

A Variational Autoencoder for Neural Temporal Point Processes with Dynamic Latent Graphs

Sikun Yang^{1, 2, 3}, Hongyuan Zha^{4, 5}

¹ School of Computing and Information Technology, Great Bay University, 523000 Dongguan, China

² Great Bay Institute for Advanced Study, Great Bay University

³ Dongguan Key Laboratory for Data Science and Intelligent Medicine, Great Bay University

⁴ Shenzhen Institute of Artificial Intelligence and Robotics for Society

⁵ The Chinese University of Hong Kong, Shenzhen

sikunyang@gbu.edu.cn, zhahy@cuhk.edu.cn

Abstract

Continuously-observed event occurrences, often exhibit self- and mutually-exciting effects, which can be well modeled using temporal point processes. Beyond that, these event dynamics may also change over time, with certain periodic trends. We propose a novel variational auto-encoder to capture such a mixture of temporal dynamics. More specifically, the whole time interval of the input sequence is partitioned into a set of sub-intervals. The event dynamics are assumed to be stationary within each sub-interval, but could be changing across those sub-intervals. In particular, we use a sequential latent variable model to learn a dependency graph between the observed dimensions, for each sub-interval. The model predicts the future event times, by using the learned dependency graph to remove the non-contributing influences of past events. By doing so, the proposed model demonstrates its higher accuracy in predicting inter-event times and event types for several real-world event sequences, compared with existing state of the art neural point processes.

Introduction

There has been growing interests in modeling and understanding temporal dynamics in event occurrences. For instance, modeling customer behaviors and interactions, is crucial for recommendation systems and online social media, to improve the resource allocation and customer experience (Farajtabar et al. 2014, 2016). These event occurrences usually demonstrate heterogeneous dynamics. On one aspect, individuals usually reciprocate in their interactions with each other (*reciprocity*). For example, if *Alice* sends an email to *Bob*, then *Bob* is more likely to send an email to *Alice* soon afterwards. On the other aspect, long event sequences often exhibit a certain amount of *periodic* trends. For instance, during working time, individuals are more likely to reciprocate in the email interactions with their colleagues, but such mutually exciting effects will become weaker during non-working time, as illustrated in Fig. 1. Temporal point processes (TPPs), such as Hawkes processes (HPs) (Hawkes 1971), are in particular well-fitted to capture the *reciprocal* and *clustering* effects in event dynamics. Nonetheless, the conventional HPs cannot adequately




Figure 1: An example illustrates the email interactions between three individuals, *Alice* (A), *Bob* (B), *Jane* (J). During working time, *Alice* sends emails (black sticks) frequently with *Bob*, while they interact less frequently, during non-working time. We treat the sequence of emails from one individual to another as an observed dimension, corresponding to a vertex of the dependency graph. The dynamic graph between the four dimensions, aligned with three subintervals, are shown on the top.

capture the latent state-transition dynamics. Recently, neural temporal point processes (neural TPPs) demonstrate a strong capability in capturing long-range dependencies in event sequences using neural networks (Du et al. 2016; Xiao et al. 2017, 2019; Omi, Ueda, and Aihara 2019), attention mechanisms (Zhang et al. 2020; Zuo et al. 2020) and neural density estimation (Shchur, Bilos, and Günnemann 2019). These neural TPPs often use all the past events to predict a future event’s occurring time, and thus cannot remove the disturbances of the non-contributing events. To mitigate this defect, some recent works (Zhang, Lipani, and Yilmaz 2021; Lin et al. 2021) formulate the neural temporal point processes by learning a *static* graph to explicitly capture dependencies among event types. They hence can improve the accuracy in predicting future event time, by removing the non-contributing influences of past events via the learned graph. Nonetheless, the dependencies between event types, may also change over time. Using a *static* graph, the neural TPPs will retrieve a dependency graph averaged over time. To fill the gap, some main contributions are made in this paper: (i) We propose to learn a dynamic graph between the event types of an input sequence, from a novel variational auto-encoder (VAE) perspective (Kingma and Welling 2014; Rezende, Mohamed, and Wierstra 2014). More specifically, we use regularly-spaced intervals to capture the different

states, and assume *stationary* dynamics within each sub-interval. In particular, the dependencies between two event types, is captured using a latent variable, which allows to evolve over sub-intervals. We formulate the variational auto-encoder framework, by encoding a latent dynamic graph among event types, from observed sequences. The inter-event waiting times, are decoded using log-normal mixture distributions. Via the learned graph, the non-contributing influences of past events, can be effectively removed. (ii) The final experiments demonstrate the improved accuracy of the proposed method in predicting event time and types, compared against the existing closely-related methods. The interpretability of the dynamic graph estimated by the proposed method, is demonstrated with both New York Motor Vehicle Collision data.

Background

Multivariate Point Processes. Temporal point processes (TPPs) are concerned with modeling random event sequences in continuous time domain. Let $\mathcal{S} \equiv \{(t_i, v_i)\}_{i=1}^L$ denote a sequence of events, with $t_i \geq 0$ being the time-stamp and $v_i \in [1, \dots, U]$ being the type of i -th event. In addition, $\mathcal{H}_t = \{(t_i, v_i) \mid t_i < t, (t_i, v_i) \in \mathcal{S}\}$ denotes the sequence of historical events occurring up to time t . Multivariate Hawkes processes (MHPs) capture mutually-excitations among event types using the conditional intensity function specified by

$$\lambda_v^*(t) = \mu_v + \sum_{u=1}^U \sum_{\{j: t_j^u < t\}} \alpha_{(v,u)} \exp \left[-\frac{(t - t_j^u)}{\eta_{(v,u)}} \right], \quad (1)$$

where μ_v is the base rate of v -th event type, $\alpha_{(v,u)} > 0$ captures the instantaneous boost to the intensity due to event t_j^u 's arrival, and $\eta_{(v,u)} > 0$ determines the influence decay of that event over time. The stationary condition for MHPs requires $\alpha_{(v,u)}\eta_{(v,u)} < 1$. In contrast to MHPs, a mutually regressive point process (MRPP) (Apostolopoulou et al. 2019) is to capture both *excitatory* and *inhibitory* effects among event types. These parametric point processes capture a certain form of dependencies on the historical events by designing the conditional intensity functions accordingly. Despite being simple and useful, these parametric point processes either suffer from certain approximation errors caused by the model misspecifications in practice, or lack the ability to capture long-range dependencies. To address these limitations, some recent advancements (Du et al. 2016; Omi, Ueda, and Aihara 2019; Shchur, Bilos, and Günnemann 2019; Zhang et al. 2020; Zuo et al. 2020) combine temporal point processes and deep learning approaches to model complicated dependency structure behind event sequences. At a high level, these neural temporal point processes treat each event as a feature, and encode a sequence of events into an history embedding using various deep learning methods including recurrent neural nets (RNNs), gated recurrent units (GRUs), or long-short term memory (LSTM) nets. Du et al. (2016) uses a recurrent neural net to extract history embedding from observed past events, and then use the history embedding to parameterize its condi-

tional intensity function. The exponential form of this intensity, admits a closed-form integral $\Lambda^*(t) = \int_0^t \lambda^*(s) ds$, and thus leads to a tractable log-likelihood. Mei and Eisner (2017) studied a more sophisticated conditional intensity function, while calculating the log-likelihood involves approximating the integral $\Lambda^*(t)$ using Monte Carlo methods. Omi, Ueda, and Aihara (2019) proposes to model the cumulative conditional intensity function using neural nets, and hence allows to compute the log-likelihood exactly and efficiently. However, sampling using this approach is expensive, and the derived probability density function does not integrate to one. To remedy these issues, Shchur, Bilos, and Günnemann (2019) proposes to directly model the inter-event times using normalizing flows. The neural density estimation method (Shchur, Bilos, and Günnemann 2019) not only allows to perform sampling and likelihood computation analytically, but also shows competitive performance in various applications, compared with the other neural TPPs.

Variational Auto-Encoder. We briefly introduce the definition of variational auto-encoder (VAE), and refer the readers to (Kingma and Welling 2014; Rezende, Mohamed, and Wierstra 2014) for more properties. VAE is one of the most successful generative models, which allows to straightforwardly sample from the data distribution $p(\mathcal{S})$. It is in particular useful to model high-dimensional data distribution, for which sampling with Markov chain Monte Carlo is notoriously slow. More specifically, we aim at maximizing the data log-likelihood $p(\mathcal{S})$ under the generative process specified by $p(\mathcal{S}) = \int p(\mathcal{S} \mid \mathbf{z}, \theta) p(\mathbf{z}) d\mathbf{z}$, where \mathbf{z} is the latent variable, and $p(\mathbf{z})$ denotes the prior distribution, and the observation component $p(\mathcal{S} \mid \mathbf{z}, \theta)$ is parameterized by θ . Under the VAE framework, the posterior distribution $q_\phi(\mathbf{z} \mid \mathcal{S})$ can be defined as $q_\phi(\mathbf{z} \mid \mathcal{S}) \equiv \mathcal{N}(\mathbf{z}; f^\mu(\mathcal{S}; \phi), f^\Sigma(\mathcal{S}; \phi))$ where $\mathcal{N}(\cdot)$ refer to a normal distribution, the mean $f^\mu(\mathcal{S}; \phi)$ and covariance $f^\Sigma(\mathcal{S}; \phi)$ are parameterized by neural networks with parameter ϕ . To learn the model parameters, we maximize the evidence lower bound (ELBO) given by

$$\mathcal{L}(\phi, \theta) = \mathbb{E}_{q_\phi(\mathbf{z} \mid \mathcal{S})} \left[\log p(\mathcal{S} \mid \theta) \right] - \mathcal{D}_{\text{KL}} \left[q_\phi(\mathbf{z} \mid \mathcal{S}) \parallel p(\mathbf{z}) \right],$$

where \mathcal{D}_{KL} denotes the Kullback-Leibler (KL) divergence. The first term is to make the approximate posterior to produce latent variables \mathbf{z} that can reconstruct data \mathcal{S} as well as possible. The second term is to match the approximate posterior of the latent variables to the prior distribution of the latent variables. Using the reparameterization trick (Kingma and Welling 2014), we learn ϕ and θ by maximizing the ELBO using stochastic gradient descent aided by automatic differentiation.

Models

Given a sequence of events $\mathcal{S} \equiv \{(t_i, v_i)\}_{i=1}^L$, we aim to capture the complicated dependence between event types, using a dynamic graph-structured neural point process. The whole time-interval of the sequence is partitioned into K regularly-spaced sub-intervals with K specified *a priori*, to approximately represent the different states. We assume the latent graph among the event types, is changing over states, but *stationary* within each sub-interval, as illustrated in Fig. 2(a). Specifically, let $[t_k^L, t_k^R)$ stand for the k -th sub-




Figure 2: (a) The sequence consists of event types and timestamps, which are encoded by the history embedding $\{y_{t_i}^u\}$ with u being the event type. (b) The decoder captures the inter-event time for a future event, using the log-normal mixture model that is parameterized by the history embeddings. The dynamic graph captures the dependencies between those history embeddings.

interval with t_k^L being the start point and t_k^R being the end point. Let the latent variable $z_{(v,u)}^k$ capture the dependence of v -th event type on u -th event type within k -th sub-interval. For ease of exposition, we denote the set of events occurring within the k -th sub-interval by $\mathbf{s}^k \equiv \{(t_i, v_i) \mid t_i \in [t_k^L, t_k^R]\}$. Note that we model the inter-event time for each event type using log-normal mixtures. Hence, we represent the event sequence of u -th type by $\mathcal{S}_u = \{(t_i^u, \tau_i^u)\}_{i=1}^{n^u}$, where t_i^u is i -th event observed in the sequence of u -th event type, $\tau_i^u = t_i^u - t_{i-1}^u$ denotes the corresponding inter-event time, and the total number of events $L = \sum_{u=1}^U n^u$. Next we shall explain each component of the variational autoencoder in the following subsections.

Prior. We assume that the dependency graph among event types, evolves over sub-intervals. Hence, we use an autoregressive model to capture the prior probabilities of the latent variables $\{z_{(v,u)}^k\}_{v,u,k}$. More specifically, the prior distribution of $z_{(v,u)}^k$ depends on its previous state $z_{(v,u)}^{k-1}$ and the event sequence up to time t_k^L (all the events in first $k-1$ sub-intervals):

$$p_\phi(\mathbf{z} \mid \mathcal{S}) \equiv \prod_{k=1}^K p_\phi(\mathbf{z}^k \mid \mathbf{z}^{1:k-1}, \mathbf{s}^{1:k}).$$

The prior component is specified as follows: for each event of v -th type (t_i^v, m_i^v) , where m_i^v represents the aux-

iliary event mark if available, we embed t_i^v and m_i^v into a fixed-dimensional vector $\mathbf{y}_v^{t_i} \in \mathbb{R}^D$. Then, we pass the event embedding $\mathbf{y}_v^{t_i}$ through a fully-connected graph neural network (GNN) to obtain the relation embedding $\mathbf{h}_{(v,u),\text{emb}}^{t_i}$ between event types v and u :

$$\begin{aligned} \mathbf{h}_{v,1}^{t_i} &= f_{\text{emb}}^1(\mathbf{y}_v^{t_i}), \\ v \rightarrow e : \mathbf{h}_{(v,u),1}^{t_i} &= f_e^1([\mathbf{h}_{v,1}^{t_i}, \mathbf{h}_{u,1}^{t_i}]), \\ e \rightarrow v : \mathbf{h}_{v,2}^{t_i} &= f_v^1\left(\sum_{u \neq v} \mathbf{h}_{(v,u),1}^{t_i}\right), \\ v \rightarrow e : \mathbf{h}_{(v,u),\text{emb}}^{t_i} &= f_e^2([\mathbf{h}_{v,2}^{t_i}, \mathbf{h}_{u,2}^{t_i}]), \end{aligned}$$

where $f(\cdot)$ denotes a multilayer perceptron (MLP) for each layer of the GNN, $\mathbf{h}_{v,\ell}^{t_i}$ and $\mathbf{h}_{(v,u),\ell}^{t_i}$ represent the node-wise and edge-wise hidden states of the ℓ -th intermediate layer, respectively. The final output of the GNN $\mathbf{h}_{(v,u),\text{emb}}^{t_i}$ models the relation at time t_i . The GNN architecture is illustrated in Fig. 3(a).

We need to concatenate all the relation variable, $\{\mathbf{h}_{(v,u),\text{emb}}^{t_i}\}_{t_i \in [t_k^L, t_k^R]}$, and transform them into the relation state $\mathbf{h}_{(v,u),\text{emb}}^k$ of k -th sub-interval, using a MLP:

$$\mathbf{h}_{(v,u),\text{emb}}^k = f_{\text{emb}}^2([\mathbf{h}_{(v,u),\text{emb}}^{t_i}]) \quad \text{for } t_i \in [t_k^L, t_k^R].$$

A forward recurrent neural network (RNN) is used to capture the dependence of a relation state $\mathbf{h}_{(v,u),\text{fwd}}^k$ on its cur-

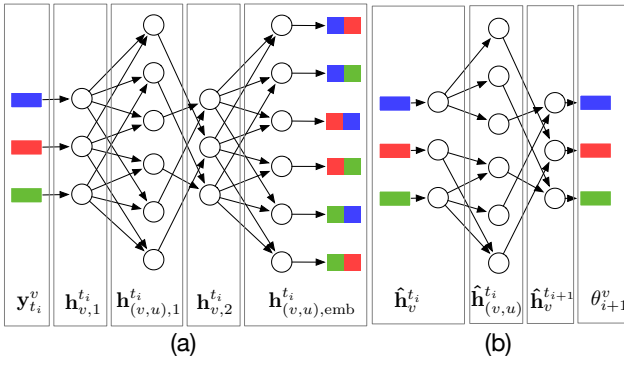


Figure 3: (a) A fully-connected GNN is used to transform event embeddings into relation embeddings between event types at each timestamp. (b) A GRNN is used to transform the past influences of those related event types into the current embedding according to the dynamic graphs.

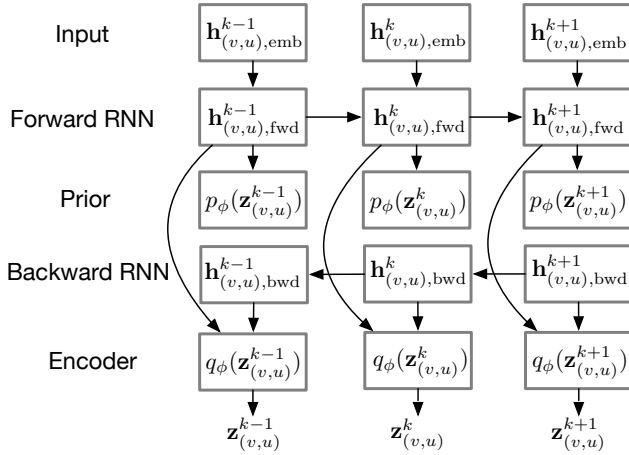


Figure 4: The relation embeddings are fed into a forward RNN and a backward RNN to encode the influences from the past and future relations, respectively.

rent embedding $\mathbf{h}_{(v,u),emb}^k$, and its previous state $\mathbf{h}_{(v,u),fwd}^{k-1}$:

$$\mathbf{h}_{(v,u),fwd}^k = \text{RNN}_{fwd}(\mathbf{h}_{(v,u),emb}^k, \mathbf{h}_{(v,u),fwd}^{k-1}).$$

Finally, we encode $\mathbf{h}_{(v,u),fwd}^k$ into the logits of the prior distribution for $\mathbf{z}_{(v,u)}^k$, using a MLP:

$$p_\phi(\mathbf{z}_{(v,u)}^k | \mathbf{z}^{1:k-1}, \mathbf{s}^{1:k}) = \text{softmax}(f_{\text{prior}}(\mathbf{h}_{(v,u),fwd}^k)).$$

Fig. 4 shows the prior distribution built upon the forward RNN.

Encoder. The posterior distribution of the latent variables $q_\phi(\mathbf{z} | \mathcal{S})$ depends on both the past and future events:

$$q_\phi(\mathbf{z} | \mathcal{S}) \equiv \prod_{k=1}^K q_\phi(\mathbf{z}^k | \mathcal{S}).$$

Hence, the encoder is designed to approximate the distribution of the relation variables using the whole event sequence. To that end, a backward GNN is used to propagate

the hidden states $\mathbf{h}_{(v,u),bwd}^k$ reversely:

$$\mathbf{h}_{(v,u),bwd}^k = \text{RNN}_{bwd}(\mathbf{h}_{(v,u),emb}^k, \mathbf{h}_{(v,u),bwd}^{k+1}).$$

Finally, we concatenate both the forward state $\mathbf{h}_{(v,u),fwd}^k$ and the backward state $\mathbf{h}_{(v,u),bwd}^k$, and transform them into the logits of the approximate posterior, using a MLP:

$$q_\phi(\mathbf{z}_{(v,u)}^k | \mathcal{S}) = \text{softmax}(f_{\text{enc}}([\mathbf{h}_{(v,u),bwd}^k, \mathbf{h}_{(v,u),fwd}^k])).$$

Note that the prior and encoder share parameters, and thus the parameters of the two components are denoted by ϕ .

Decoder. The role of the decoder is to predict the inter-event times $\{\tau_i^u\}_{i=1}^{\tau_u}$ for each event type u . In particular, we capture the latent dynamics $\{\hat{\mathbf{h}}_v^{t_i}\}$ behind these inter-event times using a graph recurrent neural network (GRNN) specified by

$$v \rightarrow e : \hat{\mathbf{h}}_{(v,u)}^{t_i} = z_{(v,u)}^k f_e^1([\hat{\mathbf{h}}_v^{t_i}, \hat{\mathbf{h}}_u^{t_i}]), \quad \text{for } t_i \in [t_k^L, t_k^R],$$

$$e \rightarrow v : \tilde{\mathbf{h}}_v^{t_i} = \sum_{u \neq v} \hat{\mathbf{h}}_{(v,u)}^{t_i},$$

$$\hat{\mathbf{h}}_v^{t_{i+1}} = \text{GRU}(\tilde{\mathbf{h}}_v^{t_i}, \hat{\mathbf{h}}_v^{t_i}),$$

where $z_{(v,u)}^k$ determines how the u -th event type $\hat{\mathbf{h}}_u^{t_i}$ influences the v -th event type $\hat{\mathbf{h}}_v^{t_{i+1}}$ through the relation $z_{(v,u)}^k$ at time t_{i+1} . The latent embedding $\hat{\mathbf{h}}_v^{t_i}$ itself evolves over time, using a gated recurrent unit (GRU).

Given the dynamic embeddings $\{\hat{\mathbf{h}}_v^{t_i}\}$, we model the inter-event time $p(\tau_i^u)$ using a log-normal mixture model (Shchur, Bilos, and Günnemann 2019),

$$p(\tau | \omega, \mu, \sigma) = \sum_{c=1}^C \omega_c \frac{1}{\tau \sigma_c \sqrt{2\pi}} \exp\left(-\frac{(\log \tau - \mu_c)^2}{2\sigma_c^2}\right),$$

where $\omega_c, \mu_c, \sigma_c$ represent the mixture weights, the mean and the standard deviations of c -th mixture component, respectively. In particular, the parameters of the distribution for each inter-event time τ_i^u is constructed as

$$\omega_i^u = \text{softmax}(V_\omega \hat{\mathbf{h}}_u^{t_i} + \beta_\omega), \quad \sigma_i^u = \exp(V_\sigma \hat{\mathbf{h}}_u^{t_i} + \beta_\sigma),$$

$$\mu_i^u = V_\mu \hat{\mathbf{h}}_u^{t_i} + \beta_\mu,$$

where $\{V_\omega, V_\sigma, V_\mu, \beta_\omega, \beta_\sigma, \beta_\mu\}$ refer to the learnable parameters. We use softmax and \exp transformations to impose the sum-to-one and positive constraints on the distribution parameters accordingly. Fig. 3(b) presents the GNN architecture used to construct the parameters of the log-normal mixtures in the decoder part. We assume the inter-event time τ_i^u is conditionally independent of the past events, given the model parameters. Hence, the distribution of the inter-event time under the decoder, factorizes as

$$p_\theta(\tau | \mathbf{z}) = \prod_{u=1}^U \prod_{i=1}^{\tau_u} p(\tau_i^u | \theta_i^u).$$

The decoder part of the VAE framework is illustrated in Fig. 2 (b). Hence, we can naturally make the next event time prediction using

$$\hat{t}_{i+1}^u = t_i^u + \int_0^\infty \tau p(\tau_{i+1}^u | \theta_{i+1}^u) d\tau.$$

Training. We next explain how to learn the parameters of the VAE framework for dynamic graph-structured neural point processes. The event sequences \mathcal{S} are passed through the GNN in the encoder, to obtain the relation embedding $\mathbf{h}_{(v,u),\text{emb}}^{t_i}$ for all the timestamps $\{t_i\}_{i=1}^L$ and each pair of two event types (v, u) . We then concatenate all the relation embeddings, and transform them into a relation state $\mathbf{h}_{(v,u),\text{emb}}^k$ for each sub-interval k . The relational states $\{\mathbf{h}_{(v,u),\text{emb}}^k\}$ are fed into the forward and backward RNNs to compute the prior distribution $p_\phi(\mathbf{z} | \mathcal{S})$ and posterior distribution $q_\phi(\mathbf{z} | \mathcal{S})$. We then sample $\{\mathbf{z}_{(v,u)}^k\}$ from the concrete reparameterizable approximation of the posterior distribution. The hidden states $\{\hat{\mathbf{h}}_v^{t_i}\}$ evolve through a GRNN, in which the messages can only pass through the non-zero edges hinted by $\{\mathbf{z}_{(v,u)}^k\}$. These hidden states $\{\hat{\mathbf{h}}_v^{t_i}\}$ are used to parameterize the log-normal mixture distribution for the inter-event times. To learn the model parameters, we calculate the evidence lower bound (ELBO) as

$$\begin{aligned} \mathcal{L}^{\text{ELBO}}(\phi, \theta) = & \mathbb{E}_{q_\phi(\mathbf{z}^k | \mathcal{S})} \left[\sum_{u=1}^U \sum_{i=1}^{n^u} \log p(\tau_i^u | \theta_i^u) \right] \quad (2) \\ & - \sum_{k=1}^K \mathcal{D}_{\text{KL}} \left[q_\phi(\mathbf{z}^k | \mathcal{S}) || p_\phi(\mathbf{z}^k | \mathbf{z}^{1:k-1}, \mathbf{s}^{1:k}) \right]. \end{aligned}$$

As we draw samples $\{\mathbf{z}_{(v,u)}^k\}$ using an reparameterizable approximation, we can calculate gradients using backpropagation and optimize the ELBO. Hereafter, we denote the proposed model as variational autoencoder temporal point process (VAETPP).

Related Work

Wasserman (1980); Huisman and Snijders (2003) proposed to capture the non-stationary network dynamics using continuous-time Markov chains. **Graph-Structured Temporal Point Processes.** Bhattacharjya, Subramanian, and Gao (2018) proposed a proximal graphical event model to infer the relationships among event types, and thus assumes that the occurrence of an event only depends on the occurring of its parents shortly before. Shang and Sun (2019) developed a geometric Hawkes process to capture correlations among multiple point processes using graph convolutional neural networks although the inferred graph is undirected. Zuo et al. (2020) developed a graph-structured transformer Hawkes process between multiple point processes. It assumes that each point process is associated with a vertex of a *static* graph, and model the dependency among those point processes by incorporating that graph into the attention module design. Zhang, Lipani, and Yilmaz (2021) formulated a graph-structured neural temporal point process (NTPP) by sampling a latent graph using a generator. The model parameters of the graph-structured NTPP and its associated graph, can be simultaneously optimized using an efficient bi-level programming. Lin et al. (2021) recently also developed a variational framework for graph-structured neural point processes by generating a latent graph using intra-type history embeddings. The latent graph is then used to govern the

message passing between intra-type embedding to decode the type-wise conditional intensity function. Zhang and Yan (2021) proposed to learn a static conditional graph between event types, by incorporating prior knowledge. Additionally, Linderman and Adams (2014) proposed to learn the latent graph structure among events via a probabilistic model. Wu et al. (2020) tries to model event propagation using a graph-biased temporal point process, with a graph specified by the following relationships among individuals on social media. Pan et al. (2020) studied a variational auto-encoder framework for modeling event sequences but haven't consider capturing latent graphs among dimensions. **Dynamic Latent Graphs behind Time Series.** Kipf et al. (2018) developed an variational autoencoder to learn a latent *static* graph among the entities that interact in a physical dynamical system. Following this success, Graber and Schwing (2020) used a sequential latent variable model to infer dynamic latent graphs among the entities from the discrete-time observations. Although Graber and Schwing (2020) and the proposed method both aim to learn dynamic graphs behind multivariate sequential observations, the main differences lie in that our input is *asynchronous* event sequences, for which we need to determine the regularly-spaced sub-intervals, and simultaneously learn a dynamic graph for each sub-interval. An interesting yet challenging direction, is to automatically infer the regularly-spaced time intervals and the dynamic graph structure, which we leave to future research.

Experiments

The proposed variational autoencoder temporal point process, is evaluated on the task of event time and type prediction. We used four real-world data to demonstrate the proposed method, compared with existing related methods. **New York Motor Vehicle Collisions(NYMVC):** This data contains a collection of vehicle collision events occurring at New York city since April, 2014. Each crash event (t_i, v_i) records a motor vehicle collision occurring in district v_i , at time t_i . Specifically, during peak periods, a vehicle collision may give rise to a sequence of collisions in the same or nearby districts, in short time. Hence, it is well-fitted to model and predict the occurrences of these events using multivariate point processes. In addition, the influence relations among the districts may change over time as the aforementioned triggering effects become weaker during night time, compared to day time. We created each event sequence using the motor vehicle collision records between 8:00 and 23:00, and treated each three hours as a sub-interval. We considered the five districts, *Manhattan, Brooklyn, Bronx, Queens* and *Staten Island* as the event types. **Stack Exchange Data:** The three stack exchange data from different sources, are included in the experiments: *MathOF*, *AskUbuntu*, and *SuperUser*. The stack exchange data consists of various interactions among the participant. Each event (v_i, u_i, t_i) means that at timestamp t_i , user v_i may post an answer or comment to u_i 's questions or comments. These interaction events among users usually exhibit a certain amount of *clustering* effects and periodic trends. For instance, some questions about popular technology, may quickly lead to a lot of answers or comments from the others who share similar

Methods	MathOF	AskUbuntu	SuperUser	NYMVC
Exponential	2.549 ± 0.074	2.584 ± 0.029	2.517 ± 0.018	2.474 ± 0.043
RMTTP	1.912 ± 0.087	1.981 ± 0.014	2.025 ± 0.054	1.944 ± 0.012
FullyNN	1.652 ± 0.062	1.884 ± 0.073	1.777 ± 0.023	1.473 ± 0.024
LogNormMix	-0.859 ± 0.121	0.303 ± 0.037	-0.868 ± 0.018	-2.578 ± 0.032
THP	-2.531 ± 0.024	-2.235 ± 0.028	-2.349 ± 0.051	-1.889 ± 0.037
VAETPP (static)	-2.632 ± 0.028	-2.312 ± 0.026	-2.466 ± 0.021	-2.016 ± 0.032
VAETPP	-3.501 ± 0.068	-2.867 ± 0.032	-3.812 ± 0.057	-5.952 ± 0.046

Table 1: Negative log-likelihood for inter-event time prediction on the real-world data.

Datasets	# sequences	# events	# types
MathOF	1453	590836	15
AskUbuntu	1561	65960	11
SuperUser	1240	84627	10
NYMVC	2000	863624	5

Table 2: Statistics of the datasets.

interests. Additionally, these triggering effects demonstrate periodic trends: users are more inclined to respond to technical topics during working days, than weekends/holidays. We consider the user who makes the action toward the other as the event type. Hence, we derived each sequence from events occurring within one week, and consider each day as a sub-interval. These datasets are detailed in Table 2.

Experimental Setup. We compared the ability of the models in predicting the inter-event times τ_i^u using the historical events $\mathcal{H}_{t_i^u}$, for each event type $u \in [1, \dots, U]$, as illustrated in Fig. 2(b). Each real-world data is split into multiple event sequences. For each real-world data, we choose the 60% of the sequences for training, 20% for validation, and 20% for testing. For training, we maximize the ELBO in Eq. 2 for the proposed model, and the expected log-likelihood for the other models. With the learned parameters, we measure the predictive performance of each model, using its obtained negative log-likelihood (NLL) on the validation set. Hence, the model configuration that achieves best predictive performance, can be chose using the validation set. Finally, the NLL loss on the test set, are used to compare the ability of the models in predicting inter-event times. We report the results averaged over ten random training/validation/testing splits. For the developed VAETPP, the dimension of the input embedding $\mathbf{y}_{t_i^u}^u$ is 64. For the fully-connected GNN of the encoder, f_{emb}^1, f_e^1, f_v^1 , and f_e^1 are two-layer MLPs with 64 units for each layer and Exponential Linear Unit (ELU) activations. We use f_{emb}^2 to transform the concatenated hidden states within each sub-interval into one hidden state, and thus parameterize f_{emb}^2 using a one-layer MLP with 64 hidden units and ReLU activations. Both the forward RNN and backward RNN have 64 hidden units. We parameterize f_{prior} and f_{enc} by a one-layer MLP with 64 hidden units and Rectified Linear Unit (ReLU) activations. We set the number of edge types of the dynamic graph among event types to be two, and specify the first edge type to indicate no dependency. For the decoder part, we parameterize f_e^1 using a separate two-layer MLP with 64 hidden/output units, for each of the two edge types. The

GRU has 64 hidden units. We chose the number of mixing components used in the log-normal mixture distribution for the VAETPP, using the validation data. In the experiments, we used 16 mixing components in the experiments. We also consider restricting the VAETPP with a static latent graph, and denote it as VAETPP (static), to validate the importance of learning dynamic graphs in capturing period trends in event sequences. We compared the proposed method against the following baselines: **Exponential**. The conditional intensity function of the constant intensity model (Upadhyay, De, and Gomez-Rodriguez 2018) is defined as $\lambda^*(t_i) = \exp(\mathbf{v}^T \mathbf{h}_i + \mathbf{b})$, in which \mathbf{h}_i denotes the event history embedding learned by a RNN, \mathbf{v} and \mathbf{b} refer to the model parameters. The probability density function (PDF) of the constant intensity model is an exponential distribution, as given by $p^*(\tau) = \gamma \exp(-\gamma)$, where $\gamma = \exp(\mathbf{v}^T \mathbf{h}_i + \mathbf{b})$. **Recurrent Marked Temporal Point Processes(RMTTP)** (Du et al. 2016). The method encodes past events into historical embeddings using a RNN, and models exponential-distributed conditional intensity. **Fully Neural Networks (FullyNN)** (Omi, Ueda, and Aihara 2019). It captures the cumulative distributions of inter-event times using a neural network. **Log Normal Mixture (LogNormMix)** (Shchur, Bilos, and Günnemann 2019). The method encodes the event history into embedding vectors with a RNN, and decodes the inter-waiting time with a log-normal mixture distribution. **Transformer Hawkes Process (THP)** (Zuo et al. 2020). It leverages the self-attention mechanism to capture long-term dependencies in observed event sequences.

Negative Log-likelihood Comparison. Table 1 compares the negative log-likelihood loss of all the methods in modeling inter-event times. As expected, the LogNormMix admits more flexibility compared with simple model using unimodal distributions (Gompertz/RMTTP, Exponential), and thus shows improved performance by large margins. Transformer Hawkes process (THP) can effectively learn long-range dependencies among events, and thus achieves lower NLL loss. The proposed VAETPP not only can capture the complicated inter-event time distribution using a log-normal mixture decoder. It can also improve the inter-event time prediction in further by effectively removing the non-contributing effects of irrelevant past events, via the dynamic dependency graph. Hence, the proposed VAETPP achieves the best NLL loss values consistently on all the datasets.

Event Prediction Comparison. We also consider the tasks of event time and type prediction in the experiments. In particular, following (Zuo et al. 2020), we make the next event time prediction using a linear predictor as $\hat{t}_{i+1}^u =$

Methods	MathOF	AskUbuntu	SuperUser	NYMVC
RMTTP	0.952 ± 0.008	0.983 ± 0.025	1.103 ± 0.068	1.135 ± 0.095
LogNormMix	0.673 ± 0.082	0.969 ± 0.043	0.708 ± 0.098	0.798 ± 0.038
THP	0.693 ± 0.024	0.791 ± 0.057	0.779 ± 0.032	0.859 ± 0.027
VAETPP(static)	0.632 ± 0.020	0.812 ± 0.032	0.788 ± 0.038	0.864 ± 0.046
VAETPP	0.569 ± 0.018	0.642 ± 0.012	0.674 ± 0.046	0.775 ± 0.064

Table 3: Root mean square error (RMSE) comparison for event time prediction.

Methods	MathOF	AskUbuntu	SuperUser	NYMVC
RMTTP	0.154 ± 0.022	0.189 ± 0.021	0.208 ± 0.008	0.251 ± 0.025
LogNormMix	0.206 ± 0.027	0.225 ± 0.011	0.235 ± 0.013	0.276 ± 0.038
THP	0.242 ± 0.019	0.261 ± 0.016	0.248 ± 0.006	0.294 ± 0.028
VAETPP(static)	0.286 ± 0.031	0.286 ± 0.022	0.256 ± 0.018	0.278 ± 0.026
VAETPP	0.321 ± 0.016	0.318 ± 0.012	0.288 ± 0.018	0.301 ± 0.022

Table 4: Event type prediction accuracy comparison.

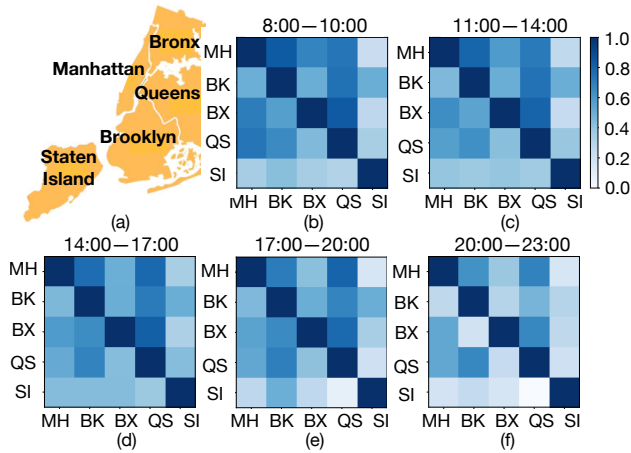


Figure 5: (a) shows the relative locations of the five boroughs of New York city, Manhattan (MH), Bronx (BX), Brooklyn (BK), Queens (QS), Staten Island (SI); (b-f) show the latent dynamic graph between the five boroughs over five time-intervals, estimated by the VAETPP.

$\mathbf{W}^{\text{time}} \theta_i^u$, where θ_i^u is the history embedding updated by the VAETPP after observing of i -th event of u -th type, and $\mathbf{W}^{\text{time}} \in \mathbb{R}^{1 \times D}$ denotes the event time predictor’s parameter. The next event type prediction is

$$\hat{\mathbf{p}}_{i+1} = \text{softmax}(\mathbf{W}^{\text{type}} \theta_i^u), \hat{m}_{i+1}^u = \arg \max_j \hat{\mathbf{p}}_{i+1}(j),$$

where $\mathbf{W}^{\text{type}} \in \mathbb{R}^{J \times D}$ denotes the event type predictor’s parameter, and $\hat{\mathbf{p}}_{i+1}(j)$ refers to the j -th entry of $\hat{\mathbf{p}}_{i+1} \in \mathbb{R}^J$. The loss functions for event time and type prediction are defined as

$$\tilde{\mathcal{L}}(\mathcal{S}; \theta) = \sum_u \sum_{i=1}^{n_u} (t_i^u - \hat{t}_i^u)^2, \hat{\mathcal{L}}(\mathcal{S}; \theta) = - \sum_{j=2}^L \mathbf{m}_j^T \log(\hat{\mathbf{p}}_j),$$

respectively, where \mathbf{m}_j is the one-hot encoding for the type of j -th event. To learn the parameters of the event time and type predictor, we consider minimizing a composite loss function as

$\min_{\phi, \theta} -\mathcal{L}^{\text{ELBO}}(\mathcal{S}; \phi, \theta) + \tilde{\mathcal{L}}(\mathcal{S}; \theta) + \hat{\mathcal{L}}(\mathcal{S}; \theta)$, where $\mathcal{L}^{\text{ELBO}}(\mathcal{S}; \phi, \theta)$ is the evidence lower bound derived in Eq.2. We used the training data to learn the model parameters, and choose the best configuration according to the predictive performance on the validation set. Finally, we evaluated model performance on the test set. Specifically, we predicted each held-out event (t_j, m_j) given its history. We evaluated event type prediction by accuracy and event time prediction by Root Mean Square Error (RMSE). Tab. 3 and 4 shows the results for event time and type prediction, respectively. Our VAETPP outperforms the baselines in predicting event time and types on all the data.

Model Interpretability. Fig. 5(a) shows the relative location of the five boroughs of New York city: Manhattan, Brooklyn, Bronx, Queens, Staten Island. The latent dynamic graphs for the five time-intervals, are plotted in Fig. 5(b-f). From the results, we find that the influences between Manhattan, Brooklyn, Bronx and Queens, are much stronger, compared with influences between these four areas and Staten Island. These influences between all the areas, gradually become weaker during night time, compared against day time. The results not only show the high model interpretability, but also explains why VAETPP obtains better prediction accuracy by effectively removing non-contributing historical events’ influences via its estimated graphs.

Conclusion

We have presented a novel variational auto-encoder for modeling asynchronous event sequences. To capture the *periodic* trends behind long sequences, we use regularly-spaced intervals to capture the different states behind the sequences, and assume stationary dynamics within each sub-interval. The dependency structure between event types, is captured using latent variables, which allow to be evolving over time, to capture the time-varying graphs. Hence, the proposed model can effectively remove the influences from the irrelevant past event types, and achieves better accuracy in predicting inter-event times and types, compared with other neural point processes. We plan to generalize the work to capture non-stationary network dynamics (Yang and Koepl 2018b,a, 2020; Yang and Zha 2023) in the future research.

Acknowledgments

The work is partially funded by Shenzhen Science and Technology Program (JCYJ20210324120011032) and Shenzhen Institute of Artificial Intelligence and Robotics for Society.

References

- Apostolopoulou, I.; Linderman, S.; Miller, K.; and Dubrawski, A. 2019. Mutually Regressive Point Processes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1–12.
- Bhattacharjya, D.; Subramanian, D.; and Gao, T. 2018. Proximal Graphical Event Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1–10.
- Du, N.; Dai, H.; Trivedi, R. S.; Upadhyay, U.; Gomez-Rodriguez, M.; and Song, L. 2016. Recurrent Marked Temporal Point Processes: Embedding Event History to Vector. In *SIGKDD*, 1555–1564. New York, NY, USA.
- Farajtabar, M.; Du, N.; Gomez-Rodriguez, M.; Valera, I.; Zha, H.; and Song, L. 2014. Shaping Social Activity by Incentivizing Users. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2474–2482.
- Farajtabar, M.; Ye, X.; Harati, S.; Song, L.; and Zha, H. 2016. Multistage Campaigning in Social Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2–9.
- Graber, C.; and Schwing, A. G. 2020. Dynamic Neural Relational Inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8513–8522.
- Hawkes, A. G. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1): 83–90.
- Huisman, M.; and Snijders, T. 2003. Statistical analysis of longitudinal network data with changing composition. *Sociological Methods & Research*, 32(2): 253–287.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 1–14.
- Kipf, T.; Fetaya, E.; Wang, K.-C.; Welling, M.; and Zemel, R. 2018. Neural Relational Inference for Interacting Systems. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2688–2697.
- Lin, H.; Tan, C.; Wu, L.; Gao, Z.; and Li, S. Z. 2021. An Empirical Study: Extensive Deep Temporal Point Process. *CoRR*.
- Linderman, S. W.; and Adams, R. P. 2014. Discovering Latent Network Structure in Point Process Data. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1413–1421. Beijing, China.
- Mei, H.; and Eisner, J. 2017. The Neural Hawkes Process: A Neurally Self-Modulating Multivariate Point Process. In *Advances in Neural Information Processing Systems (NeurIPS)*, 6757–6767.
- Omi, T.; Ueda, N.; and Aihara, K. 2019. Fully Neural Network based Model for General Temporal Point Processes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1–11.
- Pan, Z.; Huang, Z.; Lian, D.; and Chen, E. 2020. A Variational Point Process Model for Social Event Sequences. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 173–180.
- Rezende, D. J.; Mohamed, S.; and Wierstra, D. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1278–1286. Beijing, China.
- Shang, J.; and Sun, M. 2019. Geometric Hawkes Processes with Graph Convolutional Recurrent Neural Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 4878–4885.
- Shchur, O.; Bilos, M.; and Günnemann, S. 2019. Intensity-Free Learning of Temporal Point Processes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 1–21.
- Upadhyay, U.; De, A.; and Gomez-Rodriguez, M. 2018. Deep Reinforcement Learning of Marked Temporal Point Processes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 3172–3182.
- Wasserman, S. 1980. Analyzing Social Networks as Stochastic Processes. *Journal of the American Statistical Association*, 75(370): 280–294.
- Wu, W.; Liu, H.; Zhang, X.; Liu, Y.; and Zha, H. 2020. Modeling Event Propagation via Graph Biased Temporal Point Process. *IEEE Transactions on Neural Networks and Learning Systems*, 1–11.
- Xiao, S.; Farajtabar, M.; Ye, X.; Yan, J.; Yang, X.; Song, L.; and Zha, H. 2017. Wasserstein Learning of Deep Generative Point Process Models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Xiao, S.; Yan, J.; Farajtabar, M.; Song, L.; Yang, X.; and Zha, H. 2019. Learning Time Series Associated Event Sequences With Recurrent Point Process Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 30(10): 3124–3136.
- Yang, S.; and Koepl, H. 2018a. Dependent Relational Gamma Process Models for Longitudinal Networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 5551–5560.
- Yang, S.; and Koepl, H. 2018b. A Poisson Gamma Probabilistic Model for Latent Node-Group Memberships in Dynamic Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 4366–4373.
- Yang, S.; and Koepl, H. 2020. The Hawkes Edge Partition Model for Continuous-time Event-based Temporal Networks. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, 460–469.
- Yang, S.; and Zha, H. 2023. Estimating Latent Population Flows from Aggregated Data via Inverting Multi-Marginal Optimal Transport. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, 181–189.
- Zhang, Q.; Lipani, A.; Kirnap, O.; and Yilmaz, E. 2020. Self-Attentive Hawkes Process. In *Proceedings of the International Conference on Machine Learning (ICML)*, 11183–11193.

Zhang, Q.; Lipani, A.; and Yilmaz, E. 2021. Learning Neural Point Processes with Latent Graphs. In *Proceedings of the international conference on World Wide Web (WWW)*, 1495–1505. New York, NY, USA.

Zhang, Y.; and Yan, J. 2021. Neural Relation Inference for Multi-dimensional Temporal Point Processes via Message Passing Graph. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 3406–3412.

Zuo, S.; Jiang, H.; Li, Z.; Zhao, T.; and Zha, H. 2020. Transformer Hawkes Process. In *Proceedings of the International Conference on Machine Learning (ICML)*, 11692–11702.