

---

# Fast and Flexible Temporal Point Processes with Triangular Maps

---

Oleksandr Shchur, Nicholas Gao, Marin Bilos, Stephan Günnemann  
Technical University of Munich, Germany  
{shchur, gaoni, bilos, guennemann}@in.tum.de

## Abstract

Temporal point process (TPP) models combined with recurrent neural networks provide a powerful framework for modeling continuous-time event data. While such models are flexible, they are inherently sequential and therefore cannot benefit from the parallelism of modern hardware. By exploiting the recent developments in the field of normalizing flows, we design TriTPP— a new class of non-recurrent TPP models, where both sampling and likelihood computation can be done in parallel. TriTPP matches the flexibility of RNN-based methods but permits orders of magnitude faster sampling. This enables us to use the new model for variational inference in continuous-time discrete-state systems. We demonstrate the advantages of the proposed framework on synthetic and real-world datasets.

## 1 Introduction

Temporal data lies at the heart of many high-impact machine learning applications. Electronic health records, financial transaction ledgers and server logs contain valuable information. A common challenge encountered in all these settings is that both the number of events and their times are variable. The framework of temporal point processes (TPP) allows us to naturally handle data that consists of variable-number events in continuous time. Du et al. [1] have shown that the flexibility of TPPs can be improved by combining them with recurrent neural networks (RNN). While such models are expressive and can achieve good results in various prediction tasks, they are poorly suited for sampling: sequential dependencies preclude parallelization. We show that it's possible to overcome the above limitation and design flexible TPP models without relying on RNNs. For this, we use the framework of triangular maps [2] and recent developments in the field of normalizing flows [3].

Our main contributions are: **(1)** We propose a new parametrization for several classic TPPs. This enables efficient parallel likelihood computation and sampling, which was impossible with existing parametrizations. **(2)** We propose TriTPP— a new class of non-recurrent TPPs. TriTPP matches the flexibility of RNN-based methods, while allowing orders of magnitude faster sampling. **(3)** We derive a differentiable relaxation for non-differentiable sampling-based TPP losses. This allows us to design a new variational inference scheme for Markov jump processes.

## 2 Background

**Temporal point processes (TPP)** [4] are stochastic processes that model the distribution of discrete events on some continuous time interval  $[0, T]$ . A realization of a TPP is a *variable-length* sequence of strictly increasing arrival times  $\mathbf{t} = (t_1, \dots, t_N), t_i \in [0, T]$ . We make the standard assumption and focus our discussion on regular finite TPPs [4]. One way to specify such a TPP is by using the (strictly positive) conditional intensity function  $\lambda^*(t) := \lambda(t|\mathcal{H}_t)$  that defines the rate of arrival of new events

---

Code and datasets are available under [www.daml.in.tum.de/triangular-tpp](http://www.daml.in.tum.de/triangular-tpp)

given the history  $\mathcal{H}_t = \{t_j : t_j < t\}$ . The  $*$  symbol reminds us of the dependence on the history [5]. Equivalently, we can consider the *cumulative* conditional intensity  $\Lambda^*(t) := \Lambda(t|\mathcal{H}_t) = \int_0^t \lambda^*(u)du$ , also known as the compensator.<sup>1</sup> We can compute the likelihood of a realization  $\mathbf{t}$  on  $[0, T]$  as

$$p(\mathbf{t}) = \left( \prod_{i=1}^N \lambda^*(t_i) \right) \exp \left( - \int_0^T \lambda^*(u)du \right) = \left( \prod_{i=1}^N \frac{\partial}{\partial t_i} \Lambda^*(t_i) \right) \exp(-\Lambda^*(T)) \quad (1)$$

For example, we can use a TPP to model the online activity of a user in a 24-hour interval. In this case, each realization  $\mathbf{t}$  could correspond to the timestamps of the posts by the user on a specific day.

**Triangular maps** [2] provide a framework that connects autoregressive models, normalizing flows and density estimation. Bogachev et al. [6] have shown that any density  $p(\mathbf{x})$  on  $\mathbb{R}^N$  can be equivalently represented by another density  $\tilde{p}(\mathbf{z})$  on  $\mathbb{R}^N$  and an increasing differentiable triangular map  $\mathbf{F} = (f_1, \dots, f_N) : \mathbb{R}^N \rightarrow \mathbb{R}^N$  that pushes forward  $p$  into  $\tilde{p}$ .<sup>2</sup> A map  $\mathbf{F}$  is called triangular if each component function  $f_i$  depends only on  $(x_1, \dots, x_i)$  and is an increasing function of  $x_i$ . Intuitively, we can think of  $\mathbf{F}$  as converting a random variable  $\mathbf{x} \sim p$  into a random variable  $\mathbf{z} := \mathbf{F}(\mathbf{x})$  with a density  $\tilde{p}$ . We can compute the density  $p(\mathbf{x})$  using the change of variables formula

$$p(\mathbf{x}) = |\det J_{\mathbf{F}}(\mathbf{x})| \tilde{p}(\mathbf{F}(\mathbf{x})) = \left( \prod_{i=1}^N \frac{\partial}{\partial x_i} f_i(x_1, \dots, x_i) \right) \tilde{p}(\mathbf{F}(\mathbf{x})) \quad (2)$$

where  $\det J_{\mathbf{F}}(\mathbf{x})$  is the Jacobian determinant of  $\mathbf{F}$  at  $\mathbf{x}$ . Here, we used the fact that  $J_{\mathbf{F}}(\mathbf{x})$  is a positive-definite lower-triangular matrix. To specify a complex density  $p(\mathbf{x})$ , we can pick some simple density  $\tilde{p}(\mathbf{z})$  and learn the triangular map  $\mathbf{F}$  that pushes  $p$  into  $\tilde{p}$ . It's important that  $\mathbf{F}$  and its Jacobian determinant can be evaluated efficiently if we are learning  $p(\mathbf{x})$  via maximum likelihood. We can sample from  $p(\mathbf{x})$  by applying the inverse map  $\mathbf{F}^{-1}$  to the samples drawn from  $\tilde{p}(\mathbf{z})$ . Note that  $\mathbf{F}^{-1} : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is also an increasing differentiable triangular map. Fast computation of  $\mathbf{F}^{-1}$  is important when learning  $p(\mathbf{x})$  via sampling-based losses (e.g., in variational inference).

### 3 Defining temporal point processes using triangular maps

We can notice the similarity between the right-hand sides of Equations 1 and 2, which seems to suggest some connection between TPPs and triangular maps. Indeed, it turns out that triangular maps can also be used to specify densities of point processes. Let  $\mathbf{t} = (t_1, \dots, t_N)$  be a realization of a TPP on  $[0, T]$  with compensator  $\Lambda^*$  (i.e. with density  $p(\mathbf{t})$ ). The random time change theorem states that in this case  $\mathbf{z} = (\Lambda^*(t_1), \dots, \Lambda^*(t_N))$  is a realization of a homogeneous Poisson process (HPP) with unit rate on the interval  $[0, \Lambda^*(T)]$  [4, Theorem 7.4.I][5, Proposition 4.1] (Figure 1).

The transformation  $\mathbf{F} = (f_1, \dots, f_N) : \mathbf{t} \mapsto \mathbf{z}$  is an increasing triangular map. Each component function  $f_i(\mathbf{t}) = \Lambda(t_i|t_1, \dots, t_{i-1})$  only depends on  $(t_1, \dots, t_i)$  and is increasing in  $t_i$  since  $\frac{\partial}{\partial t_i} \Lambda^*(t_i) = \lambda^*(t_i) > 0$ . The number  $N$  of the component functions  $f_i$  depends on the length of the specific realization  $\mathbf{t}$ . Notice that the term  $\prod_{i=1}^N \frac{\partial}{\partial t_i} \Lambda^*(t_i)$  in Equation 1 corresponds to the Jacobian determinant of  $\mathbf{F}$ . Similarly, the second term,  $\tilde{p}(\mathbf{z}) = \tilde{p}(\mathbf{F}(\mathbf{t})) = \exp(-\Lambda^*(T))$ , corresponds to the density of a HPP with unit rate on  $[0, \Lambda^*(T)]$  for any realization  $\mathbf{z}$ . This demonstrates that all TPP densities (Equation 1) correspond to increasing triangular maps (Equation 2). As for the converse of this statement, every increasing triangular map that is bijective on the space of increasing sequences defines a valid TPP (see Appendix C.3).

Our main idea is to define TPP densities  $p(\mathbf{t})$  by directly specifying the respective maps  $\mathbf{F}$ . In Section 3.1, we show how maps that satisfy certain properties allow us to efficiently compute density and generate samples. We demonstrate this by designing a new parametrization for several established models in Section 3.2. Finally, we propose a new class of fast and flexible TPPs in Section 3.3.

#### 3.1 Requirements for efficient TPP models

**Density evaluation.** The time complexity of computing the density  $p(\mathbf{t})$  for various TPP models can be understood by analyzing the respective map  $\mathbf{F}$ . For a general triangular map  $\mathbf{F} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ , computing  $\mathbf{F}(\mathbf{t})$  takes  $\mathcal{O}(N^2)$  operations. For example, this holds for Hawkes processes with

<sup>1</sup>For convenience, we provide a list of abbreviations and notation used in the paper in Appendix A.

<sup>2</sup>Note that some other works instead define  $\mathbf{F}$  as the map that pushes the density  $\tilde{p}(\mathbf{z})$  into  $p(\mathbf{x})$ .

arbitrary kernels [7]. If the compensator  $\Lambda^*$  has Markov property, the complexity of evaluating  $F$  can be reduced to  $\mathcal{O}(N)$  *sequential* operations. This class of models includes Hawkes processes with exponential kernels [8, 9] and RNN-based autoregressive TPPs [1, 10, 11]. Unfortunately, such models do not benefit from the parallelism of modern hardware. Defining an efficient TPP model will require specifying a forward map  $F$  that can be computed in  $\mathcal{O}(N)$  *parallel* operations.

**Sampling.** As a converse of the random time change theorem, we can sample from a TPP density  $p(t)$  by first drawing  $z$  from an HPP on  $[0, \Lambda^*(T)]$  and applying the inverse map,  $t = F^{-1}(z)$  [4]. There are, however, several caveats to this method. Not all parametrizations of  $F$  allow computing  $F^{-1}(z)$  in closed form. Even if  $F^{-1}$  is available, its evaluation for most models is again sequential [1, 9]. Lastly, the number of points  $N$  that will be generated (and thus  $\Lambda^*(T)$  for HPP) is not known in advance. Therefore, existing methods typically resort to generating the samples one by one [5, Algorithm 4.1]. We show that it's possible to do better than this. If the inverse map  $F^{-1}$  can be applied in parallel, we can produce large batches of samples  $t_i$ , and then discard the points  $t_i > T$  (Figure 2). Even though this method may produce samples that are later discarded, it is much more efficient than sequential generation on GPUs (Section 6.1).

To summarize, defining a TPP efficient for both density computation and sampling requires specifying a triangular map  $F$ , such that both  $F$  and its inverse  $F^{-1}$  can be evaluated analytically in  $\mathcal{O}(N)$  *parallel* operations. We will now show that maps corresponding to several classic TPP models can be defined to satisfy these criteria.

### 3.2 Fast temporal point process models

**Inhomogeneous Poisson process (IPP)** [4] is a TPP whose conditional intensity doesn't depend on the history,  $\Lambda(t|\mathcal{H}_t) = \Lambda(t)$ . The corresponding map is  $F = \Lambda$ , where  $\Lambda$  simply applies the function  $\Lambda : [0, T] \rightarrow \mathbb{R}_+$  elementwise to the sequence  $(t_1, \dots, t_N)$ .

**Renewal process (RP)** [12] is a TPP where each inter-event time  $t_i - t_{i-1}$  is sampled i.i.d. from the same distribution with the cumulative hazard function  $\Phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ . The compensator of an RP is  $\Lambda(t|\mathcal{H}_t) = \Phi(t - t_i) + \sum_{j=1}^i \Phi(t_j - t_{j-1})$ , where  $t_i$  is the last event before  $t$ . The triangular map of an RP can be represented as a composition  $F = C \circ \Phi \circ D$ , where  $D \in \mathbb{R}^{N \times N}$  is the pairwise difference matrix,  $C \equiv D^{-1} \in \mathbb{R}^{N \times N}$  is the cumulative sum matrix, and  $\Phi$  applies  $\Phi$  elementwise.

**Modulated renewal process (MRP)** [13] generalizes both inhomogeneous Poisson and renewal processes. The cumulative intensity is  $\Lambda(t|\mathcal{H}_t) = \Phi(\Lambda(t) - \Lambda(t_i)) + \sum_{j=1}^i \Phi(\Lambda(t_j) - \Lambda(t_{j-1}))$ . Again, we can represent the triangular map of an MRP as a composition,  $F = C \circ \Phi \circ D \circ \Lambda$ .

All three above models permit fast density evaluation and sampling. Since  $\Phi$  and  $\Lambda$  (as well as their inverses  $\Phi^{-1}$  and  $\Lambda^{-1}$ ) are elementwise transformations, they can obviously be applied in  $\mathcal{O}(N)$  parallel operations. Same holds for multiplication by the matrix  $D$ , as it is bidiagonal. Finally, the cumulative sum defined by  $C$  can also be computed in parallel in  $\mathcal{O}(N)$  [14]. Therefore, by reformulating IPP, RP and MRP using triangular maps, we can satisfy our efficiency requirements.

**Parametrization for  $\Phi$  and  $\Lambda$**  must satisfy several conditions. First, to define a valid TPP,  $\Phi$  and  $\Lambda$  have to be positive, strictly increasing and differentiable. Next, both functions, their derivatives (for density computation) and inverses (for sampling) must be computable in closed form to meet the efficiency requirements. Lastly, we want both functions to be highly flexible. Constructing such functions is not trivial. While IPP, RP and MRP are established models, none of their existing parametrizations satisfy all the above conditions simultaneously. Luckily, the same properties are necessary when designing normalizing flows [15]. Recently, Durkan et al. [3] used rational quadratic splines (RQS) to define functions that satisfy our requirements. We propose to use RQS to define  $\Phi$  and  $\Lambda$  for (M)RP and IPP. This parametrization is flexible, while also allowing efficient density evaluation and sampling — something that existing approaches are unable to provide (see Section 5).




Figure 1: Triangular map  $F(t) = (\Lambda^*(t_1), \dots, \Lambda^*(t_N))$  is used for computing  $p(t)$ .




Figure 2: Sampling is done by applying  $F^{-1}$  to a sample  $z$  from a HPP with unit rate.




Figure 3: TriTPP defines an expressive map  $F$  as a composition of easy-to-invert transformations.

### 3.3 Defining more flexible triangular maps

Even though the splines can make the functions  $\Phi$  and  $\Lambda$  arbitrarily flexible, the overall expressiveness of MRP is still limited. Its conditional intensity  $\lambda^*(t)$  depends only on the global time and the time since the last event. This means, MRP cannot capture, e.g., self-exciting [7] or self-correcting [16] behavior. We will now construct a model that is more flexible without sacrificing the efficiency.

The efficiency of the MRP stems from the fact that the respective triangular map  $F$  is defined as a composition of easy-to-invert transformations. More specifically, we are combining *learnable* element-wise nonlinear transformations  $\Phi$  and  $\Lambda$  with *fixed* lower-triangular matrices  $D$  and  $C$ . We can make the map  $F$  more expressive by adding *learnable* lower-triangular matrices into the composition. Using full  $N \times N$  lower triangular matrices would be inefficient (multiplication and inversion are  $\mathcal{O}(N^2)$ ), and also would not work for variable-length sequences (i.e., arbitrary values of  $N$ ). Instead, we define block-diagonal matrices  $B_l$ , where each block is a repeated  $H \times H$  lower-triangular matrix with strictly positive diagonal entries. Computing  $B_l^{-1}$  takes  $\mathcal{O}(H^2)$ , and multiplication by  $B_l$  or  $B_l^{-1}$  can be done in  $\mathcal{O}(NH)$  in parallel. We stack  $L$  such matrices  $B_l$  and define the triangular map  $F = C \circ \Phi_2 \circ B_L \circ \dots \circ B_1 \circ \Phi_1 \circ D \circ \Lambda$ . The blocks in every other layer are shifted by an offset  $H/2$  to let the model capture long-range dependencies. Note that now we use two element-wise learnable splines  $\Phi_1$  and  $\Phi_2$  before and after the block-diagonal layers. Figure 3 visualizes the overall sequence of maps and the Jacobians of each transformation. We name the temporal point process densities defined by the triangular map  $F$  as TriTPP.

Both the forward map  $F$  and its inverse  $F^{-1}$  can be evaluated in parallel in linear time, making TriTPP efficient for density computation and sampling. Our insight that TPP densities can be represented by increasing triangular maps was crucial for arriving at this result. Alternative representations of TriTPP, e.g., in terms of the compensator  $\Lambda^*$  or the conditional intensity  $\lambda^*$ , are cumbersome and do not emphasize the parallelism of the model. TriTPP and our parametrizations of IPP, RP, MRP can be efficiently implemented on GPU to handle batches of variable-length sequences (Appendix C).

## 4 Differentiable sampling-based losses for temporal point processes

Fast parallel sampling allows us to efficiently answer prediction queries such as "How many events are expected to happen in the next hour given the history?". More importantly, it enables us to efficiently train TPP models using objective functions of the form  $\mathbb{E}_p[g(\mathbf{t})]$ . This includes using  $p(\mathbf{t})$  to specify the policy in reinforcement learning [17], to impute missing data during training [11] or to define an approximate posterior in variational inference (Section 4.2). In all but trivial cases the expression  $\mathbb{E}_p[g(\mathbf{t})]$  has no closed form, so we need to estimate its gradients w.r.t. the parameters of  $p(\mathbf{t})$  using Monte Carlo (MC). Recall that we can sample from  $p(\mathbf{t})$  by applying the map  $F^{-1}$  to  $\mathbf{z}$  drawn from an HPP with unit rate. This enables the so-called reparametrization trick [18]. Unfortunately, this is not enough. Sampling-based losses for TPPs are in general not differentiable. This is a property of the loss functions that is independent of the parametrization of  $p(\mathbf{t})$  or the sampling method. In the following, we provide a simple example and a solution to this problem.

### 4.1 Entropy maximization

Consider the problem of maximizing the entropy of a TPP. An entropy penalty can be used as a regularizer during density estimation [19] or as a part of the ELBO in variational inference. Let  $p_\lambda(\mathbf{t})$  be a homogeneous Poisson process on  $[0, T]$  with rate  $\lambda > 0$ . It is known that the entropy is maximized when  $\lambda = 1$  [20], but for sake of example assume that we want to learn  $\lambda$  that maximizes the entropy  $-\mathbb{E}_p[\log p_\lambda(\mathbf{t})]$  with gradient ascent. We sample from  $p_\lambda(\mathbf{t})$  by drawing a sequence  $\mathbf{z} = (z_1, z_2, \dots)$  from a HPP with unit rate and applying the inverse map  $\mathbf{t} = F_\lambda^{-1}(\mathbf{z}) = \frac{1}{\lambda}\mathbf{z}$




Figure 4: Monte Carlo estimate of the entropy.




Figure 5: Maximizing the entropy with different values of  $\gamma$ .




Figure 6: Markov modulated Poisson process with 2 states.

(Figure 2). We obtain an MC estimate of the entropy using a single such sample  $\mathbf{t} = (t_1, t_2, \dots)$  as

$$-\mathbb{E}_p[\log p_\lambda(\mathbf{t})] \approx \lambda T - \sum_{i=1}^{\infty} \mathbf{1}(t_i \leq T) \log \lambda = \lambda T - \sum_{i=1}^{\infty} \mathbf{1}\left(\frac{1}{\lambda} z_i \leq T\right) \log \lambda \quad (3)$$

Here, the indicator function  $\mathbf{1}(\cdot)$  discards all the events  $t_i > T$ . We can see that for any sample  $\mathbf{z}$  the right-hand side of Equation 3 is not continuous w.r.t.  $\lambda$  at points  $\lambda = \frac{1}{T} z_i$ . At such points, decreasing  $\lambda$  by an infinitesimal amount will "push" the sample  $t_i = \frac{1}{\lambda} z_i$  outside the  $[0, T]$  interval, thus increasing  $\log p_\lambda(\mathbf{t})$  by a constant  $\log \lambda$ . We plot the right-hand side of Equation 3 as a function of  $\lambda$  in Figure 4, estimated with 5 MC samples. Clearly, such function cannot be optimized with gradient ascent. Increasing the number of MC samples almost surely adds more points of discontinuity and does not fix the problem. In general, non-differentiability arises when estimating expectations of a function  $g(\mathbf{t})$  that depends on the events  $t_i$  inside  $[0, T]$ . For any TPP density  $p(\mathbf{t})$ , the discontinuities occur at the parameter values that map the HPP realizations  $z_i$  exactly to the interval boundary  $T$ .

**Relaxation.** We obtain a differentiable approximation to Equation 3 by relaxing the indicator functions as  $\mathbf{1}(t_i \leq T) \approx \sigma_\gamma(T - t_i)$ , where  $\sigma_\gamma(x) = 1/(1 + \exp(-x/\gamma))$  is the sigmoid function with a temperature parameter  $\gamma > 0$ . Decreasing the temperature  $\gamma$  makes the approximation more accurate, but complicates optimization, similarly to the Gumbel-softmax trick [21]. Figure 5 shows convergence plots for different values of  $\gamma$ . Our relaxation applies to MC estimation of any function  $g(\mathbf{t})$  that can be expressed in terms of the indicator functions. This method also enables differentiable sampling with reparametrization from a Poisson distribution, which might be of independent interest.

## 4.2 Variational inference for Markov jump processes

Combining fast sampling (Section 3) with the differentiable relaxation opens new applications for TPPs. As an example, we design a variational inference scheme for Markov jump processes.

**Background.** A Markov jump process (MJP)  $\{s(t)\}_{t \geq 0}$  is a piecewise-constant stochastic process on  $\mathbb{R}_+$ . At any time  $t$ , the process occupies a discrete state  $s(t) \in \{1, \dots, K\}$ . The times when the state changes are called jumps. A trajectory of an MJP on an interval  $[0, T]$  with  $N$  jumps can be represented by a tuple  $(\mathbf{t}, \mathbf{s})$  of jump times  $\mathbf{t} = (t_1, \dots, t_N)$  and the visited states  $\mathbf{s} = (s_1, \dots, s_{N+1})$ . Note that  $N$  may vary for different trajectories. The prior over the trajectories  $p(\mathbf{t}, \mathbf{s} | \boldsymbol{\pi}, \mathbf{A})$  of an MJP is governed by an initial state distribution  $\boldsymbol{\pi}$  and a  $K \times K$  generator matrix  $\mathbf{A}$  (see Appendix B.1).

MJPs are commonly used to model the unobserved (latent) state of a system. In a latent MJP, the state  $s(t)$  influences the behavior of the system and indirectly manifests itself via some observations  $\mathbf{o}$ . For concreteness, we consider the Markov-modulated Poisson process (MMPP) [22]. In an MMPP, each of the  $K$  states of the MJP has an associated observation intensity  $\lambda_k$ . An MMPP is an inhomogeneous Poisson process where the intensity depends on the current MJP state as  $\lambda(t) = \lambda_{s(t)}$ . For instance, a 2-state MMPP can model the behavior of a social network user, who switches between an "active" (posting a lot) and "inactive" (working or sleeping) states (Figure 6). Given the observations  $\mathbf{o}$ , we might be interested in inferring the trajectory  $(\mathbf{t}, \mathbf{s})$ , the model parameters  $\boldsymbol{\theta} = \{\boldsymbol{\pi}, \mathbf{A}, \boldsymbol{\lambda}\}$ , or both.

**Variational inference.** The posterior distribution  $p(\mathbf{t}, \mathbf{s} | \mathbf{o}, \boldsymbol{\theta})$  of MMPP is intractable, so we approximate it with a variational distribution  $q(\mathbf{t}, \mathbf{s}) = q(\mathbf{t})q(\mathbf{s} | \mathbf{t})$ . Note that this is *not* a mean-field approximation used in other works [23]. We model the distribution over the jump times  $q(\mathbf{t})$  with TriTPP (Section 3.3). We find the best approximate posterior by maximizing the ELBO [24]

$$\max_{q(\mathbf{t})} \max_{q(\mathbf{s} | \mathbf{t})} \mathbb{E}_{q(\mathbf{t})} [\mathbb{E}_{q(\mathbf{s} | \mathbf{t})} [\log p(\mathbf{o} | \mathbf{t}, \mathbf{s}, \boldsymbol{\theta}) + \log p(\mathbf{t}, \mathbf{s} | \boldsymbol{\theta}) - \log q(\mathbf{t}, \mathbf{s})]] \quad (4)$$

Given jump times  $\mathbf{t}$ , the true posterior over the states  $p(\mathbf{s}|\mathbf{t}, \mathbf{o}, \boldsymbol{\theta})$  is just the posterior of a discrete hidden Markov model (HMM). This means that we only need to model  $q(\mathbf{t})$ ; the optimal  $q^*(\mathbf{s}|\mathbf{t})$ , i.e.

$$q^*(\mathbf{s}|\mathbf{t}) = \arg \max_{q(\mathbf{s}|\mathbf{t})} \mathbb{E}_{q(\mathbf{s}|\mathbf{t})} [\log p(\mathbf{o}|\mathbf{t}, \mathbf{s}, \boldsymbol{\theta}) + \log p(\mathbf{t}, \mathbf{s}|\boldsymbol{\theta}) - \log q(\mathbf{s}|\mathbf{t})] = p(\mathbf{s}|\mathbf{t}, \mathbf{o}, \boldsymbol{\theta}) \quad (5)$$

can be found by doing inference in an HMM — doable efficiently via the forward-backward algorithm [25]. The inner expectation w.r.t.  $q(\mathbf{s}|\mathbf{t})$  in Equation 4 can be computed analytically. We approximate the expectation w.r.t.  $q(\mathbf{t})$  with Monte Carlo. Since all terms of Equation 4 are not differentiable, we apply our relaxation from Section 4.1. We provide a full derivation of the ELBO and the implementation details in Appendix B.2.

The proposed framework is not limited to approximating the posterior over the trajectories. With small modifications (Appendix B.3), we can simultaneously learn the parameters  $\boldsymbol{\theta}$ , either obtaining a point estimate  $\boldsymbol{\theta}^*$  or a full approximate posterior  $q(\boldsymbol{\theta})$ . Our variational inference scheme can also be extended to other continuous-time discrete-state models, such as semi-Markov processes [26].

## 5 Related work

**Triangular maps** [2] can be seen as a generalization of autoregressive normalizing flows [27, 28, 15]. Existing normalizing flow models are either limited to fixed-dimensional data [29, 30] or are inherently sequential [31, 32]. Our model proposed in Section 3.3 can handle variable-length inputs, and allows for both  $\mathbf{F}$  and  $\mathbf{F}^{-1}$  to be evaluated efficiently in parallel.

**Sampling from TPPs.** Inverse method for sampling from inhomogeneous Poisson processes can be dated back to Çinlar [33]. However, traditional inversion methods for IPPs are different from our approach (Section 3). First, they are typically sequential. Second, existing methods either use extremely basic compensators  $\Lambda(t)$ , such as  $\lambda t$  or  $e^{\alpha t}$ , or require numerical inversion [34]. As an alternative to inversion, thinning approaches [35] became the dominant paradigm for generating IPPs, and TPPs in general. Still, sampling via thinning has a number of disadvantages. Thinning requires a piecewise-constant upper bound on  $\lambda(t)$ , which might not always be easy to find. If the bound is not tight, a large fraction of samples will be rejected. Moreover, thinning is not differentiable, doesn't permit reparametrization, and is hard to express in terms of parallel operations on tensors [36]. Our inversion-based sampling addresses all the above limitations. It's also possible to generate an IPP by first drawing  $N \sim \text{Poisson}(\Lambda(T))$  and then sampling  $N$  points  $t_i$  i.i.d. from a density  $p(t) = \lambda(t)/\Lambda(T)$  [37]. Unlike inversion, this method is only applicable to Poisson processes. Also, the operation of sampling  $N$  is not differentiable, which limits the utility of this approach.

**Inhomogeneous Poisson processes** are commonly defined by specifying the intensity function  $\lambda(t)$  via a latent Gaussian process [38]. Such models are flexible, but highly intractable. It's possible to devise approximations by, e.g., bounding the intensity function [39, 40]. Our spline parametrization of IPP compares favorably to the above models: it is also highly flexible, has a tractable likelihood and places no restrictions on the intensity. Importantly, it is much easier to implement and train. If uncertainty is of interest, we can perform approximate Bayesian inference on the spline coefficients [24]. Recently, Morgan et al. [41] used splines to model the intensity function of IPPs. Since  $\Lambda^{-1}$  cannot be computed analytically for their model, sampling via thinning is the only available option.

**Modulated renewal processes** have been known for a long time [13, 42], but haven't become as popular as IPPs among practitioners. This is not surprising, since inference and sampling in MRPs are even more challenging than in Cox processes [43, 44]. Our proposed parametrization addresses the shortcomings of existing approaches and makes MRPs straightforward to apply in practice.

**Neural TPPs.** Du et al. [1] proposed a TPP model based on a recurrent neural network. Follow-up works improved the flexibility of RNN-based TPPs by e.g. changing the RNN architecture [45], using more expressive conditional hazard functions [10, 46] or modeling the inter-event time distribution with normalizing flows [11]. All the above models are inherently sequential and therefore inefficient for sampling (Section 6.1). Recently, Turkmen et al. [36] proposed to speed up RNN-based *marked* TPPs by discretizing the interval  $[0, T]$  into a regular grid. Samples within each grid cell can be produced in parallel for each mark, but the cells themselves still must be processed sequentially.

**Latent space models.** TPPs governed by latent Markov dynamics have intractable likelihoods that require approximations [47, 48]. For MJPs, the state-of-the-art approach is the Gibbs sampler by Rao & Teh [49]. It allows to exactly sample from the posterior  $p(\mathbf{t}, \mathbf{s}|\mathbf{o}, \boldsymbol{\theta})$ , but is known to converge

slowly if the parameters  $\theta$  are to be learned as well [50]. Existing variational inference approaches for MJPs can only learn a fixed time discretization [23] or estimate the marginal statistics of the posterior [51, 52]. In contrast, our method (Section 4.2) produces a full distribution over the jump times.

## 6 Experiments

### 6.1 Scalability

**Setup.** The key feature of TriTPP is its ability to compute likelihood and generate samples in parallel, which is impossible for RNN-based models. We quantify this difference by measuring the runtime of the two models. We implemented TriTPP and RNN models in PyTorch [53]. The architecture of the RNN model is nearly identical to the ones used in [1, 10, 11], except that the cumulative conditional hazard function is parametrized with a spline [3] to enable closed-form sampling. Appendix E contains the details for this and other experiments. We measure the runtime of (a) computing the log-likelihood (and backpropagate the gradients) for a batch of 100 sequences of varying lengths and (b) sample sequences of the same sizes. We used a machine with an Intel Xeon E5-2630 v4 @ 2.20 GHz CPU, 256GB RAM and an Nvidia GTX1080Ti GPU. The results are averaged over 100 runs.

**Results.** Figure 7 shows the runtimes for varying sequence lengths. Training is rather fast for both models, on average taking 1-10ms per iteration. RNN is slightly faster for short sequences, but is outperformed by TriTPP on sequences with more than 400 events. Note that during training we used a highly optimized RNN implementation based on custom CUDA kernels (since all the event times  $t_i$  are already known). In contrast, TriTPP is implemented using generic PyTorch operations. When it comes to sampling, we notice a massive gap in performance between TriTPP and the RNN model. This happens because RNN-based TPPs are defined autoregressively and can only produce samples  $t_i$  one by one: to obtain  $p(t_i|t_1, \dots, t_{i-1})$  we must know all the past events. Recently proposed transformer TPPs [54, 55] are defined in a similar autoregressive way, so they are likely to be as slow for sampling as RNNs. TriTPP generates all the events in a sequence in parallel, which makes it more than 100 times faster than the recurrent model for longer sequences.




Figure 7: Scalability analysis. Standard devs. are below 1ms.

### 6.2 Density estimation

**Setup.** A fast TPP model is of little use if it cannot accurately learn the data distribution. The main goal of this experiment is to establish whether TriTPP can match the flexibility of RNN-based TPPs. As baselines, we use the IPP, RP and MRP models from Section 3.2 and Hawkes process [56].

**Datasets.** We use 6 synthetic datasets from Omi et al. [10]: Hawkes1&2 [7], self-correcting (SC) [16], inhomogeneous Poisson (IPP), renewal (RP) and modulated renewal (MRP) processes. Note that the data generators for IPP, RP and MRP by Omi et al. are *not* parametrized using splines, so these datasets are not guaranteed to be fitted perfectly by our models. We also consider 7 real-world datasets: PUBG (online gaming), Reddit-Comments, Reddit-Submissions (online discussions), Taxi (customer pickups), Twitter (tweets) and Yelp1&2 (check-in times). See Appendix D for more details.

**Metrics.** The standard metric for comparing generative models, including TPPs, is negative log-likelihood (NLL) on a hold-out set [36, 10, 11]. We partitioned the sequences in each dataset into train/validation/test sequences (60%/20%/20%). We trained the models by minimizing the NLL of the train set using Adam [57]. We tuned the following hyperparameters:  $L_2$  regularization  $\{0, 10^{-5}, 10^{-4}, 10^{-3}\}$ , number of spline knots  $\{10, 20, 50\}$ , learning rate  $\{10^{-3}, 10^{-2}\}$ , hidden size  $\{32, 64\}$  for RNN, number of blocks  $\{2, 4\}$  and block size  $\{8, 16\}$  for TriTPP. We used the validation set for hyperparameter tuning, early stopping and model development. We computed the results for the test set only once before including them in the paper. All results are averaged over 5 runs.

While NLL is a popular metric, it has known failure modes [58]. For this reason, we additionally computed maximum mean discrepancy (MMD) [59] between the test sets and the samples drawn from each model after training. To measure similarity between two realizations  $\mathbf{t}$  and  $\mathbf{t}'$ , we use a Gaussian kernel  $k(\mathbf{t}, \mathbf{t}') = \exp(-d(\mathbf{t}, \mathbf{t}')/2\sigma^2)$ , where  $d(\mathbf{t}, \mathbf{t}')$  is the "counting measure" distance from [60,



Table 1: Average test set NLL on synthetic and real-world datasets (lower is better). Best NLL in **bold**, second best underlined. Results with standard deviations can be found in Appendix F.1.

	Hawkes1	Hawkes2	SC	IPP	MRP	RP	PUBG	Reddit-C	Reddit-S	Taxi	Twitter	Yelp1	Yelp2
IPP	1.06	1.03	1.00	<b>0.71</b>	0.70	0.89	-0.06	-1.59	-4.08	<b>-0.68</b>	1.60	<u>0.62</u>	-0.05
RP	0.65	0.08	0.94	0.85	0.68	<b>0.24</b>	0.12	-2.08	-4.00	-0.58	1.20	0.67	-0.02
MRP	0.65	0.07	0.93	<b>0.71</b>	0.36	0.25	-0.83	-2.13	-4.38	<b>-0.68</b>	1.23	<b>0.61</b>	<b>-0.10</b>
Hawkes	<b>0.51</b>	0.06	1.00	0.86	0.98	0.39	0.11	<b>-2.40</b>	-4.19	-0.64	<b>1.04</b>	0.69	0.01
RNN	<u>0.52</u>	<b>-0.03</b>	<b>0.79</b>	0.73	<u>0.37</u>	<b>0.24</b>	<u>-1.96</u>	<b>-2.40</b>	<b>-4.89</b>	-0.66	1.08	0.67	-0.08
TriTPP	0.56	<u>0.00</u>	<u>0.83</u>	<b>0.71</b>	<b>0.35</b>	<b>0.24</b>	<b>-2.41</b>	-2.36	<u>-4.49</u>	-0.67	<u>1.06</u>	0.64	<u>-0.09</u>

Table 2: MMD between the hold-out test set and the generated samples (lower is better).

	Hawkes1	Hawkes2	SC	IPP	MRP	RP	PUBG	Reddit-C	Reddit-S	Taxi	Twitter	Yelp1	Yelp2
IPP	0.08	0.09	0.58	<b>0.02</b>	0.15	0.07	<b>0.01</b>	0.10	0.21	0.10	0.16	0.15	<u>0.16</u>
RP	0.06	0.06	1.13	0.34	1.24	<b>0.01</b>	0.46	0.07	0.18	0.57	0.14	0.16	0.23
MRP	0.05	0.06	0.50	<b>0.02</b>	<u>0.11</u>	0.02	<u>0.12</u>	0.09	0.20	<u>0.09</u>	0.13	<u>0.13</u>	<u>0.16</u>
Hawkes	<u>0.02</u>	0.04	0.58	0.36	0.65	0.05	0.16	<b>0.04</b>	0.35	0.20	0.20	0.20	0.32
RNN	<b>0.01</b>	<b>0.02</b>	<b>0.19</b>	0.09	0.17	<b>0.01</b>	0.23	<b>0.04</b>	<b>0.09</b>	0.13	<b>0.08</b>	0.19	0.18
TriTPP	0.03	<u>0.03</u>	<u>0.23</u>	<b>0.02</b>	<b>0.08</b>	<b>0.01</b>	0.16	0.07	<u>0.16</u>	<b>0.08</b>	<b>0.08</b>	<b>0.12</b>	<b>0.14</b>

Equation 3]. For completeness, we provide the definitions in Appendix E.2. MMD quantifies the dissimilarity between the true data distribution  $p^*(t)$  and the learned density  $p(t)$  — lower is better.

**Results.** Table 1 shows the test set NLLs for all models and datasets. We can see that the RNN model achieves excellent scores and outperforms the simpler baselines, which is consistent with earlier findings [1]. TriTPP is the only method that is competitive with the RNN — our method is within 0.05 nats of the best score on 11 out of 13 datasets. TriTPP consistently beats MRP, RP and IPP, which confirms that learnable block-diagonal transformations improve the flexibility of the model. The gap get larger on the datasets such as Hawkes, SC, PUBG and Twitter, where the inability of MRP to learn self-exciting and self-correcting behavior is especially detrimental. While Hawkes process is able to achieve good scores on datasets with "bursty" event occurrences (Reddit, Twitter), it is unable to adequately model other types of behavior (SC, MRP, PUBG).

Table 2 reports the MMD scores. The results are consistent with the previous experiment: models with lower NLL typically obtain lower MMD. One exception is the Hawkes process that achieves low NLL but high MMD on Taxi and Twitter. TriTPP again consistently demonstrates excellent performance. Note that MMD was computed using the test sequences that were unseen during training. This means that TriTPP models the data distribution better than other methods, and does not just simply overfit the training set. In Appendix F.1, we provide additional experiments for quantifying the quality of the distributions learned by different models. Overall, we conclude that TriTPP is flexible and able to model complex densities, in addition to being significantly more efficient than RNN-based TPPs.

### 6.3 Variational inference

**Setup.** We apply our variational inference method (Section 4.2) for learning the posterior distribution over the latent trajectories of an MMPP. We simulate an MMPP with  $K = 3$  latent states. As a baseline, we use the state-of-the-art MCMC sampler by Rao & Teh [49].

**Results.** Figure 8 shows the true latent MJP trajectory, as well as the marginal posterior probabilities learned by our method and the MCMC sampler of Rao & Teh. We can see that TriTPP accurately recovers the true posterior distribution over the trajectories. The two components that enable our new variational inference approach are our efficient parallel sampling algorithm for TriTPP (Section 3) and the differential relaxation (Section 4). Appendix F.2 contains an additional experiment on real-world data, where we both learn the parameters  $\theta$  and infer the posterior over the trajectories.




Figure 8: Posterior distributions over the latent trajectory of an MMPP learned using our VI approach & MCMC.



## 7 Future work & conclusions

**Future work & limitations.** We parametrized the nonlinear transformations of our TPP models with splines. Making a spline more flexible requires increasing the number of knots, which increases the number of parameters and might lead to overfitting. New deep *analytically invertible* functions will improve both our models, as well as normalizing flows in general. Currently, TriTPP is not applicable to marked TPPs [5]. Extending our model to this setting is an important task for future work.

**Conclusions.** We have shown that TPP densities can be represented with increasing triangular maps. By directly parametrizing the respective transformations, we are able to construct TPP models, for which both density evaluation and sampling can be done efficiently in parallel. Using the above framework, we defined TriTPP— a new class of flexible probability distributions over variable-length sequences. In addition to being highly efficient thanks to its parallelism, TriTPP shows excellent performance on density estimation, as shown by our experiments. High flexibility and efficiency of TriTPP allow it to be used as a plug-and-play component of other machine learning models.

### Broader impact

Existing works have applied TPPs and MJPs for analyzing electronic health records [61, 62], detecting anomalies in network traffic [63, 64] and modeling user behavior on online platforms [65, 66]. Thanks to fast sampling, our model can be used for solving new prediction tasks on such data, and the overall improved scalability allows practitioners to work with larger datasets. We do not find any of the above use cases ethically questionable, though, general precautions must be implemented when handling sensitive personal data. Since our model exploits fast parallel computations, has fewer parameters and converges in fewer iterations, it is likely to be more energy-efficient compared to RNN-based TPPs. However, we haven't performed experiments analyzing this specific aspect of our model.

### Acknowledgments

This research was supported by the German Federal Ministry of Education and Research (BMBF), grant no. 01IS18036B, the Software Campus Project Deep-RENT and by the BMW AG. The authors of this work take full responsibilities for its content.

### References

- [1] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. Recurrent marked temporal point processes: Embedding event history to vector. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- [2] Priyank Jaini, Kira A Selby, and Yaoliang Yu. Sum-of-squares polynomial flow. *International Conference on Machine Learning*, 2019.
- [3] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. In *Advances in Neural Information Processing Systems*, 2019.
- [4] Daryl J Daley and David Vere-Jones. An introduction to the theory of point processes. Vol. I. Probability and its applications, 2003.
- [5] Jakob Gulddahl Rasmussen. Temporal point processes: the conditional intensity function. *Lecture Notes, Jan*, 2011.
- [6] Vladimir Bogachev, Aleksandr Kolesnikov, and Kirill Medvedev. Triangular transformations of measures. *Sbornik: Mathematics*, 196(3):309, 2005.
- [7] Alan G Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.
- [8] David Oakes. The Markovian self-exciting process. *Journal of Applied Probability*, 12(1):69–77, 1975.

- [9] Angelos Dassios and Hongbiao Zhao. Exact simulation of Hawkes process with exponentially decaying intensity. *Electronic Communications in Probability*, 18, 2013.
- [10] Takahiro Omi, Naonori Ueda, and Kazuyuki Aihara. Fully neural network based model for general temporal point processes. In *Advances in Neural Information Processing Systems*, 2019.
- [11] Oleksandr Shchur, Marin Biloš, and Stephan Günnemann. Intensity-free learning of temporal point processes. *International Conference on Learning Representations*, 2020.
- [12] David Roxbee Cox. *Renewal Theory*. Methuen, 1962.
- [13] David Roxbee Cox. The statistical analysis of dependencies in point processes. *Stochastic Point Processes*. Wiley: New York, pages 55–66, 1972.
- [14] Guy E Blelloch. Prefix sums and their applications. Technical report, 1990.
- [15] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.
- [16] Valerie Isham and Mark Westcott. A self-correcting point process. *Stochastic Processes and Their Applications*, 8(3):335–347, 1979.
- [17] Utkarsh Upadhyay, Abir De, and Manuel Gomez Rodriguez. Deep reinforcement learning of marked temporal point processes. In *Advances in Neural Information Processing Systems*, 2018.
- [18] Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. Monte Carlo gradient estimation in machine learning. *arXiv preprint arXiv:1906.10652*, 2019.
- [19] Yves Grandvalet and Yoshua Bengio. Entropy regularization. *Semi-supervised learning*, pages 151–168, 2006.
- [20] François Baccelli and Jae Oh Woo. On the entropy and mutual information of point processes. In *IEEE International Symposium on Information Theory*, 2016.
- [21] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-softmax. *International Conference on Learning Representations*, 2017.
- [22] Wolfgang Fischer and Kathleen Meier-Hellstern. The Markov-modulated Poisson process cookbook. *Performance Evaluation*, 18, 1993.
- [23] Boqian Zhang, Jiangwei Pan, and Vinayak A Rao. Collapsed variational Bayes for Markov jump processes. In *Advances in Neural Information Processing Systems*, 2017.
- [24] Cheng Zhang, Judith Bütepage, Hedvig Kjellström, and Stephan Mandt. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):2008–2026, 2018.
- [25] Pengyu Wang and Phil Blunsom. Collapsed variational Bayesian inference for hidden Markov models. In *Artificial Intelligence and Statistics*, 2013.
- [26] William Feller. On semi-markov processes. *Proceedings of the National Academy of Sciences of the United States of America*, 51(4):653, 1964.
- [27] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. MADE: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, 2015.
- [28] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, 2016.
- [29] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. *International Conference on Learning Representations*, 2017.
- [30] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, 2017.

- [31] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.
- [32] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with PixelCNN decoders. In *Advances in Neural Information Processing Systems*, 2016.
- [33] Erhan Çinlar. *Introduction to Stochastic Processes*. Prentice-Hall, 1975.
- [34] Raghu Pasupathy. Generating nonhomogeneous Poisson processes. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- [35] PA W Lewis and Gerald S Shedler. Simulation of nonhomogeneous Poisson processes by thinning. *Naval research logistics quarterly*, 26(3):403–413, 1979.
- [36] Ali Caner Türkmen, Yuyang Wang, and Alexander J Smola. Fastpoint: Scalable deep point processes. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2019.
- [37] David Roxbee Cox. The statistical analysis of series of events. *Monographs on Applied Probability and Statistics*, 1966.
- [38] David R Cox. Some statistical methods connected with series of events. *Journal of the Royal Statistical Society: Series B (Methodological)*, 17(2):129–157, 1955.
- [39] Ryan Prescott Adams, Iain Murray, and David JC MacKay. Tractable nonparametric Bayesian inference in Poisson processes with Gaussian process intensities. In *International Conference on Machine Learning*, 2009.
- [40] Christian Donner and Manfred Opper. Efficient Bayesian inference of sigmoidal Gaussian Cox processes. *The Journal of Machine Learning Research*, 19, 2018.
- [41] Lucy Morgan, Barry Nelson, Andrew Titman, and David Worthington. A spline-based method for modelling and generating a nonhomogeneous poisson process. In *Winter Simulation Conference*, 2019.
- [42] Mark Berman. Inhomogeneous and modulated gamma processes. *Biometrika*, 68(1):143–152, 1981.
- [43] Vinayak Rao and Yee W Teh. Gaussian process modulated renewal processes. In *Advances in Neural Information Processing Systems*, 2011.
- [44] Thomas A Lasko. Efficient inference of Gaussian-process-modulated renewal processes with application to medical event data. In *Conference on Uncertainty in Artificial Intelligence*, 2014.
- [45] Hongyuan Mei and Jason M Eisner. The neural Hawkes process: A neurally self-modulating multivariate point process. In *Advances in Neural Information Processing Systems*, 2017.
- [46] Marin Biloš, Bertrand Charpentier, and Stephan Günnemann. Uncertainty on asynchronous time event prediction. In *Advances in Neural Information Processing Systems*, 2019.
- [47] Marcel Hirt and Petros Dellaportas. Scalable bayesian learning for state space models using variational inference with smc samplers. *International Conference on Artificial Intelligence and Statistics*, 2019.
- [48] Jing Wu, Owen Ward, James Curley, and Tian Zheng. Markov-modulated Hawkes processes for sporadic and bursty event occurrences. *arXiv preprint arXiv:1903.03223*, 2019.
- [49] Vinayak Rao and Yee Whye Teh. Fast MCMC sampling for Markov jump processes and extensions. *The Journal of Machine Learning Research*, 14, 2013.
- [50] Boqian Zhang and Vinayak Rao. Efficient parameter sampling for Markov jump processes. *arXiv preprint arXiv:1704.02369*, 2019.

- [51] Manfred Opper and Guido Sanguinetti. Variational inference for Markov jump processes. In *Advances in Neural Information Processing Systems*, 2008.
- [52] Christian Wildner and Heinz Koepl. Moment-based variational inference for Markov jump processes. *International Conference on Machine Learning*, 2019.
- [53] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.
- [54] Qiang Zhang, Aldo Lipani, Omer Kirnap, and Emine Yilmaz. Self-attentive hawkes process. *International Conference on Machine Learning*, 2020.
- [55] Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. Transformer hawkes process. *International Conference on Machine Learning*, 2020.
- [56] E. Bacry, M. Bompain, S. Gaïffas, and S. Poulsen. tick: a Python library for statistical learning, with a particular emphasis on time-dependent modeling. *arXiv preprint arXiv:1707.03003*, 2017.
- [57] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- [58] Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative models. *International Conference on Learning Representations*, 2016.
- [59] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13, 2012.
- [60] Shuai Xiao, Mehrdad Farajtabar, Xiaojing Ye, Junchi Yan, Le Song, and Hongyuan Zha. Wasserstein learning of deep generative point process models. In *Advances in Neural Information Processing Systems*, 2017.
- [61] Ahmed M Alaa and Mihaela Van Der Schaar. A hidden absorbing semi-markov model for informatively censored temporal data: Learning and inference. *The Journal of Machine Learning Research*, 19(1):108–169, 2018.
- [62] Qi Cao, Erik Buskens, Talitha Feenstra, Tiny Jaarsma, Hans Hillege, and Douwe Postmus. Continuous-time semi-markov models in health economic decision making: an illustrative example in heart failure disease management. *Medical Decision Making*, 36(1):59–71, 2016.
- [63] Alexander Ihler, Jon Hutchins, and Padhraic Smyth. Learning to detect events with markov-modulated poisson processes. *ACM Transactions on Knowledge Discovery from Data*, 1(3), 2007.
- [64] Ping Yan, Timothy Schoenharl, Alec Pawling, and Greg Madey. Anomaly detection in the WIPER system using Markov modulated Poisson process. 2007.
- [65] Manuel Gomez Rodriguez, David Balduzzi, and Bernhard Schölkopf. Uncovering the temporal dynamics of diffusion networks. *International Conference on Machine Learning*, 2011.
- [66] Charalampos Mavroforakis, Isabel Valera, and Manuel Gomez-Rodriguez. Modeling the dynamics of learning activity on the web. In *International Conference on World Wide Web*, 2017.
- [67] Vinayak Rao and Yee W. Teh. MCMC for continuous-time discrete-state systems. In *Advances in Neural Information Processing Systems*, 2012.