of all the methods and the experiments is publicly available.¹

4.1. Baseline Method

We briefly describe the baseline method LEN. For training, the lengths of all the inter-event time intervals of the target events, $L = \{l_i : l_i = t_{i+1} - t_i, (t_i, x), (t_{i+1}, x) \in S_x\}$ are collected. Then, an empirical distribution of the inter-event time can be formulated as $\hat{F}(l) = \frac{1}{|L|} \sum_{i=1}^{|L|} \mathbb{I}[l_i \leq l]$. Here, for simplicity, we describe the method as if we only had one sequence in the training data, but it is easy to see how it works for multiple sequences, which is the case in our experiments.

For testing, LEN outputs a commission outlier score for a target event at time $t_n \in S_x$ as

$$s_c(t_n) = -\min\{\hat{F}(t_n - t_{n-1}), 1 - \hat{F}(t_n - t_{n-1})\}, \quad (19)$$

where $t_n - t_{n-1}$ is the inter-event time between the current and previous target events. Intuitively, if the inter-event time is too small ($\hat{F}(\cdot)$ is small) or too big ($1 - \hat{F}(\cdot)$ is small), it is likely that t_n has occurred at an abnormal time (too early or too late) and therefore is a commission outlier. The negation makes sure that a higher score indicates that it is more likely to be an outlier. For a blank interval B , LEN outputs an omission outlier score as the length of B ,

$$s_o(B) = |B|. \quad (20)$$

Intuitively, the longer the blank interval, the more likely it is to contain omission outliers.

4.2. Experiments on Synthetic Event Sequences

We generate synthetic event sequences using two different types of point processes. One is the inhomogeneous Poisson

process. The other is the Gamma process. For each type of processes, there is a set of parameters that determine the distribution of the points. We allow the parameters to vary according to a context state x , which can take two distinct values $x \in \{0, 1\}$.

For the inhomogeneous Poisson process, the CIF is a piecewise constant function with the value $\lambda = f(x)$, where x is the context state. In the experiments, we set $f(0) = 0.1$ and $f(1) = 1$. For the Gamma process, the inter-event time follows a Gamma distribution $Gam(a_x, b_x)$ (a_x shape, b_x rate), where x is the context state. In the experiments, we set $(a_0, b_0) = (10, 10)$ and $(a_1, b_1) = (100, 10)$.

The changes of the context state x are driven by a continuous-time Markov chain (Nodelman et al., 2002) with a transition matrix $Q = \begin{bmatrix} -0.05 & 0.05 \\ 0.05 & -0.05 \end{bmatrix}$ such that

$$p(x(t+dt) = j | x(t) = i) = \mathbb{I}[i = j] + Q_{ij} dt,$$

where dt is infinitesimal time. Each change of the state generates a context event.

For each point process type, we simulate 40 sequences. All sequences are simulated in the same time span $\mathcal{T} = [0, 1000]$. We use 50 percent of the sequences for training and the others for testing.

Outlier Simulation Commission outliers are simulated by adding points generated from a separate point process. Omission outliers are simulated by random removals of existing points. There is a parameter $\alpha(t)$ that defines the rate of outliers relative to normal points and can change over time (see Appendix E for details). To evaluate the performance of the proposed methods under different conditions, we conduct three types of simulations.

- **Constant rate.** In this case, the rate of omission and commission is a constant, $\alpha(t) = \alpha_0$, relative to the normal points, where $\alpha_0 = 0.1$. We also changed α_0 , but it did not affect the relative performance of each method in almost all cases (see Appendix I).
- **Periodic rate** (denoted as [sin]). In this case, the rate of omission and commission is a periodic function defined by $\alpha(t) = \alpha_0(1 + \sin(2\pi t/p))/2$ relative to the normal points, where α_0 controls the overall scale, t is the time, and p is the period. We set $\alpha_0 = 0.2$ and $p = 100$.
- **Piecewise-constant rate** (denoted as [pc]). In this case, we randomly generate a piecewise-constant function $g(t) : \mathcal{T} \rightarrow [0, 1]$ with a step size $s = 10$ for each sequence. The value of the function at each step is randomly generated from a uniform distribution. Then the rate of omission and commission is defined as

¹<https://github.com/siqil/CPPOD>

$\alpha(t) = \alpha_0 g(t)$ relative to the normal points, where $\alpha_0 = 0.2$ controls the overall scale.

We note that the first type satisfies the constant rate assumptions we made when developing the methods, but the second and third types do not. By using different simulations, we try to verify the robustness of our methods in different cases when the constant rate assumptions may not be satisfied.

Outlier Detection We apply the methods to detect outliers in an online manner, i.e., a commission (omission) outlier score is generated whenever a new target event is observed. For omission outliers, we also try to detect outliers at additional checkpoints within long blank intervals. See Appendix F for details of the approach and Appendix G for the outlier ratios of the datasets.

Results Table 1 shows the results in Area Under the Receiver Operating Characteristic curve (AUROC). CPPOD achieves the best performance for both commission and omission outliers, showing the effectiveness of our outlier scoring methods. PPOD being worse than CPPOD shows the importance of the context events in these cases. LEN performs much better than RND but worse than our proposed methods, which shows (1) it is an effective method for detecting outliers; (2) although intuitive, it is less effective than our proposed methods. Meanwhile, it does not have rigorous theoretical justifications.

As we can see, changing the outlier simulation mechanism does not affect the relative performance of each method. Even when the rates are not constants, i.e., $\alpha = [\sin]$ or $[pc]$, CPPOD and PPOD are still performing similarly as when α is a constant and much better than the baselines. We also checked the performance of our outlier scoring methods combined with the ground-truth model (see Appendix I), and the performance is similar to CPPOD.

Empirical Verification of the Bounds on FDR and FPR

To empirically verify the bounds on FDR and FPR presented in Section 3.5, we randomly repeat the experiments using our outlier scoring methods with the ground-truth model (replacing the learned LSTM model with the true model in CPPOD) on the synthetic data 10 times with different test data. Each time we calculate the FDR and FPR for different thresholds on the scores. For verifying FPR, we only test the inter-event time intervals for omission outliers. Their means and standard deviations over all repetitions are shown with the theoretical bounds in Figure 2 for Gamma process. For FPR, the bounds overlap the empirical rates. See Appendix H for more details and the results for Poisson process.




Figure 2. FDR (commission outlier) and FPR (omission outlier) on synthetic data (Gamma process).

4.3. Experiments on Real-World Clinical Data

In this part, we use real-world clinical data extracted from the MIMIC III dataset (Johnson et al., 2016). The dataset consists of de-identified electronic health records of ICU patients. We use four types of events corresponding to commonly used medications and lab tests as our targets and form four datasets by collecting the target events and corresponding context events. The target events and their context events are listed in Table 2. The medical category (medication, lab, or vital sign) of each type of events is in brackets following the type. For example, Potassium Chloride is a type of medication, and Potassium (Blood) is a type of lab test. The latter is used as the context for the former, as the administration of the medication can be triggered by observing an abnormally low value in the lab test.

We record in the data every type of events in the table. However, for Potassium (Blood) and Total Calcium (Blood), we further split the context events into three subtypes depending on whether the value in the lab test is low, normal, or high. For INR(PT) (previous state), we create context events of two subtypes based on whether the value of the *previous* event is normal or abnormal (with a 1-second delay). For Arterial Blood Pressure systolic (ABPs) and Non-invasive Blood Pressure systolic (NBPs), we split the events into two subtypes depending on whether the value is normal or low. These subtypes help to define better the contexts influencing the target events, since depending on their values, the target events can be more/less likely to occur.

All target and context events for one patient admission form one event sequence. We randomly select 2000 sequences for the first three datasets and 500 sequences for the last one, as each sequence in the last one contains much more events than the first three. For each dataset, we use 50 percent of the sequences for training and the others for testing.

Simulation Evaluation We generate and detect commission and omission outliers on top of the existing data using the same procedures for synthetic data (Section 4.2) except that we set $p = 24 \times 7$ and $s = 12$. This allows us to obtain

Table 1. AUROC on synthetic data. Dataset: name abbreviation (C=commission, O=omission) [α].

Dataset	Poi (C) [0.1]	Poi (C) [sin]	Poi (C) [pc]	Poi (O) [0.1]	Poi (O) [sin]	Poi (O) [pc]
RND	.500 (\pm .010)	.493 (\pm .007)	.512 (\pm .009)	.503 (\pm .008)	.498 (\pm .013)	.491 (\pm .007)
LEN	.601 (\pm .008)	.575 (\pm .006)	.584 (\pm .011)	.650 (\pm .006)	.659 (\pm .007)	.652 (\pm .011)
PPOD	.684 (\pm .010)	.661 (\pm .016)	.664 (\pm .009)	.737 (\pm .006)	.741 (\pm .012)	.734 (\pm .013)
CPPOD	.711 (\pm .012)	.707 (\pm .017)	.697 (\pm .014)	.778 (\pm .005)	.791 (\pm .010)	.784 (\pm .010)
Dataset	Gam (C) [0.1]	Gam (C) [sin]	Gam (C) [pc]	Gam (O) [0.1]	Gam (O) [sin]	Gam (O) [pc]
RND	.485 (\pm .007)	.493 (\pm .008)	.506 (\pm .007)	.505 (\pm .012)	.503 (\pm .010)	.515 (\pm .010)
LEN	.754 (\pm .006)	.762 (\pm .008)	.757 (\pm .005)	.799 (\pm .005)	.809 (\pm .006)	.813 (\pm .005)
PPOD	.816 (\pm .008)	.817 (\pm .006)	.813 (\pm .005)	.901 (\pm .007)	.902 (\pm .006)	.905 (\pm .006)
CPPOD	.871 (\pm .006)	.886 (\pm .004)	.870 (\pm .007)	.956 (\pm .003)	.956 (\pm .004)	.955 (\pm .004)

Table 2. Names of target and context events from MIMIC. INR=international normalized ratio; PT=prothrombin time.

Target	Context
INR(PT) [Lab]	INR(PT) [Lab] (previous state); Heparin [Medication]; Warfarin [Medication]
Calcium Gluconate [Medication]	Total Calcium (Blood) [Lab]
Potassium Chloride [Medication]	Potassium (Blood) [Lab]
Norepinephrine [Medication]	Arterial Blood Pressure systolic [Vital Sign]; Non-invasive Blood Pressure systolic [Vital Sign]

ground-truth labels for analyses. Table 3 shows the AUROC results. The results have more variations across different datasets in this case, which can be seen by examining the performance of LEN. Omission outliers appear to be more challenging than commission outliers except for INR(PT). CPPOD and PPOD outperform RND and LEN on all the datasets for both types of outliers.

In all cases, CPPOD is the best or very close to it. In the latter cases, the best is always PPOD, and the differences are very small. These are the cases where the additional context events are not as influential as the history of the target events themselves for the occurrences of the target events, so PPOD is as good as but simpler than CPPOD. However, for Potassium Chloride and Calcium Gluconate, we can see a clear advantage of CPPOD over PPOD by using additional context events.

Manual Evaluation of Clinical Relevance Finally, we apply CPPOD to the original data (without any simulated outliers). For each dataset, we select three commission outliers and three omission outliers with the highest outlier scores. We manually examine the electronic health record

of the patient for each outlier to determine whether the event commissions and omissions identified by CPPOD are clinically relevant and assign the outliers to three classes: N (outlier is not relevant), Y (outlier is relevant), and P (outlier is probably relevant given the available data but cannot be determined due to insufficient data recorded in MIMIC).

Specifically, we retrieve medical notes and related data for each case across time. We examine the notes and data to find evidence to support the decisions made by the physicians (giving a medication / lab test or not) that are detected as outliers. If there is evidence supporting the decision (e.g., given the condition of the patient, the medication should be given, and it was given), then the outlier is not clinically relevant (N). If the evidence is against the decision, it is clinically relevant (Y). If we do not have enough data, but the data we have do not support the decision, it is probably relevant (P).

Table 4 shows the results of the manual evaluation. Overall, CPPOD is performing well, as many detected outliers are clinically relevant. For both false outliers in Cal (C), there was an apparent change in the schedule of the medication (previous pattern being interrupted), so for the model, they do appear to be outliers. For Pot (C), one patient had rare conditions, for which we do not have the corresponding context events. The other had a drop in the value of the lab test (context), but it did not reach the *abnormally* low level, so the model does not have the exact context information that supported the human decision of giving the medication.

We note that for this evaluation, we do not have any guarantees on any of the assumptions we made when developing the methods in theory. Nonetheless, CPPOD is still effective. Moreover, as we have discovered, the original data may already contain outliers, so even in the previous evaluation, the models are unlikely to be trained on clean data, but Table 3 still shows the benefits of PPOD and CPPOD. This demonstrates that the assumptions do not have to be rigorously satisfied in practice for the methods to work reasonably well, although we may lose the theoretical guarantees on the performance, which depend on the assumptions.

Table 3. AUROC on MIMIC data. Dataset: name abbreviation (C=commission, O=omission) [α].

Dataset	INR (C) [0.1]	INR (C) [sin]	INR (C) [pc]	INR (O) [0.1]	INR (O) [sin]	INR (O) [pc]
RND	.496 (\pm .010)	.508 (\pm .009)	.488 (\pm .010)	.498 (\pm .011)	.516 (\pm .012)	.508 (\pm .009)
LEN	.596 (\pm .009)	.588 (\pm .010)	.607 (\pm .010)	.726 (\pm .008)	.717 (\pm .011)	.720 (\pm .011)
PPOD	.682 (\pm .010)	.675 (\pm .009)	.673 (\pm .008)	.748 (\pm .009)	.760 (\pm .010)	.773 (\pm .009)
CPPOD	.687 (\pm .009)	.680 (\pm .009)	.681 (\pm .010)	.746 (\pm .010)	.764 (\pm .009)	.770 (\pm .009)
Dataset	Cal (C) [0.1]	Cal (C) [sin]	Cal (C) [pc]	Cal (O) [0.1]	Cal (O) [sin]	Cal (O) [pc]
RND	.504 (\pm .013)	.502 (\pm .016)	.508 (\pm .011)	.493 (\pm .016)	.518 (\pm .017)	.496 (\pm .017)
LEN	.739 (\pm .012)	.688 (\pm .015)	.742 (\pm .011)	.526 (\pm .009)	.529 (\pm .012)	.541 (\pm .010)
PPOD	.830 (\pm .010)	.797 (\pm .010)	.837 (\pm .009)	.759 (\pm .008)	.758 (\pm .009)	.759 (\pm .011)
CPPOD	.866 (\pm .006)	.835 (\pm .009)	.860 (\pm .011)	.775 (\pm .008)	.777 (\pm .010)	.780 (\pm .009)
Dataset	Pot (C) [0.1]	Pot (C) [sin]	Pot (C) [pc]	Pot (O) [0.1]	Pot (O) [sin]	Pot (O) [pc]
RND	.498 (\pm .012)	.503 (\pm .010)	.511 (\pm .010)	.495 (\pm .017)	.508 (\pm .011)	.524 (\pm .010)
LEN	.733 (\pm .013)	.691 (\pm .009)	.718 (\pm .013)	.533 (\pm .012)	.536 (\pm .014)	.552 (\pm .011)
PPOD	.839 (\pm .009)	.813 (\pm .011)	.831 (\pm .008)	.736 (\pm .011)	.744 (\pm .011)	.746 (\pm .011)
CPPOD	.878 (\pm .009)	.857 (\pm .007)	.874 (\pm .010)	.748 (\pm .011)	.759 (\pm .012)	.761 (\pm .011)
Dataset	Nor (C) [0.1]	Nor (C) [sin]	Nor (C) [pc]	Nor (O) [0.1]	Nor (O) [sin]	Nor (O) [pc]
RND	.494 (\pm .014)	.536 (\pm .012)	.524 (\pm .012)	.510 (\pm .010)	.488 (\pm .014)	.503 (\pm .012)
LEN	.864 (\pm .010)	.837 (\pm .012)	.844 (\pm .016)	.468 (\pm .016)	.462 (\pm .013)	.476 (\pm .014)
PPOD	.890 (\pm .012)	.858 (\pm .014)	.884 (\pm .014)	.835 (\pm .010)	.842 (\pm .011)	.851 (\pm .011)
CPPOD	.897 (\pm .013)	.871 (\pm .014)	.882 (\pm .013)	.832 (\pm .009)	.837 (\pm .011)	.848 (\pm .010)

Table 4. Manual evaluation on MIMIC data.

Target	INR		Cal		Pot		Nor	
	C	O	C	O	C	O	C	O
Y	1	2	1	3	1	2	1	2
P	1	1	0	0	0	0	1	0
N	1	0	2	0	2	1	1	1

5. Discussion

In this work, we have studied two new outlier detection problems: detection of commission and omission outliers in continuous-time event sequences. We proposed outlier scoring methods based on Bayesian decision theory and hypothesis testing with theoretical guarantees. The proposed methods depend on a point-process model built from the data. In this work, we experimented with a model considering context, adapted from the continuous-time LSTM. We conducted experiments on both synthetic and real-world event sequences. The results show the flexibility of the adapted model and, more importantly, the effectiveness of the proposed outlier scoring methods.

As verified by the experiments, the assumptions we made when developing the methods in theory do not have to be satisfied for them to work reasonably well in practice, although when they are satisfied, we get the theoretical justifications and guarantees on the performance. For future work, there are two interesting directions. One is to develop algorithms to learn the models robustly in presence of outliers in the

training data. The other is to allow the model to adapt in a non-stationary environment. Both are orthogonal to this work, as they can be combined with the outlier scoring methods developed in this work to improve the performance by providing a better model.

Acknowledgement

This work was supported by NIH grant R01-GM088224. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. Siqi Liu also acknowledges the support by CS50 Merit Pre-doctoral Fellowship from the Department of Computer Science, University of Pittsburgh. Finally, the authors would like to thank all the anonymous reviewers who provided helpful comments on the earlier versions of the paper.

References

- Adams, R. P., Murray, I., and MacKay, D. J. Tractable nonparametric Bayesian inference in Poisson processes with Gaussian process intensities. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 9–16. ACM, 2009.
- Aggarwal, C. C. *Outlier Analysis*. Springer New York, 2013.
- Akçay, S., Atapour-Abarghouei, A., and Breckon, T. P. GANomaly: Semi-supervised anomaly detection via ad-

- versarial training. In Jawahar, C. V., Li, H., Mori, G., and Schindler, K. (eds.), *Computer Vision – ACCV 2018*, Lecture Notes in Computer Science, pp. 622–637, Cham, 2019. Springer International Publishing.
- Apostolopoulou, I., Linderman, S., Miller, K., and Dubrawski, A. Mutually regressive point processes. In *Advances in Neural Information Processing Systems*, pp. 5116–5127, 2019.
- Chandola, V., Banerjee, A., and Kumar, V. Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3):15:1–15:58, July 2009.
- Daley, D. J. and Vere-Jones, D. *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods*. Springer, New York, 2003.
- Ding, H., Khan, M., Sato, I., and Sugiyama, M. Bayesian nonparametric Poisson-process allocation for time-sequence modeling. In *International Conference on Artificial Intelligence and Statistics*, pp. 1108–1116, 2018.
- Du, N., Dai, H., Trivedi, R., Upadhyay, U., Gomez-Rodriguez, M., and Song, L. Recurrent marked temporal point processes: Embedding event history to vector. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1555–1564. ACM, 2016.
- Fawcett, T. and Provost, F. Adaptive fraud detection. *Data Mining and Knowledge Discovery*, 1(3):291–316, 1997.
- Fiore, U., Palmieri, F., Castiglione, A., and De Santis, A. Network anomaly detection with the restricted Boltzmann machine. *Neurocomputing*, 122:13–23, December 2013.
- Garcia-Teodoro, P., Diaz-Verdejo, J., Maciá-Fernández, G., and Vázquez, E. Anomaly-based network intrusion detection: Techniques, systems and challenges. *Computers & Security*, 28(1):18–28, 2009.
- Gupta, M., Gao, J., Aggarwal, C. C., and Han, J. Outlier detection for temporal data: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 26(9):2250–2267, 2013.
- Hauskrecht, M., Valko, M., Kveton, B., Visweswaram, S., and Cooper, G. Evidence-based anomaly detection. In *Annual American Medical Informatics Association Symposium*, pp. 319–324, November 2007.
- Hauskrecht, M., Batal, I., Valko, M., Visweswaram, S., Cooper, G. F., and Clermont, G. Outlier detection for patient monitoring and alerting. *Journal of Biomedical Informatics*, 46(1):47–55, February 2013.
- Hauskrecht, M., Batal, I., Hong, C., Nguyen, Q., Cooper, G. F., Visweswaram, S., and Clermont, G. Outlier-based detection of unusual patient-management actions: an ICU study. *Journal of Biomedical Informatics*, 64:211–221, 2016.
- Hawkes, A. G. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.
- Jia, J. and Benson, A. R. Neural jump stochastic differential equations. *arXiv:1905.10403 [cs, stat]*, January 2020.
- Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L.-w. H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Anthony Celi, L., and Mark, R. G. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3: 160035, May 2016.
- Lee, Y., Lim, K. W., and Ong, C. S. Hawkes processes with stochastic excitations. In *International Conference on Machine Learning*, pp. 79–88, 2016.
- Li, S., Xiao, S., Zhu, S., Du, N., Xie, Y., and Song, L. Learning temporal point processes via reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 10781–10791, 2018.
- Liu, S. and Hauskrecht, M. Nonparametric regressive point processes based on conditional Gaussian processes. In *Advances in Neural Information Processing Systems*, pp. 1062–1072, 2019.
- Lloyd, C., Gunter, T., Osborne, M., and Roberts, S. Variational inference for Gaussian process modulated Poisson processes. In *International Conference on Machine Learning*, pp. 1814–1822, 2015.
- Lloyd, C., Gunter, T., Osborne, M., Roberts, S., and Nickson, T. Latent point process allocation. In *Artificial Intelligence and Statistics*, pp. 389–397, May 2016.
- Mei, H. and Eisner, J. M. The neural Hawkes process: A neurally self-modulating multivariate point process. In *Advances in Neural Information Processing Systems*, pp. 6757–6767, 2017.
- Mei, H., Qin, G., and Eisner, J. Imputing missing events in continuous-time event streams. In *International Conference on Machine Learning*, pp. 4475–4485. PMLR, May 2019.
- Nodelman, U., Shelton, C. R., and Koller, D. Continuous time Bayesian networks. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, UAI’02, pp. 378–387, San Francisco, CA, USA, 2002. Morgan Kaufmann Publishers Inc.

- Rao, V. and Teh, Y. W. Gaussian process modulated renewal processes. In *Advances in Neural Information Processing Systems*, pp. 2474–2482, 2011.
- Shelton, C. R., Qin, Z., and Shetty, C. Hawkes process inference with missing data. In *Thirty-Second AAAI Conference on Artificial Intelligence*, April 2018.
- Song, X., Wu, M., Jermaine, C., and Ranka, S. Conditional anomaly detection. *Knowledge and Data Engineering, IEEE Transactions on*, 19(5):631–645, 2007.
- Trouleau, W., Etesami, J., Grossglauser, M., Kiyavash, N., and Thiran, P. Learning Hawkes processes under synchronization noise. In *International Conference on Machine Learning*, pp. 6325–6334, May 2019.
- Wang, Y., Xie, B., Du, N., and Song, L. Isotonic Hawkes processes. In *International Conference on Machine Learning*, pp. 2226–2234, 2016.
- Wong, W. K., Moore, A., Cooper, G., and Wagner, M. Bayesian network anomaly pattern detection for disease outbreaks. In *International Conference on Machine Learning*, pp. 808–815, August 2003.
- Xiao, S., Farajtabar, M., Ye, X., Yan, J., Song, L., and Zha, H. Wasserstein learning of deep generative point process models. In *Advances in Neural Information Processing Systems*, pp. 3247–3257, 2017.
- Xu, H., Farajtabar, M., and Zha, H. Learning Granger causality for Hawkes processes. In *International Conference on Machine Learning*, pp. 1717–1726, 2016.
- Xu, H., Luo, D., and Zha, H. Learning Hawkes processes from short doubly-censored event sequences. In *International Conference on Machine Learning*, pp. 3831–3840, July 2017.
- Xu, J. and Shelton, C. R. Intrusion detection using continuous time Bayesian networks. *Journal of Artificial Intelligence Research*, 39:745–774, December 2010.
- Zhou, K., Zha, H., and Song, L. Learning triggering kernels for multi-dimensional Hawkes processes. In *International Conference on Machine Learning*, pp. 1301–1309, 2013.