

NEURAL MULTI-EVENT FORECASTING ON SPATIO-TEMPORAL POINT PROCESSES USING PROBABILISTICALLY ENRICHED TRANSFORMERS

Anonymous authors

Paper under double-blind review

ABSTRACT

Predicting discrete events in time and space has many scientific applications, such as predicting hazardous earthquakes and outbreaks of infectious diseases. History-dependent spatio-temporal Hawkes processes are often used to mathematically model these point events. However, previous approaches have faced numerous challenges, particularly when attempting to forecast one or multiple future events. In this work, we propose a new neural architecture for multi-event forecasting of spatio-temporal point processes, utilizing transformers, augmented with normalizing flows and probabilistic layers. Our network makes batched predictions of complex history-dependent spatio-temporal distributions of future discrete events, achieving state-of-the-art performance on a variety of benchmark datasets including the South California Earthquakes, Citibike, Covid-19, and Hawkes synthetic pinwheel datasets. More generally, we illustrate how our network can be applied to any dataset of discrete events with associated markers, even when no underlying physics is known.

1 INTRODUCTION

Predicting the occurrence of discrete events in time and space has been the focus of many scientific studies and applications. Problems such as predicting earthquake hazards (Ogata, 1998; Chen et al., 2020), infectious diseases over a population (Meyer et al., 2012; Schoenberg et al., 2019), mobility and traffic in cities (Du et al., 2016), and brain neuronal spikes (Perkel et al., 1967) fall under this category and have gained quite some interest. Over the years, many works used the history dependent spatio-temporal Hawkes process (Ozaki, 1979; Ogata, 1988; 1998; Ogata & Zhuang, 2006; Ogata et al., 1993; Nandan et al., 2017; Sornette & Werner, 2005; Zhuang, 2012; Helmstetter & Sornette, 2003; Bansal & Ogata, 2013) to model these point events. The stochasticity and the excitatory history-dependency of the Hawkes process are modeled by a conditional intensity function that varies in time and space. Assuming n events and their associated markers (t, \mathbf{x}, M) collected in the history $H_t = \{(t_i, \mathbf{x}_i, M_i) \mid t_i < t, i = 1 : n\}$, the Hawkes intensity function is defined as

$$\lambda(t, \mathbf{x} \mid H_t) := \lim_{\Delta t, \Delta \mathbf{x} \rightarrow 0} \frac{P_{\Delta t, \Delta \mathbf{x}}(t, \mathbf{x} \mid H_t)}{|B(\mathbf{x}, \Delta \mathbf{x})| \Delta t} = \mu_\theta(\mathbf{x}) + \sum_{i: t_i < t} g_\phi(t - t_i, \mathbf{x} - \mathbf{x}_i, M_i), \quad (1)$$

where $P_{\Delta t, \Delta \mathbf{x}}(t, \mathbf{x} \mid H_t)$ denotes the history-dependant probability of having an event in a small time interval $[t, t + \Delta t)$ and a small ball $B(\mathbf{x}, \Delta \mathbf{x})$ centered at $\mathbf{x} \in \mathbb{R}^d$ with the radius of $\Delta \mathbf{x} \in \mathbb{R}^d$, and H_t denotes all events happening up to but *not* including time t . The functions μ_θ and g_ϕ represent the (parametric) background intensity function and spatio-temporal kernel, respectively, forming $\lambda(t, \mathbf{x} \mid H_t)$. A variety of different parametric forms for both μ_θ and g_ϕ have been proposed (Ogata, 1998). Despite their use across different domains, Hawkes processes have several limitations. **First**, one has to predetermine parametric forms for μ_θ and g_ϕ . This limits the expressive power of the model and, more importantly, might lead to a mismatch between the proposed model and one that can better represent the observed data. **Second**, (1) imposes similar behavior across all events over time meaning that every preceding event has the same form of impact on future events' occurrence. As an example, in the classical Hawkes process with a temporal decaying kernel $g_\phi(t - t_i) = \frac{1}{\beta} \exp(-\frac{t-t_i}{\beta})$, the same parameter β is applied to all events via their corresponding time differences. **Third**, the intensity function $\lambda(t, \mathbf{x} \mid H_t)$ can be used to predict only one

step ahead via, e.g., the commonly-used first-order moment (Rasmussen, 2011; Snyder & Miller, 2012) given by

$$\mathbb{E}[t_{n+1}|H_{t_n}] = \int_{t_n}^{\infty} t \int_{\mathbf{V}} \lambda(t, \mathbf{x}|H_{t_n}) \exp\left(-\int_{\mathbf{V}} \int_{t_n}^t \lambda(u, \mathbf{v}|H_{t_n}) du d\mathbf{v}\right) d\mathbf{x} dt, \quad (2)$$

$$\mathbb{E}[\mathbf{x}_{n+1}|H_{t_n}] = \int_{\mathbf{V}} \mathbf{x} \int_{t_n}^{\infty} \lambda(t, \mathbf{x}|H_{t_n}) \exp\left(-\int_{\mathbf{V}} \int_{t_n}^t \lambda(u, \mathbf{v}|H_{t_n}) du d\mathbf{v}\right) dt d\mathbf{x}, \quad (3)$$

where solving the integral in a continuous high-dimensional space could be computationally expensive, or inaccurate in the case of using sampling methods such as Monte Carlo sampling (Hastings, 1970). **Fourth**, although (2) makes it feasible to predict the next event, the multi-event prediction demands sequential usage of the first-order moment, where at each step the history is updated via the most recent predicted event. This makes the problem of multi-event prediction challenging and prone to error accumulation.

To alleviate some of the mentioned shortcomings, previous works have proposed learning the intensity function $\lambda(t, \mathbf{x}|H_t)$ directly from data using neural networks. The early studies in this regard (Du et al., 2016; Mei & Eisner, 2017; Xiao et al., 2017) attempted to learn $\lambda(t, \mathbf{x}|H_t)$ using a recurrent neural network (RNN) and variants thereof. In a recent work, Chen et al. (2020) designed a neural ordinary differential equation (ODE) architecture that can learn a continuous intensity function. This is attained by combining jump and attentive continuous normalizing flows (CNFs) (Chen et al., 2018) for space, and an RNN architecture for time. Despite being data-driven, these works do not tackle the last three drawbacks mentioned above. Moreover, due to the memory limitations of RNNs, these works are incapable of unveiling long-term dependencies. On the other hand, RNNs and their variants suffer from vanishing and exploding gradients, especially when applied to long sequential data. Furthermore, sequential training of both RNNs and jumpCNFs as proposed by Chen et al. (2018) results in a slow and computationally expensive training process. Some other recent works (Zuo et al., 2020; Zhou et al., 2022), proposed using attention mechanisms (already present in the attentiveCNF in Chen et al. (2020)), to alleviate the second drawback mentioned earlier. The capability of discovering long-term dependencies while achieving fast training (via parallelization) is another benefit of this line of work. However, the last two drawbacks are still present in those approaches, hindering their utility for multi-event forecasting. Directly learning the temporal distribution using the variational autoencoders (VAEs) (Kingma & Welling, 2014) augmented with CNFs (Chen et al., 2018) as proposed by Mehra et al. (2019), or utilizing the deep sigmoidal flow (DSF) (Huang et al., 2018) and the sum of squares (SOS) polynomial flow (Jaini et al., 2019) as proposed by Shchur et al. (2019) could tackle the first and third limitations. However, the second and last mentioned shortcomings remained to be solved. In another recent work, Zhu et al. (2021) proposed learning the Hawkes intensity function based on Gaussian diffusion kernels while using imitation learning as a flexible model-fitting approach. However, their method could not tackle any of the aforementioned shortcomings. Therefore, in the present work, we introduce a novel neural network that is capable of *simultaneous* spatio-temporal multi-event forecasting while addressing all the above-mentioned shortcomings. To the best of our knowledge, this is the first work proposing a data-driven multi-event forecasting network for stochastic point processes. Our architecture augments the encoder and decoder blocks of a transformer architecture (Vaswani et al., 2017) from natural language processing (NLP) with probabilistic and bijective layers. This network design provides a batch of rich multi-event spatio-temporal distributions associated with multiple future events. We propose to directly learn/predict these distributions (without the need to learn the intensity function as opposed to the previous works), using self-supervised learning (Liu et al., 2021), marking a further point of departure with respect to existing work. The rest of this paper is organized as follows. In Section 2, we provide background knowledge on attention mechanisms and normalizing flows, since these are used as constituent blocks in our solution. Next, in Section 3 we provide a formal definition of the problem to be solved along with our proposed neural multi-event forecasting network. Further, in Section 4 we evaluate the performance of our network on a variety of datasets and compare it with baseline models. Finally, in Section 5 we finish the paper with conclusions and future goals. Our contributions can be summarized as follows:

- We introduce a neural architecture that is capable of simultaneous multi-event forecasting of the time and location of discrete events in continuous time and space.
- We compare the performance of our solution with state-of-the-art models through extensive experiments on a variety of datasets that represent stochastic point events.

2 BACKGROUND

2.1 ATTENTION AND TRANSFORMERS

Attention blocks play a critical role in the transformer architecture (Vaswani et al., 2017). Here we introduce their fundamental pieces, already specializing them to the problem at hand. Let $\{\kappa_i = [t_i, \mathbf{x}_i, M_i]\}_{i=1:n}$ be a sequence of (column) vectors $\kappa_i \in \mathbb{R}^{d+2}$ representing n discrete events formed by concatenating the associated time $t_i \in \mathbb{R}$ and markers $\mathbf{x}_i \in \mathbb{R}^d$ and $M_i \in \mathbb{R}$. The marker \mathbf{x}_i denotes the location of the event at time t_i in a d dimensional space, and M_i can be used to encode any other information of interest about the event, such as the magnitude of an earthquake or the biker’s age. Even though we represent here M_i as a scalar, these markers can also be higher dimensional. Moreover, the events κ_i are temporally sorted such that $t_i < t_j$ for $i < j$. Using learnable linear transformations, we form *query*, *key* and *value* vectors (denoted by \mathbf{q}_i , \mathbf{k}_i , and \mathbf{v}_i , respectively) for each event κ_i as

$$\mathbf{q}_i = \mathbf{W}_q \kappa_i, \quad \mathbf{k}_i = \mathbf{W}_k \kappa_i, \quad \mathbf{v}_i = \mathbf{W}_v \kappa_i, \quad (4)$$

where $\mathbf{W}_q \in \mathbb{R}^{d_k \times (d+2)}$, $\mathbf{W}_k \in \mathbb{R}^{d_k \times (d+2)}$, $\mathbf{W}_v \in \mathbb{R}^{d_v \times (d+2)}$. For a given event κ_i , we can build the matrix $\mathbf{K}_{(i)} = [\mathbf{k}_1 \mathbf{k}_2 \dots \mathbf{k}_i]$, which contains as columns the key vectors of all the events up to (and containing) event κ_i and, similarly, the matrix $\mathbf{V}_{(i)} = [\mathbf{v}_1 \mathbf{v}_2 \dots \mathbf{v}_i]$ values. Based on this notation, we compute the hidden representation \mathbf{h}_i of event κ_i as

$$\mathbf{h}_i = \mathbf{V}_{(i)} \text{softmax} \left(\frac{\mathbf{K}_{(i)}^\top \mathbf{q}_i}{\sqrt{d_k}} \right). \quad (5)$$

Note that in (5) we are generating \mathbf{h}_i by linearly combining the columns of $\mathbf{V}_{(i)}$, i.e., the values of all events preceding (and including) κ_i . The weights in this linear combination are given by a softmax applied to the inner products between the keys of every event and the query specific of κ_i . In this sense, \mathbf{h}_i contains a weighted combination of the values of all preceding events, where these values and the weights in the combination can be learned by tuning the transformations \mathbf{W}_q , \mathbf{W}_k , and \mathbf{W}_v .

2.2 NORMALIZING FLOWS

A normalizing flow is a transformation of a simple probability distribution to a complex one using a sequence of differentiable and invertible bijective maps (Kobyzev et al., 2020). Assume that $\kappa \in \mathbb{R}^{d_\kappa}$ has a complex unknown distribution p_κ . Having $\mathbf{z} \in \mathbb{R}^{d_\kappa}$ as a random sample drawn from a simple known distribution p_z (e.g., a normal distribution), we could use a sequence of invertible and differentiable functions $F = F_1 \circ F_2 \dots \circ F_k$ to map \mathbf{z} into κ according to $\kappa = F(\mathbf{z})$. Therefore, using the change of variables formula we have that

$$p_z(\mathbf{z}) = p_\kappa(F(\mathbf{z})) |\det(\mathbf{D}F(\mathbf{z}))|, \quad (6)$$

where $\det(\mathbf{D}F(\mathbf{z}))$ denotes the determinant of the Jacobian matrix of $F(\mathbf{z})$ and accounts for the volume change in the density as we move from the simple distribution to the complex one. Equation (6) makes it feasible to have an explicit form for the complex underlying distribution p_κ and we can effectively sample from it by generating samples from p_z and transforming those via F .

3 NEURAL MULTI-EVENT FORECASTING

Given a history of n events, we want to determine the probability distribution (in time and space) of the next L events. Formally, we want to solve the following problem.

Problem 1. Given the event history $H_{t_{n+1}} = \{(t_i, \mathbf{x}_i, M_i) \mid i = 1 : n\}$, estimate the probability distributions $\{p_l(t_l, \mathbf{x}_l \mid H_{t_l})\}_{l=n+1}^{n+L}$ of the next L events.

We refer to the above problem as *batched* or *multi-event* prediction because given the history up to event n we want to estimate the spatio-temporal probability densities of a *batch* of L events shown by p_l in Problem 1. Note that despite using the word *batched*, these distributions do not need to be independent of each other. We present here a data-driven solution of our problem where, during a training phase, we get to observe several sequences of length $n + L$ of our events of interest.

We propose a transformer-based architecture trained in a self-supervised manner to extract the relations between historical events. Furthermore, we enhance the transformer block with probabilistic and bijective layers to construct history-dependent spatio-temporal distributions from the hidden states of the transformer. Since $p_l(t_l, \mathbf{x}_l | H_{t_l}) = p_l(\mathbf{x}_l | t_l, H_{t_l})p_l(t_l | H_{t_l})$, we separately predict two batched history-dependent temporal *and* spatial distributions. Figure 1 shows a high-level view of our proposed network. Out of the $n + L$ events in the training sequences, n are used as the history information and the inputs to the encoder, whereas the remaining L are the inputs to the decoder in the training phase, and our goal (consistent with Problem 1) is to *predict* their spatio-temporal distributions in the test phase.

3.1 NETWORK ARCHITECTURE

As shown in Figure 1, the transformer consists of two main blocks known as the encoder and the decoder. The encoder is fed with input sequence batches, where each sequence has n events with their associated markers (t_i, \mathbf{x}_i, M_i) . Using multi-head attention blocks (Vaswani et al., 2017), we extract representations for all n events that capture history dependency. More precisely, via (5), we extract the hidden representation \mathbf{h}_{t_i} for time and $\mathbf{h}_{\mathbf{x}_i}$ for space, separately, which will be further used as part of the inputs to the decoder block.

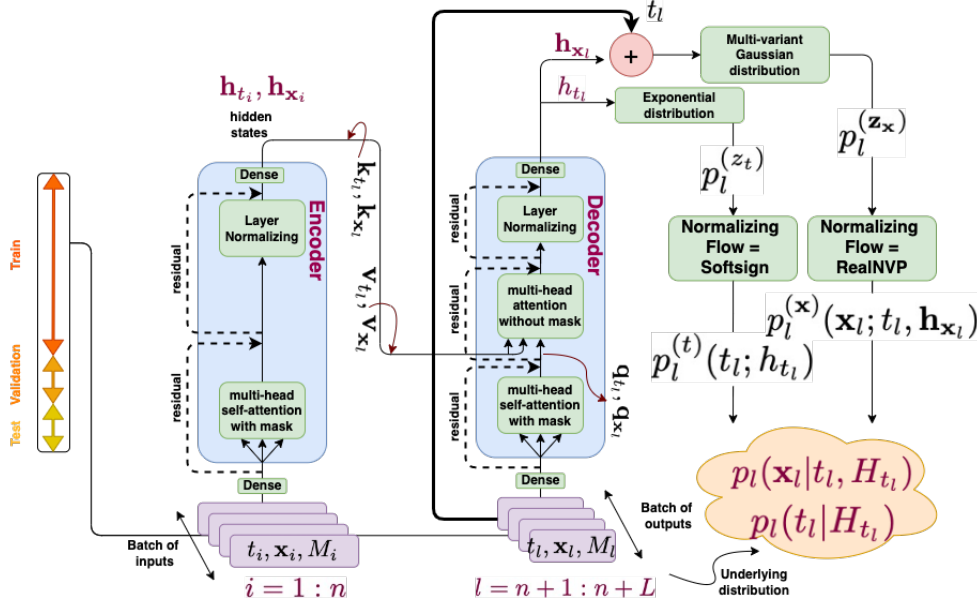


Figure 1: The framework of our proposed architecture. Our network is composed of an encoder, a decoder, and two parallel batched probabilistic layers each followed by bijective layers known normalizing flows. The input to the decoder is used to learn the output of the bijective layers using log-likelihood maximization in the training phase. In the test phase, a sequence of all zeros is input to the decoder and the output of the bijective layers are the predicted batched probability distributions, i.e., our proposed solution to Problem 1.

The decoder contains two attention blocks that extract hidden representations of the output data (events $n + 1$ through $n + L$) during the training phase. We denote these representations¹ by $\{h_{t_l}, \mathbf{h}_{\mathbf{x}_l}\}_{l=n+1}^{n+L}$. The first attention block in the decoder captures history dependency within the subsequence of events between time $n + 1$ and $n + L$. We denote the outputs of this first block by \mathbf{q}_{t_l} and $\mathbf{q}_{\mathbf{x}_l}$ since these are used as queries for the second multi-head attention block in the decoder; see Figure 1. The hidden states $\{\mathbf{h}_{t_i}, \mathbf{h}_{\mathbf{x}_i}\}_{i=1}^n$ from the encoder are used as both the key and the value inputs to the second attention block in the decoder layer. In this way, the history dependency between the input (from event 1 to n) and output (from event $n + 1$ to L) events is encoded in the output hidden representations of the decoder, i.e., in $\{h_{t_l}, \mathbf{h}_{\mathbf{x}_l}\}_{l=n+1}^{n+L}$. Despite the typical imple-

¹Notice that, as a convention, we tend to use subindex i (as in \mathbf{h}_{t_i}) to refer to the events in the interval $\{1, \dots, n\}$ and subindex l (as in h_{t_l}) for events in the interval $\{n + 1, \dots, n + L\}$. Also we use bold notation to show vector form of a variable, meaning h_{t_l} depicts 1 dimensional hidden states, whereas \mathbf{h}_{t_i} stands for dimensions of higher than 1.

mentation of the transformer decoder using look-ahead masks, due to our multi-event objective we remove this mask in the second attention block of the decoder. It is worth mentioning that, during testing, the time and markers for the events $n + 1$ through $n + L$ are input as zero, so that no information is used for prediction beyond $H_{t_{n+1}}$, thus abiding by the constraint of Problem 1. Finally, normalizing, residual, and dense layers are used in both the encoder and the decoder blocks for dimension matching and enhancing the training.

In our next step, we separately inject $\{h_{t_l}\}_{l=n+1}^{n+L}$ and $\{\mathbf{h}_{\mathbf{x}_l}\}_{l=n+1}^{n+L}$ into two separate batched probabilistic layers (Dürr et al., 2020). In essence, for the time variable t_l , we define a trainable map going from h_{t_l} into the parameter of an exponential distribution. We denote this distribution by $p_l^{(z_t)}(\cdot; h_{t_l})$. Similarly, for the spatial variable \mathbf{x}_l , we define a trainable map going from $(\mathbf{h}_{\mathbf{x}_l}, t_l)$ into the parameters of a multivariate Gaussian distribution. We denote this distribution by $p_l^{(\mathbf{z}_x)}(\cdot; t_l, \mathbf{h}_{\mathbf{x}_l})$. We learned the parameters of these distributions independently for every $l \in \{n + 1, \dots, n + L\}$. However, it should be noticed that the hidden states $\{h_{t_l}\}_{l=n+1}^{n+L}$ and $\{\mathbf{h}_{\mathbf{x}_l}\}_{l=n+1}^{n+L}$ already encode spatio-temporal dependencies, so that the parameters learned for the probability distributions of different events already capture the correlations with other events across time. Also notice that the parameters of the spatial distribution depend on the time instant t_l since we want to estimate the joint probability density over time and space (see Problem 1) as $p_l(t_l, \mathbf{x}_l | H_{t_l}) = p_l(\mathbf{x}_l | t_l, H_{t_l})p_l(t_l | H_{t_l})$.²

Although the probabilistic layers have a critical role in forming the underlying batched distributions, they can only take simple and tractable forms (in our case, exponential and multivariate Gaussian), which are not expressive enough to capture the true underlying distributions of the times and locations of future events. Thus, a set of bijective layers, known as normalizing flows (Kobyzev et al., 2020), are used to transform them into richer and more complex batched forms. As discussed in Section 2.2, we can use (6) to address the transformation of distributions $p_l^{(z_t)}$ and $p_l^{(\mathbf{z}_x)}$ into more complex forms. We use the exponential and multivariate normal distributions as the base distributions $p_l^{(z_t)}$ and $p_l^{(\mathbf{z}_x)}$ because they have been used extensively as parametric kernels to form the Hawkes intensity function in (1) (Ogata, 1998; 1988). Moreover, in Appendix A.7 we have shown that the histograms depicting the space and time distributions in various datasets are not too far from being a multivariate Gaussian distribution for space and an exponential for consecutive time intervals. In this sense, normalizing flows only need to learn to further adjust these distributions to the observed data. More precisely, we find continuous and invertible functions F_1 and F_2 such that, for every l , these functions can transform samples from $p_l^{(z_t)}$ and $p_l^{(\mathbf{z}_x)}$ (denoted by z_{t_l} and $\mathbf{z}_{\mathbf{x}_l}$) into t_l and \mathbf{x}_l as in

$$t_l = F_1(z_{t_l}), \quad \mathbf{x}_l = F_2(\mathbf{z}_{\mathbf{x}_l}). \quad (7)$$

Using (6) we can relate the known expressions of the base distributions with the distributions of t_l and \mathbf{x}_l using the Jacobians of F_1 and F_2 , respectively. We denote these learned distributions by $p_l^{(t)}(\cdot; h_{t_l})$ and $p_l^{(\mathbf{x})}(\cdot; t_l, \mathbf{h}_{\mathbf{x}_l})$. In terms of parametric functional forms, we use a softsign bijector for F_1 and a RealNVP bijector (Dinh et al., 2016) for F_2 . To train these bijections, we seek to maximize the log-likelihood of the true events $\{t_l, \mathbf{x}_l\}_{l=n+1}^{n+L}$ under the learned distributions $p_l^{(t)}(\cdot; h_{t_l})$ and $p_l^{(\mathbf{x})}(\cdot; t_l, \mathbf{h}_{\mathbf{x}_l})$. More details about the training of our network will be discussed in Section 4.3 and Appendix A.2. Going back to Problem 1, by maximizing the log-likelihood we are training our network in an end-to-end fashion such that $p_l^{(t)}(t_l; h_{t_l}) \approx p_l(t_l | H_{t_l})$ and $p_l^{(\mathbf{x})}(\mathbf{x}_l; t_l, \mathbf{h}_{\mathbf{x}_l}) \approx p_l(\mathbf{x}_l | t_l, H_{t_l})$, where, we recall, H_{t_l} contains the history of events up to (but not including) event l . From here we can obtain

$$p_l(t_l, \mathbf{x}_l | H_{t_l}) = p_l(\mathbf{x}_l | t_l, H_{t_l})p_l(t_l | H_{t_l}), \quad \text{for } l = n + 1, \dots, n + L. \quad (8)$$

Finally, be recalling that during testing every event beyond n is completed via zero padding, in (8) we are effectively obtaining a distribution for time and space of every event l in the range $n + 1$ to $n + L$ given $H_{t_{n+1}}$, as we aimed for in Problem 1.

²Note that during the test phase we only feed the decoder with $t_l = 0$ when modeling $p_l^{(z_t)}(\cdot; h_{t_l})$; However, we use the true known time values $t_{l_{\text{known}}}$ as inputs to the probabilistic layers when modeling $p_l^{(\mathbf{z}_x)}(\cdot; t_l, \mathbf{h}_{\mathbf{x}_l})$ in both train and test phases. More details are provided in Appendix A.2.

3.2 ADDRESSING THE LIMITATIONS OF EXISTING APPROACHES

As previewed in Section 1, our architecture in Figure 1 was inspired by the need of addressing four shortcomings present in classical Hawkes processes and partially shared by more modern alternatives. Here, we explicitly discuss how the different constituent blocks in our architecture address these deficiencies. **First**, we do not depend on prespecified kernels to model the probability densities of time and location. Notice that, although we adopt exponential and Gaussian distributions for the latent variables z_{t_l} and $\mathbf{z}_{\mathbf{x}_l}$ as further described in Appendix A.1, these are later transformed via normalizing flows to better represent the observed data. **Second**, the attention mechanisms in both the encoder and decoder blocks of the transformer can learn heterogeneous effects between events across time. **Third**, the explicit forms of $p_i^{(t)}(t_l; h_{t_l})$ and $p_i^{(\mathbf{x})}(\mathbf{x}_l; t_l, \mathbf{h}_{\mathbf{x}_l})$ obtained via normalizing flows facilitate the fast sampling of future events via (8) without having to rely on integration as in (2). **Fourth**, the incorporation of the decoder block of the transformer enables multi-event prediction bypassing the need of sequential one-by-one event prediction. To be more precise, the use of our decoder block implies that, *during training*, we are incorporating the future information $\{h_{t_l}, \mathbf{h}_{\mathbf{x}_l}\}_{l=n+1}^{n+L}$ along with their dependence on the encoded history $\{h_{t_i}, \mathbf{h}_{\mathbf{x}_i}\}_{i=1}^n$. This makes our training procedure self-supervised rather than unsupervised, unlike all previous works in the area. In essence, during training, the network learns to fit the distribution of multiple future events jointly. Hence, during testing, we can zero-pad the input to the decoder and still obtain an estimation of the spatio-temporal probability densities of several events in the future.

4 EXPERIMENTS

4.1 DATA

We used a variety of synthetic and real-world datasets representing point processes with discrete events. We briefly introduce these datasets here; more details can be found in Appendix A.6. In all datasets, we collected sequences of 500 events with an overlap of 498 events. We used $n = 497$ events in each sequence as the inputs and the final $L = 3$ events as the outputs that we want to predict during the testing phase. Train, validation, and test sets are formed according to the 80%–14%–6% split rule after shuffling the formed sequences.

South California Earthquakes. Earthquake events from 2008 to 2016 in South California (Ross et al., 2019) of magnitude at least 2.5. The event description includes time, location in 3 dimensions, magnitude and consecutive events’ time intervals as event markers.

Citibike. Rental events from a bike sharing service in New York City (Amazonawstripdata). The event description includes starting time of the biking, location in 2 dimensions, the biker’s birth year, and consecutive rentals’ time intervals as event markers.

Covid-19. Daily Covid-19 cases in different states of the United States (The-NewYork-Times). The event description includes day of catching Covid-19, location in 2 dimensions, the number of cases on that day, and consecutive events’ time intervals as event markers.

Pinwheel. Hawkes pinwheel dataset introduced in Chen et al. (2020). Hawkes time instances were simulated using the thinning algorithm introduced in Ogata (1981) and assigned to a cluster-based pinwheel distribution. For simplicity, we assigned the same magnitude to all spatial points sampled from the formed distribution when forming the data sequences.

4.2 BASELINES

We consider two categories of baseline methods, namely, those that try to predict the time and those that try to predict the location of the next event. Time category models are used to learn the intensity function from input times and we predict the expected time of multiple events in the future by sequentially utilizing (2). These models include the Hawkes process, self-correcting point process (Isham & Westcott, 1979), and the homogeneous Poisson process (Pasupathy, 2010). Note that the sequential prediction of the future time-points for more than one event, in cases of the Hawkes and self-correcting intensities, relies on a redefined history that incorporates the very last predicted event time in the previous step. We also use space models to sequentially learn the space distributions of multiple events in the future. In this regard, we use the conditional multi-variate Gaussian mixture model (GMM) (Murphy, 2012; Bishop & Nasrabadi, 2006) to learn the space distribution, where the Gaussian kernel parameters are learned from the historical events.

4.3 NETWORK TRAIN AND TEST

Training starts by passing batches of input and output sequences to the encoder and the decoder, respectively. The encoder outputs are also passed to the decoder as earlier explained in Section 3.1. We used inputs and outputs with all the existing markers for each dataset as mentioned in Section 4.1. A list of the hyperparameters used in this work is shown in Table 3 in Appendix A.7. We define the loss as minimizing the negative log-likelihood

$$\text{loss} = - \sum_{l=n+1}^{n+L} \log p_l^{(t)}(t_l; h_{t_l}) - \sum_{l=n+1}^{n+L} \log p_l^{(x)}(\mathbf{x}_l; t_l, \mathbf{h}_{\mathbf{x}_l}). \quad (9)$$

Thus, we are effectively learning the coefficients in the transformer blocks – which affect the hidden representations h_{t_l} and $\mathbf{h}_{\mathbf{x}_l}$ – and the coefficients in the probabilistic and bijective layers – which parameterize $p_l^{(t)}$ and $p_l^{(x)}$ – to maximize the likelihood of generating the observed data.

The baseline methods are trained to encode the history of past events and to sequentially predict the expected next event based on (2). Therefore, in each step we compute the expected time \hat{t}_l and location $\hat{\mathbf{x}}_l$ of event l and we want these to be close to the true values t_l and \mathbf{x}_l before updating the history for the next prediction step. To satisfy this condition, during the training phase, we add regularizers to the negative log-likelihood as in

$$\text{loss} = - \sum_{l=1}^{n+L} \log p_t(t_l|H_{t_l}) - \sum_{l=1}^{n+L} \log p_x(\mathbf{x}_l|t_l, H_{t_l}) + \lambda_1 \sum_{l=n+1}^{n+L} |t_l - \hat{t}_l| + \lambda_2 \sum_{l=n+1}^{n+L} \|\mathbf{x}_l - \hat{\mathbf{x}}_l\|_2, \quad (10)$$

where $p_t(\cdot|H_{t_l})$ and $p_x(\cdot|t_l, H_{t_l})$ are the learned probability densities for the baseline method at hand, and λ_1 and λ_2 encode the relative weights of the regularizers. To be more precise, if a method only predicts the time, we use only the first and third terms in (10) as the loss and if a method only predicts the location we use the second and fourth terms as the loss. For a fair comparison between the loss of our network vs. the baselines, during the testing phase we report the negative log-likelihood for events $n+1$ through $n+L$ for all the tested models. We train our network using the Adam optimizer with an initial learning rate of 0.001 and a scheduled learning rate rule. We used dropout layers to reduce overfitting and ran 1000 epochs using batches of size 32.³

Table 1: Results representing the negative log-likelihood for the events in the range $n+1$ to $n+L$ for the learned distributions (less is better). \pm indicate the standard deviation of the loss among all test batch sequences.

	Earthquake	Citibike	Covid-19	Pinwheel
Homo-Poisson ($p_l^{(t)}$)	6.91 \pm 0.0	6.83 \pm 0.0	6.32 \pm 0.0	6.82 \pm 0.0
Hawkes ($p_l^{(t)}$)	0.42 \pm 0.09	-0.53 \pm 0.05	-3.79 \pm 0.01	-3.98\pm1.54
Self-correcting ($p_l^{(t)}$)	NAN	1.95 \pm 0.11	4.6 \pm 0.02	2.03 \pm 0.73
Conditional-GMM ($p_l^{(x)}$)	4.88 \pm 0.29	3.05 \pm 1.26	3.48 \pm 1.37	4.38 \pm 0.42
Our network ($p_l^{(t)}$)	0.33\pm0.17	-6.279\pm0.123	-5.87\pm0.09	1.27 \pm 0.26
Our network ($p_l^{(x)}$)	2.03\pm0.26	-3.28\pm0.14	2.18\pm0.21	1.57\pm0.09
Our network ($p_l^{(t,x)}$)	2.36\pm0.36	-9.56\pm0.18	-3.7\pm0.16	2.84\pm0.19

4.4 RESULTS

Figure 2 shows the predicted batched spatial and temporal distributions for the next three events on the South California Earthquake data (Ross et al., 2019) using our proposed architecture. On the top right figure, we took the average of 1000 samples taken from the predicted batched history dependent time distributions to forecast the time of future events shown by the red stars. The linear

³Code is available as supplementary material and will be publicly released in github if the paper is accepted.

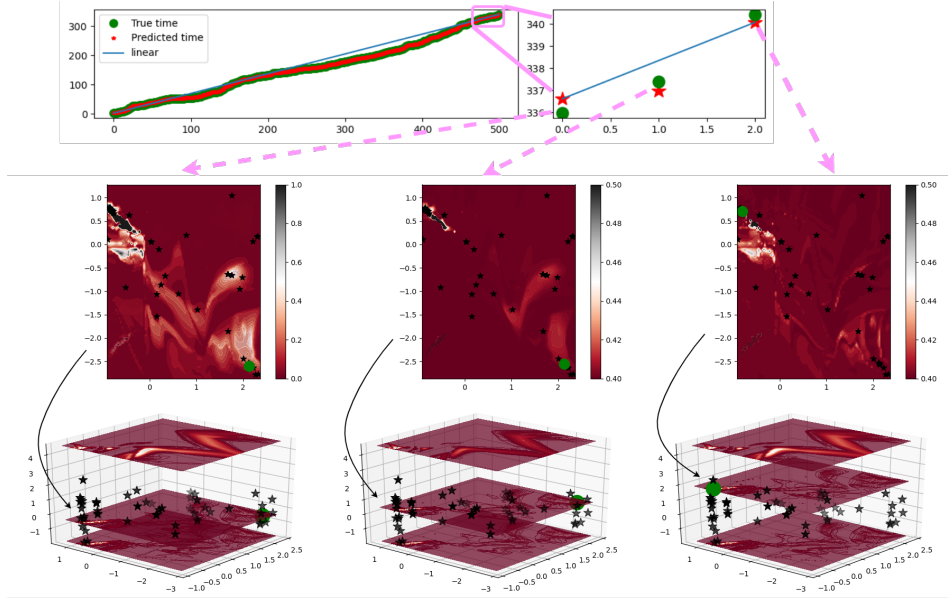


Figure 2: Visualizations of the learned history-dependent batched spatial and temporal distributions for the last three events of an unseen earthquake sequence using **our proposed architecture**. The red stars on the **top** figures are the simultaneous prediction of the true events’ time, shown by green circles, where we took an average between 1000 samples taken from the output predicted time distributions. The **bottom** figures depict the predicted batched spatial distribution associated with the last three events in 3 dimensions where the z-axis is along events’ depth. The 2-d representation of the last three events’ predicted spatial distributions mapped to the events’ depth are shown on the **middle** plots where except for the one on the left, we show the differences in the predicted distributions. The black stars on the spatial distributions represent the last 50 events of the input sequence as the history information.

trend is shown to emphasize that our network is capable of predicting time events with different in-between time-intervals. The bottom figures show the 3-dimensional representation of batched history-dependent multi-modal predicted densities along the main faults, prone to earthquake occurrence. In the 3-d visualization we show three slices along the z-axis, representing the lowest, highest, and the depth associated with the true future event, shown by the green circle. In the middle, we show the 2-dimensional density visualizations that are associated with the true event’s depth shown in the 3-d figures below them. Note that, while the bottom figures represent the predicted densities, in the middle figures we have the predicted density on the left (i.e., for the first predicted event) and the other two middle figures depict the consecutive predicted *density differences* for the two subsequent events. This is done to highlight the performance of our network in predicting simultaneous but different multi-event densities. Each predicted space density is associated with the assigned events’ predicted time shown above it. Despite seeing only slight differences between the predicted space densities, regions that are most expected to have the next event are still well-recognized. Note that while the same history is encoded for all the three future events, the differences between the predicted batched densities are highly influenced by the decoder layer which incorporates future information that was available during the training phase. Considering that consecutive outputs in most sequences occur very close to each other in time and space, the decoder layer is trained to represent the same behaviour in the test phase, predicting batched densities with slight but informative differences. The high bias on the upper left of the space density maps is caused by the highly earthquake-prone regions/faults in South California.

Figure 3 on the other hand, represents the sequential prediction performance of the baseline models on the same dataset used in Figure 2. The Hawkes model and the conditional GMM model are used to learn the time and space distributions, respectively, where we used (2) to sequentially predict the time and location of the next event. Comparing this result with Figure 2 we see the better performance of our network in predicting the time and space distribution of future events on earthquake data. Moreover, the conditional GMM model is incapable of predicting a complex and well-defined density along fault lines representing earthquake-prone regions. More details on how we perform

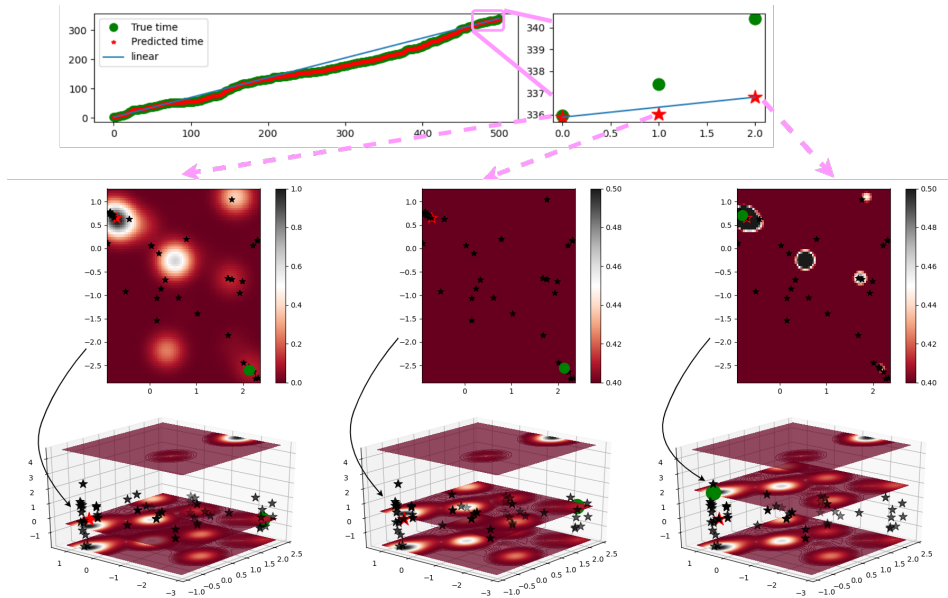


Figure 3: Same visualizations as Figure 2 for the same data using **Hawkes model** for time, augmented with **conditional-GMM** for space. Here the prediction in time and space is performed sequentially using (2) by updating the history at each prediction step.

sequential prediction is given in Appendix A.5. The black stars on spatial density plots in Figures 2 and 3 show the last 50 out of the 497 historical events that are used as the encoder inputs.

The results shown in Table 1 for all the experiments indicate that, except for the Hawkes time pinwheel synthetic dataset, our network outperforms all the baseline models. Due to a very high error accumulation associated with the sequential prediction using the conditional GMM model, we used the true rather than the predicted values at each prediction step to update the history. More visualization results are given in Appendix A.4. Since the times associated with the pinwheel dataset are synthetically simulated from the Hawkes model, Hawkes is a better candidate for forecasting the expected time distribution for this dataset as shown in Table 1. Among the given baseline models, the self-correcting and homogeneous Poisson are performing the worst, where the self-correcting with the inhibitory assumption behind history dependency never converged on the earthquake dataset. While all the provided baseline models assume certain parametric forms for the space or time distributions, our proposed network is applicable to any discrete type dataset even in the absence of underlying physics or known models that can guide the right choice of a spatio-temporal kernel.

5 CONCLUSION

We have proposed a novel probabilistic approach by enhancing a transformer architecture to conditionally learn the spatio-temporal distribution of multiple stochastic discrete events. Our network combines the transformer model and normalizing flows augmented with batched probabilistic layers to simultaneously learn the underlying distributions of multiple future events using self-supervised learning. The attention blocks of the transformer architecture assign score-based history dependency among events, which can capture both excitatory and inhibitory behaviors while being parallel and fast to train. We show that our approach achieves state-of-the-art performance on different spatio-temporal datasets collected from a wide range of domains.

For future work, we are interested in using other available sources of data as proposed by the work of Okawa et al. (2022) such as GPS and beneath-earth images to propose a context-aware batched spatio-temporal distribution forecasting tool, enhancing the performance of our network on the stochastic point processes. Another interesting direction is to use the diffusion model (Song et al., 2020; Ho et al., 2020; Nichol & Dhariwal, 2021) instead of normalizing flows as generative models of learning the probability distributions. We hope that by adding gradual noise in the reverse diffusion denoising steps, we might better learn regions associated with rare events.

REFERENCES

- Amazonawstripdata. Citibike data in the united states, 2021. URL <https://s3.amazonaws.com/tripdata/index.html>.
- AR Bansal and Y Ogata. A non-stationary epidemic type aftershock sequence model for seismicity prior to the december 26, 2004 m 9.1 sumatra-andaman islands mega-earthquake. *Journal of Geophysical Research: Solid Earth*, 118(2):616–629, 2013.
- Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- Feng Chen and Thomas Stindl. *RHawkes: Renewal Hawkes Process*, 2022. URL <https://cran.r-project.org/package=RHawkes>. R package version 1.0.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. *arXiv preprint arXiv:1806.07366*, 2018.
- Ricky TQ Chen, Brandon Amos, and Maximilian Nickel. Neural spatio-temporal point processes. *arXiv preprint arXiv:2011.04583*, 2020.
- Joshua V Dillon, Ian Langmore, Dustin Tran, Eugene Brevdo, Srinivas Vasudevan, Dave Moore, Brian Patton, Alex Alemi, Matt Hoffman, and Rif A Saurous. Tensorflow distributions. *arXiv preprint arXiv:1711.10604*, 2017.
- Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- 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 *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1555–1564, 2016.
- Oliver Dürr, Beate Sick, and Elvis Murina. *Probabilistic Deep Learning: With Python, Keras and TensorFlow Probability*. Manning Publications, 2020.
- Beno Gutenberg and Charles F Richter. Frequency of earthquakes in california. *Bulletin of the Seismological society of America*, 34(4):185–188, 1944.
- W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970.
- Agnes Helmstetter and Didier Sornette. Predictability in the epidemic-type aftershock sequence model of interacting triggered seismicity. *Journal of Geophysical Research: Solid Earth*, 108 (B10), 2003.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- Chin-Wei Huang, David Krueger, Alexandre Lacoste, and Aaron Courville. Neural autoregressive flows. In *International Conference on Machine Learning*, pp. 2078–2087. PMLR, 2018.
- Valerie Isham and Mark Westcott. A self-correcting point process. *Stochastic processes and their applications*, 8(3):335–347, 1979.
- Priyank Jaini, Kira A Selby, and Yaoliang Yu. Sum-of-squares polynomial flow. In *International Conference on Machine Learning*, pp. 3009–3018. PMLR, 2019.
- DP Kingma and M Welling. Auto-encoding variational bayes. in international conference on learning representations (iclr). 2014.
- Leon Knopoff. The magnitude distribution of declustered earthquakes in southern california. *Proceedings of the National Academy of Sciences*, 97(22):11880–11884, 2000.

- Ivan Kobyzev, Simon JD Prince, and Marcus A Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3964–3979, 2020.
- Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- Nazanin Mehrasa, Ruizhi Deng, Mohamed Osama Ahmed, Bo Chang, Jiawei He, Thibaut Durand, Marcus Brubaker, and Greg Mori. Point process flows. *arXiv preprint arXiv:1910.08281*, 2019.
- Hongyuan Mei and Jason M Eisner. The neural hawkes process: A neurally self-modulating multivariate point process. *Advances in neural information processing systems*, 30, 2017.
- Sebastian Meyer, Johannes Elias, and Michael Höhle. A space–time conditional intensity model for invasive meningococcal disease occurrence. *Biometrics*, 68(2):607–616, 2012.
- Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- Shyam Nandan, Guy Ouillon, Stefan Wiemer, and Didier Sornette. Objective estimation of spatially variable parameters of epidemic type aftershock sequence model: Application to California. *Journal of Geophysical Research: Solid Earth*, 122(7):5118–5143, 2017.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pp. 8162–8171. PMLR, 2021.
- Yosihiko Ogata. On lewis’ simulation method for point processes. *IEEE transactions on information theory*, 27(1):23–31, 1981.
- Yosihiko Ogata. Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical association*, 83(401):9–27, 1988.
- Yosihiko Ogata. Space-time point-process models for earthquake occurrences. *Annals of the Institute of Statistical Mathematics*, 50(2):379–402, 1998.
- Yosihiko Ogata and Jiancang Zhuang. Space–time etas models and an improved extension. *Tectonophysics*, 413(1-2):13–23, 2006.
- Yosihiko Ogata, Ritsuko S Matsu’ura, and Koichi Katsura. Fast likelihood computation of epidemic type aftershock–sequence model. *Geophysical research letters*, 20(19):2143–2146, 1993.
- Maya Okawa, Tomoharu Iwata, Yusuke Tanaka, Takeshi Kurashima, Hiroyuki Toda, and Hisashi Kashima. Context-aware spatio-temporal event prediction via convolutional hawkes processes. *Machine Learning*, pp. 1–22, 2022.
- Tohru Ozaki. Maximum likelihood estimation of hawkes’ self-exciting point processes. *Annals of the Institute of Statistical Mathematics*, 31(1):145–155, 1979.
- George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. *Advances in neural information processing systems*, 30, 2017.
- Raghu Pasupathy. Generating homogeneous poisson processes. *Wiley encyclopedia of operations research and management science*, 2010.
- Donald H Perkel, George L Gerstein, and George P Moore. Neuronal spike trains and stochastic point processes: Ii. simultaneous spike trains. *Biophysical journal*, 7(4):419–440, 1967.
- Jakob Gulddahl Rasmussen. Temporal point processes: the conditional intensity function. *Lecture Notes, Jan*, 2011.
- DA Rhoades. Estimation of the gutenbergrichter relation allowing for individual earthquake magnitude uncertainties. *Tectonophysics*, 258(1-4):71–83, 1996.
- Zachary E Ross, Daniel T Trugman, Egill Hauksson, and Peter M Shearer. Searching for hidden earthquakes in southern California. *Science*, 364(6442):767–771, 2019.

- Frederic Paik Schoenberg, Marc Hoffmann, and Ryan J Harrigan. A recursive point process model for infectious diseases. *Annals of the Institute of Statistical Mathematics*, 71(5):1271–1287, 2019.
- Oleksandr Shchur, Marin Biloš, and Stephan Günnemann. Intensity-free learning of temporal point processes. *arXiv preprint arXiv:1909.12127*, 2019.
- Donald L Snyder and Michael I Miller. *Random point processes in time and space*. Springer Science & Business Media, 2012.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- Didier Sornette and Maximilian J Werner. Constraints on the size of the smallest triggering earthquake from the epidemic-type aftershock sequence model, bâth’s law, and observed aftershock sequences. *Journal of Geophysical Research: Solid Earth*, 110(B8), 2005.
- The-NewYork-Times. Coronavirus (covid-19) data in the united states, 2021. URL <https://github.com/nytimes/covid-19-data>.
- Stefano Tinti and Francesco Mulargia. Effects of magnitude uncertainties on estimating the parameters in the gutenbergrichter frequency-magnitude law. *Bulletin of the Seismological Society of America*, 75(6):1681–1697, 1985.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- Shuai Xiao, Junchi Yan, Xiaokang Yang, Hongyuan Zha, and Stephen Chu. Modeling the intensity function of point process via recurrent neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- Zihao Zhou, Xingyi Yang, Ryan Rossi, Handong Zhao, and Rose Yu. Neural point process for learning spatiotemporal event dynamics. In *Learning for Dynamics and Control Conference*, pp. 777–789. PMLR, 2022.
- Shixiang Zhu, Shuang Li, Zhigang Peng, and Yao Xie. Imitation learning of neural spatio-temporal point processes. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- Jiancang Zhuang. Long-term earthquake forecasts based on the epidemic-type aftershock sequence (etas) model for short-term clustering. *Research in Geophysics*, 2(1):e8–e8, 2012.
- Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. Transformer hawkes process. In *International conference on machine learning*, pp. 11692–11702. PMLR, 2020.

A APPENDIX

A.1 FURTHER DETAILS ON NETWORK ARCHITECTURE

In Section 3.1 we showed the usage of two individual batched probabilistic layers that are separately fed with $\{h_{t_l}\}_{l=n+1}^{n+L}$ and $\{(\mathbf{h}_{\mathbf{x}_l}, t_l)\}_{l=n+1}^{n+L}$, where $\{h_{t_l}, \mathbf{h}_{\mathbf{x}_l}\}_{l=n+1}^{n+L}$ are the outputs of the decoder block and $\{t_l\}_{l=n+1}^{n+L}$ are the known time values. We denoted the outputs of these layers by $p_l^{(z_t)}(\cdot; h_{t_l})$ and $p_l^{(z_{\mathbf{x}})}(\cdot; t_l, \mathbf{h}_{\mathbf{x}_l})$ as the conditional exponential and multivariate Gaussian distributions, respectively, whose parameters are learned from $\{h_{t_l}\}_{l=n+1}^{n+L}$ and $\{(\mathbf{h}_{\mathbf{x}_l}, t_l)\}_{l=n+1}^{n+L}$. To further expand the mathematical description of these layers’ outputs we have

$$\begin{aligned}
p_l^{(z_t)}(\cdot; h_{t_l}) &= \frac{1}{\beta_l} \exp\left(-\frac{z_{t_l}}{\beta_l}\right), \\
p_l^{(z_x)}(\cdot; t_l, \mathbf{h}_{x_l}) &= ((2\pi)^{n_{z_x}} |\Sigma_l|)^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\mathbf{z}_{x_l} - \mu_l)^\top \Sigma_l^{-1} (\mathbf{z}_{x_l} - \mu_l)\right),
\end{aligned} \tag{11}$$

where β_l is learned as a function of h_{t_l} , and $\{\Sigma_l, \mu_l\}$ are learned as functions of $\{t_l, \mathbf{h}_{x_l}\}$. $p_l^{(z_t)}(\cdot; h_{t_l})$ and $p_l^{(z_x)}(\cdot; t_l, \mathbf{h}_{x_l})$ in (11) are individually followed by a softsign bijector F_1 and a RealNVP bijector F_2 (Dinh et al., 2016), respectively, as two separate bijective layers to model the desired batched conditional temporal and spatial distributions $p_l^{(t)}(\cdot; h_{t_l})$ and $p_l^{(x)}(\cdot; t_l, \mathbf{h}_{x_l})$. Using (6), we have

$$\begin{aligned}
p_l^{(t)}(\cdot; h_{t_l}) &= p_l^{(z_t)}(\cdot; h_{t_l}) |\det(\mathbf{D}F_1(z_{t_l}))|^{-1}, \\
p_l^{(x)}(\cdot; t_l, \mathbf{h}_{x_l}) &= p_l^{(z_x)}(\cdot; t_l, \mathbf{h}_{x_l}) |\det(\mathbf{D}F_2(\mathbf{z}_{x_l}))|^{-1}.
\end{aligned} \tag{12}$$

Notice that the usage of the softsign bijective function $F_1(y) = \frac{y}{|y|+1}$ after the exponential probabilistic layer restricts the outputs between 0 and 1, i.e. $F_1 : (0, \infty) \rightarrow (0, 1)$. This is consistent with the fact that we always normalize our time inputs $\{t_i, t_l\}$ to stay between 0 and 1, so that our time intervals also always stay between 0 and 1. Therefore, this bijective map (Dillon et al., 2017) will ensure that the outputs from the exponential probabilistic layer will stay in our desired range. Moreover, as shown in the time-interval histograms associated with all datasets in Figure 12, all follow an approximate exponential behavior. This indicates that we do not require much flexibility in choosing the temporal bijective layer. Also, note that since time is 1-dimensional, it is not applicable to utilize the well-known bijective functions, such as the RealNVP (Dinh et al., 2016), NICE Dinh et al. (2014), the Masked Autoregressive normalizing flow (Papamakarios et al., 2017), to model the batched temporal distribution, mainly because these flow architectures are designed for dimensions higher than 1.

A.2 DETAILS ON NETWORK TRAIN AND TEST

As pointed out earlier, the training and testing phases of our network come with slight differences which we better clarify in this section. As shown in Figure 1, during training, both batches of inputs $\{(t_i, \mathbf{x}_i, M_i)\}_{i=1}^n$ and outputs $\{(t_l, \mathbf{x}_l, M_l)\}_{l=n+1}^{n+L}$ are fed into the encoder and the decoder, respectively. Using (9), the goal is to learn a parametric description of a batched probability distribution that seeks to explain future events in only time and space, meaning $p_l^{(t)}$ and $p_l^{(x)}$, respectively. The reason behind considering only these markers (time and space) is that these are the main markers covered by the original Hawkes process as shown in (1). Learning batched distributions associated with other markers may need the contribution of models other than Hawkes. As an example, the Gutenberg-Richter (GR) frequency-magnitude law in earth science (Gutenberg & Richter, 1944; Knopoff, 2000; Tinti & Mulargia, 1985; Rhoades, 1996) is needed to be considered when modeling/learning the batched magnitude distribution p_{M_l} associated with future earthquake events. Also notice that, as indicated by the Hawkes intensity function in (1), using other markers M as historical information has no contradiction with learning only the batched spatio-temporal distributions.

On the other hand, during testing, we only use the input batches $\{(t_i, \mathbf{x}_i, M_i)\}_{i=1}^n$ to feed the encoder, whereas a batch of zeros as outputs, meaning $\{(t_l, \mathbf{x}_l, M_l)\}_{l=n+1}^{n+L} = 0$, are used to feed the decoder. We expect that in the testing phase, the trained network could predict the batched spatio-temporal distributions $p_l^{(t)}$ and $p_l^{(x)}$, associated with the future L events, by only relying on the historical information $\{(t_i, \mathbf{x}_i, M_i)\}_{i=1}^n$. This might cause confusion about passing zero queries to the second attention layer of the decoder, due to using zero inputs. However, this can be addressed by relying on the biases associated with the decoder block, which were tuned in the training phase. Another question that might be raised due to feeding zero inputs to the decoder in the testing phase, is how the decoder can distinguish between different inputs. Notice that we used positional embedding (PE) proposed by Vaswani et al. (2017) in both the training and testing phases. However, the decoder is also fed by the output from the encoder, which is computed during testing using the true history $\{(t_i, \mathbf{x}_i, M_i)\}_{i=1}^n$ that we use for prediction. Consequently, the decoder can distinguish between different event histories via the output of the decoder in addition to having the preserved positions using PE.

As illustrated in Figure 1, we designed our model to jointly predict both time and location distributions by first drawing a time sample \hat{t}_l from the learned $p_l^{(t)}$ and then fixing and using \hat{t}_l as an input to the batched spatial probabilistic layer, to model/predict $p_l^{(x)}(\cdot; \hat{t}_l, \mathbf{h}_{\mathbf{x}_l})$. Since our baselines are for either time prediction or location prediction, we decided to compare them with our method in such a way that we simultaneously learn/predict both $p_l^{(t)}(\cdot; h_{t_l})$ and $p_l^{(x)}(\cdot; t_l, \mathbf{h}_{\mathbf{x}_l})$, where for $p_l^{(x)}(\cdot; t_l, \mathbf{h}_{\mathbf{x}_l})$ we assume that the time t_l is known and given. In this way, the error is comparable with those that predict solely location since we are controlling for the error that could be introduced by a wrong time prediction \hat{t}_l .

A.3 ABLATION STUDY

To further leverage the importance of each layer/block we used in our proposed model, in addition to what we described in Sections 3.1 and 3.2, we also performed two ablation studies on our network using all the available datasets.

Table 2: Results representing the negative log-likelihood for the events in the range $n + 1$ to $n + L$ for the learned distributions including ablation studies (less is better). \pm indicates the standard deviation of the loss among all test batch sequences.

	Earthquake	Citibike	Covid-19	Pinwheel
Our network ($p_l^{(t)}$)	0.33 \pm 0.17	-6.279 \pm 0.123	-5.87 \pm 0.09	1.27 \pm 0.26
Our network ($p_l^{(x)}$)	2.03 \pm 0.26	-3.28 \pm 0.14	2.18 \pm 0.21	1.57 \pm 0.09
Our network ($p_l^{(t,x)}$)	2.36 \pm 0.36	-9.56 \pm 0.18	-3.7 \pm 0.16	2.84 \pm 0.19
History independent ($p_l^{(t)}$)	1.26 \pm 0.26	-6.3 \pm 0.15	-5.9 \pm 0.22	1.27 \pm 0.27
History independent ($p_l^{(x)}$)	2.49 \pm 0.32	-3.27 \pm 0.08	1.64 \pm 0.05	7.06 \pm 0.03
History independent ($p_l^{(t,x)}$)	3.75 \pm 0.46	-9.58 \pm 0.14	-4.26 \pm 0.09	8.334 \pm 0.26
No decoder ($p_l^{(t)}$)	0.37 \pm 0.18	-6.076 \pm 0.13	-5.41 \pm 0.12	1.29 \pm 0.25
No decoder ($p_l^{(x)}$)	2.21 \pm 0.3	-3.1 \pm 0.09	2.2 \pm 0.32	0.61 \pm 0.06
No decoder ($p_l^{(t,x)}$)	2.58 \pm 0.41	-9.17 \pm 0.14	-3.21 \pm 0.28	1.9 \pm 0.18

A.3.1 HISTORY INDEPENDENCY

In this study, we assume that $p_l^{(x)}$ and $p_l^{(t)}$ are completely independent of the history, meaning that we have $p_l^{(x)}(\cdot; t_l, \mathbf{h}_{\mathbf{x}_l}) = p_l^{(x)}(\cdot; t_l)$ and $p_l^{(t)}(\cdot; h_{t_l}) = p_l^{(t)}$. For this matter we provide batches of zero inputs to the encoder $\{(t_i, \mathbf{x}_i, M_i)\}_{i=1}^n = 0$ in both the training and testing phase while keeping all other steps exactly as before. Results for this study on different datasets are given in Table 2.

South California Earthquakes: As shown by Table 2, removing history dependency from our network causes worse performance of multi-event forecasting in both time and space for this dataset. This is in line with the well-known epidemic-type-aftershock-sequence (ETAS) model (Ogata, 1998) suggesting long-term and long-range dependencies in both time and space between seismic events.

Citibike: Removing history dependency on this data, did not highly affect the results. Our interpretation of this dataset is that many other factors other than the previous bikers' markers included in this dataset might influence the time or location of riding a bike. Therefore, removing history dependency does not highly influence forecasting the time or location of multiple bikers in the future when no other influencing factors are considered.

Covid-19: We realized that removing the history dependency did not affect time distribution forecasting, but instead enhanced predicting space distributions. Despite the disease being very contagious, the space distribution forecasting results are not surprising at all. This is due to the fact that we are taking into account the long-range history dependency over many different states of the

United States. Unless we consider trips between states, the patients’ time and location in one state might even provide false information influencing the chances of others getting the disease in another state. Therefore, removing this information might even enhance predicting space distributions. Our interpretation of very little changes in forecasting time distribution in the case of having no historical information is that covid-19 is among the fastest disease to catch, therefore not very long-term dependency in time is needed to forecast the time of catching the disease. We guess that the short-term dependency captured by the decoder during training is enough to predict when the next person will catch the disease.

Pinwheel: As shown on Figure 11 top, the sequences of pinwheel events are formed in a clock-wise way. Therefore, historical positional and temporal information should have a big affect on predicting the location of next events in the future. This is in line with the results shown in Table 2, indicating how removing historical information worsens spatial multi-event distribution forecasting. On the other hand, despite using synthetic Hawkes time (Chen & Stindl, 2022) associated with pinwheel events, we saw no change in time forecasting when removing history dependency. For this outcome, we infer that the thinning process (Ogata et al., 1993) used to extract the times (Chen & Stindl, 2022) assigned to the events does not consider long-term and long-range dependency. Therefore, the short-term history dependency that exists among the output events $\{(t_l, \mathbf{x}_l)\}_{l=n+1}^{n+L}$ that is captured during training by the decoder layer, suffices predicting multiple events’ time distributions with no reliance on further look onto the history.

A.3.2 NO DECODER

Here, we no longer feed $\{(t_l, \mathbf{x}_l, M_l)\}_{l=n+1}^{n+L}$ to the decoder during the training phase. This is similar to the test phase of our model where we input $\{(t_l, \mathbf{x}_l, M_l)\}_{l=n+1}^{n+L} = 0$ to the already trained network. However, in this study we input $\{(t_l, \mathbf{x}_l, M_l)\}_{l=n+1}^{n+L} = 0$ during both training and testing phases. Results in Table 2 indicate that except for the pinwheels dataset’s spatial distribution forecasting, removing the decoder does not significantly change the prediction results. Nonetheless, results associated with the full network are slightly better. As earlier mentioned in Section 3.2 during training, the network learns to fit the distribution of multiple future events jointly, in a self-supervised setting. Hence, during testing, zero-padding the input to the decoder still obtains an estimation of the spatio-temporal probability densities of several events in the future. This happens due to the reliance of the network on the historical input data, as well as the learned bias associated with the decoder layer. However, we accept that the learned bias might only provide little history dependency information between output events, especially in the case of a small number of events to be predicted ($L = 3$ in our case). Since the setup when having no decoder differs from the main setup of our proposed network only in terms of the learned bias associated with the decoder layer, we do not expect to see much difference in the results. However, we suspect that having larger output sequences will highlight this difference between the two settings. In that case, we could consider a larger distance in time and space among the output events, which results in differently tuning the decoder layers’ bias during the training phase.

In the case of the pinwheel dataset, we realized that removing the decoder enhanced the multi-event spatial distribution forecasting results. This is when the batched temporal distribution forecasting results are negatively affected. In line with the outcomes of the previous ablation study, we inferred that the pinwheel events are influenced massively by the long-term history dependency in space. Therefore, removing the short-term dependency among the output events will allow the events to focus more on long-range effects of the historical data. However, the results associated with predicting the temporal distributions got worse. As expected from the outcome of the study in Appendix A.3.1, this might be due to the importance of short-term dependency between the time of the events, modeled by the thinning process (Ogata et al., 1993).

A.3.3 CONCLUSION ON ABLATION STUDY

Through these two studies and what we discussed in the main body of this paper, we showed the critical role of each layer forming our network. Mainly, we observed that each layer is designed for a specific task, driven by the physics behind the formation of the used dataset. Unlike other methods that solely rely on the Hawkes process, our network is designed to fully consider any sort of short- or long-term history dependencies desired by the physics behind data.

In general, from these studies, we concluded that the performance of our proposed network is highly data-driven. Predicting batched distributions associated with datasets that are formed based on both long- and short-range history dependencies, such as the earthquake dataset, demands the usage of both the encoder and decoder blocks of our model.

When events are more influenced by short-term dependencies, or there’s a lack of the presence of other factors influencing the events’ occurrence, the encoder might negatively affect the prediction results. This was shown as the outcome of the first study on the Covid-19 dataset. The reverse applies to events dependent on long-range history dependencies, such as the pinwheel and earthquake datasets. There might be cases where removing the encoder has no influence on the prediction results, due to no significant history dependency, such as the Citibike dataset.

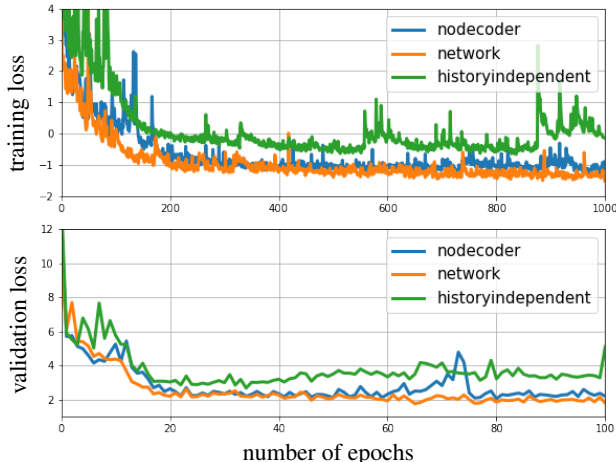


Figure 4: Comparing the performance of the proposed network with two ablation studies applied to the earthquake data. Both removing the decoder and the history dependency worsen the prediction results. This highlights the importance of both long- and short-range history dependencies among seismic events.

On the other hand, capturing the short-term dependencies, mainly done by the decoder layer, is important when events are also influenced by these dependencies. This conclusion is based on the outcomes of applying the second ablation study to many datasets, such as the earthquake dataset, Citibike, Covid-19, and the temporal events forming the Pinwheel dataset.

Since all datasets used in this work, except for the pinwheel dataset, represent real-world phenomena, we believe no single model can be used to learn/predict distributions associated with them. Therefore, it is critical to have a network that can be driven by the data’s underlying physics, whilst being capable of simultaneous multi-event prediction.

A.4 MORE RESULTS AND ANALYSIS

To show the performance of our network on synthetic Hawkes time simulated data vs. the baseline model, we used the pinwheel dataset as introduced in Section 4.1, where the events are distributed with less bias towards a specific region. As depicted in Figure 5 on top, not only the history dependent batched spatial distributions for the next three events are well predicted using our network, but also regions prone to events occurring further away in time, shown in a clock-wise pattern in Figure 11, are as well predicted. On the other hand, the conditional GMM model was not successful in extracting the complex distribution associated with next events’ occurrences. Note that using the conditional GMM model to perform sequential next event prediction, at each step, we incorporate the true event shown by the green circle, rather than the newly predicted event shown by the red star, to redefine the history for next step prediction. This is due to the high error accumulation caused by using the sequentially predicted events.

Figure 6 shows the predicted time of the next three events applying **our proposed network and the time Hawkes model** to the simulated Hawkes pinwheel dataset. While the result on the top figure are indicative of simultaneous multiple future time event prediction, the bottom figures depict utilizing (2) sequentially while updating the history via the most recently predicted event. These results are in agreement with Table 1 indicating the outperformance of using the Hawkes model for this dataset. As mentioned earlier, this is due to the fact that the time assigned to the synthetic Hawkes dataset are simulated by performing thinning algorithm on a time Hawkes process. As mentioned earlier, the expected times extracted by applying our network to this and all other datasets are measured

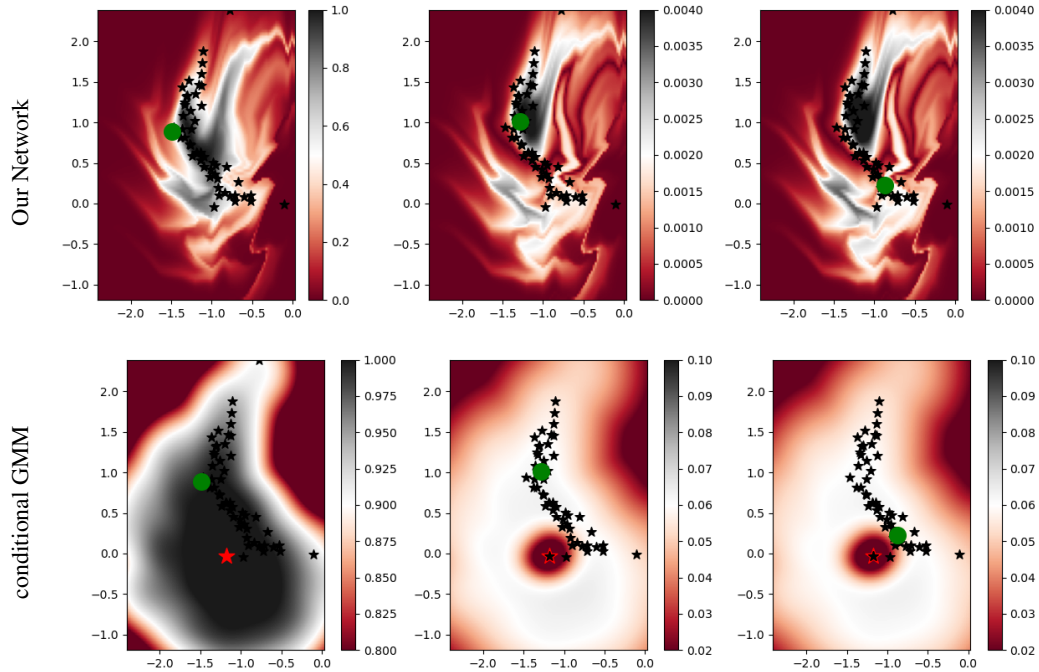


Figure 5: Visualizations of the learned history dependent batched spatial distributions for the last three events of simulated Hawkes pinwheel sequence using **our proposed architecture** vs. **conditional GMM**. The first figure on the left hand side depicts the predicted distribution associated with the first event, where the middle and the right hand side figures only depict the consecutive differences between the predicted densities. The green circle shows the true event that is expected to occur. The black stars represent the last 50 events of the input sequence as history information. The red star on bottom figures shows the sequentially predicted next event using (2) which needs to be injected back to the history for the next step prediction. Due to high error accumulation, we instead use the green circle to update history on the bottom figures.

by taking the average of 1000 samples taken from the batched history dependent predicted time distributions.




Figure 6: The true (green circles) vs. predicted (red stars) time of the last three events of an earthquake sequence using **our model** on top and the **time Hawkes model** on the bottom. The results on top are generated simultaneously, whereas the bottom results are predicted sequentially via updating history based on the most recently predicted event time. The linear line is indicative of predicting same time-intervals between events.

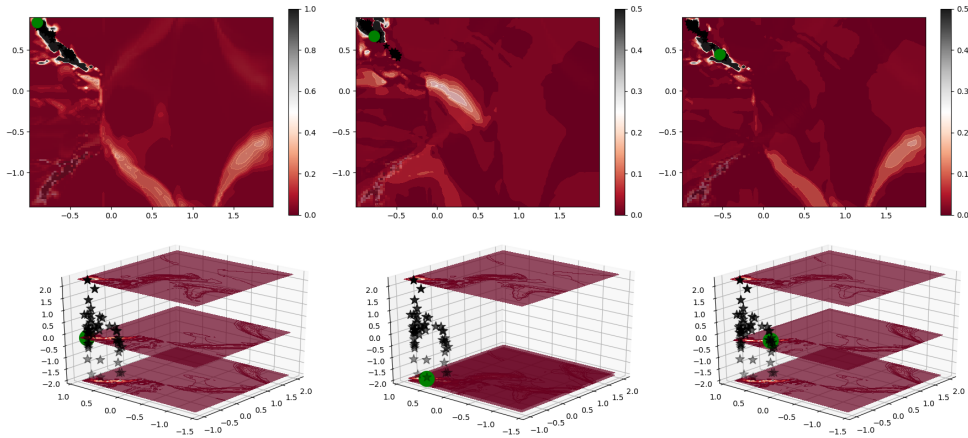


Figure 7: Visualizations of the learned history dependent batched spatial distributions for the last three events of an unseen earthquake sequence using **our proposed architecture**. The **bottom** figures depict the predicted batched spatial distribution associated with the last three events in 3 dimensions where the z axis is along events’ depth. The 2-d representation of the last three events’ predicted spatial distributions mapped to the events’ depth are shown on the **top** plots where except for the one on the right, we show the differences in the predicted distributions. The black stars on the spatial distributions represent the last 50 events of the input sequence as the history information, where in this figure the history events are very much aligned along the main fault in South California.

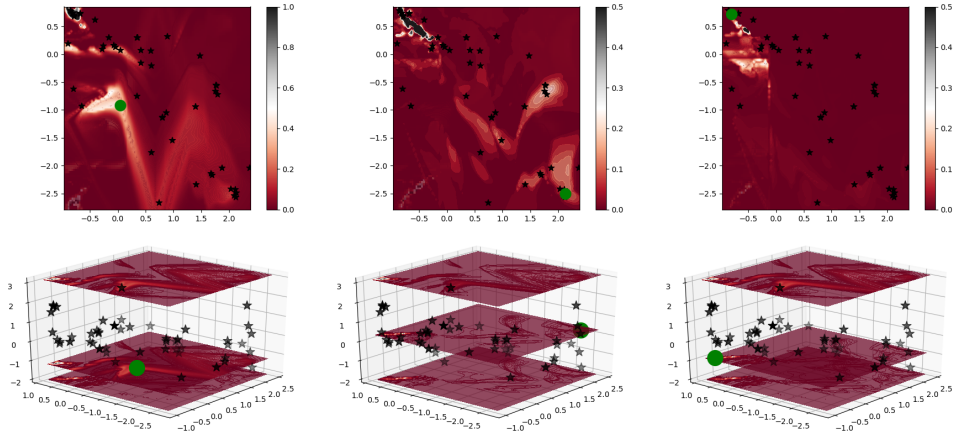


Figure 8: Same visualization as in Figure 7 on another earthquake sequence of length 500 where the history events, shown by black stars, are more distributed along all fault lines rather than just one.

A.5 BASELINES

We use baselines to separately model the time and spatial events distributions, conditioned on the history. Earlier in (10), we showed that the loss defined for the benchmarks have some differences with the loss we defined for our proposed architecture in the training phase. As already discussed in the main body of the paper, the baseline models only provide the expected next event using (2), which unlike our proposed method, gives the predicted event explicitly rather than a distribution. For further prediction in the future, the already predicted event has to be incorporated in the history, which will accordingly update the intensity function $\lambda(t_i|H_{t_i})$ as well as the space distribution $p(\mathbf{x}_i|t_i, H_{t_i})$ to be used for further predictions. To have a precise updated history, we need our predicted points to be as close to the true values as possible. Therefore, we expand (10) to provide




Figure 9: Visualization of the first 3 out of 6 encoder-decoder attention layers, where we show 3 out of 6 heads per layer. These figures are the results of applying **our proposed network** to the pinwheel dataset. The x axis represent historical events, whereas the y axis show the three expected events in the future. The colors show how weak or strong the past events have an influence on the future events' occurrence.

further detail as

$$\begin{aligned}
 \text{loss} = & \underbrace{-\sum_{i=1}^n \log p_t(t_i|H_{t_i}) - \sum_{i=1}^n \log p_{\mathbf{x}}(\mathbf{x}_i|t_i, H_{t_i})}_{\text{negative log-likelihood encoding history data}} \\
 & \underbrace{-\sum_{l=n+1}^{n+L} \log p_t(t_l|\tilde{H}_{t_l}) - \sum_{l=n+1}^{n+L} \log p_{\mathbf{x}}(\mathbf{x}_l|t_l, \tilde{H}_{t_l})}_{\text{negative log-likelihood of predicted events conditioned on the updated history}} \\
 & \underbrace{+\lambda_1 \sum_{l=n+1}^{n+L} |t_l - \hat{t}_l| + \lambda_2 \sum_{l=n+1}^{n+L} \|\mathbf{x}_l - \hat{\mathbf{x}}_l\|_2}_{\text{restricting the predicted values to match the true values}},
 \end{aligned} \tag{13}$$

where \tilde{H}_{t_l} is the updated history including the most recent predicted event $\{\hat{t}_l, \hat{\mathbf{x}}_l\}$.

Conditional Gaussian mixture model: We implemented two forms for the Conditional Gaussian mixture models as the baselines to learn the spatial distributions $p(\mathbf{x}_i|t_i, H_{t_i})_{i=1:n}$ mentioned in

Section 4.2. In the first method we compute the pairwise Gaussian log-likelihood among all \mathbf{x}_i and \mathbf{x}_j associated with events $1 \leq j \leq i \leq n$, as well as the pairwise log-timedecay Δt_{ij} . We minimize the loss associated with adding up the pairwise Gaussian log-likelihood and log-timedecay to learn the history dependent Gaussian mixture model.

The second method is the general K-cluster Gaussian mixture model introduced in Murphy (2012); Bishop & Nasrabadi (2006), where the learned parameters are functions of the historical events markers. Having a mixture of K Gaussian densities

$$p(\mathbf{x}_l | \mu, \Sigma) = \prod_{k=1}^K \frac{1}{(2\pi)^{\frac{n_k}{2}} |\Sigma_k|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\mathbf{x}_{l_k} - \mu_k)^\top \Sigma_k^{-1}(\mathbf{x}_{l_k} - \mu_k)\right), \quad 1 \leq k \leq K \quad (14)$$

where n_k represents the number of samples associated with cluster k, we learn μ_k and Σ_k from $\{(t_i, \mathbf{x}_i, M_i)\}_{i=1}^n$ as the time, space and the extra marker associated with all events assigned to cluster k. In this case, $p(\mathbf{x}_l | \mu, \Sigma)$ in (14), denotes the conditional spatial probability associated with event $\mathbf{x}_{l=n+1}$.

Although the second method is capable of conditionally being dependent on any other markers associated with events, the first method had a better performance on all datasets. Therefore, in this work we only reported the results based on the first method, where both implementations are given in our publicly available code.

A.6 PREPROCESSING OF DATA SEQUENCES

South California Earthquakes. We use earthquake discrete events distributed over South of California with latitude, longitude and depth. The distribution of all events along the main faults as well as four sequences of 500-length size, projected on a 2-d map of South California faults, are shown in Figure 11. Since we don't work with in-between distances, there is no need to convert to the Cartesian coordinate system. After forming sequences as mentioned in Section 4.1, we normalize

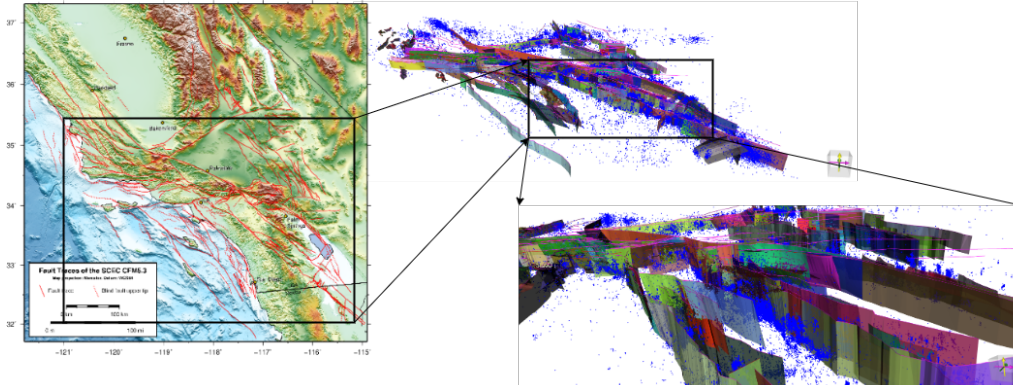


Figure 10: Visualization of distribution of discrete earthquake events on the South California main faults covering events from 2008 to 2016.

the space and magnitude of each sequence with respect to the mean and standard deviation gathered from all events. The bias caused by normalization towards more earthquake prone regions as previously shown in Figure 2 is remained to be eased in later work as discussed in Section 5.

Citibike. We use the data from April to August of 2019. No preprocessing was performed for this dataset. This resulted in 3500 training sequences, 900 validation sequences, and 600 test sequences of same length.

Covid-19. We use the data from March of 2020 to May of 2022. No preprocessing was performed for this dataset. This resulted in 1400 training sequences, 360 validation sequences, and 240 test sequences of same length.

Pinwheel. We sample from a multivariate process with 15 clusters, each containing 150 events. Due to the space Hawkes synthetic dataset not being available, we used the same spatial distribution provided in Chen et al. (2020) with the assigned Hawkes timings using the RHawkes package Chen & Stindl (2022). Since we have sequences of 500 events, each sequence will contain a few number of clusters as shown in Figure 11. This resulted in 613 training sequences, 157 validation sequences, and 106 test sequences of same length.

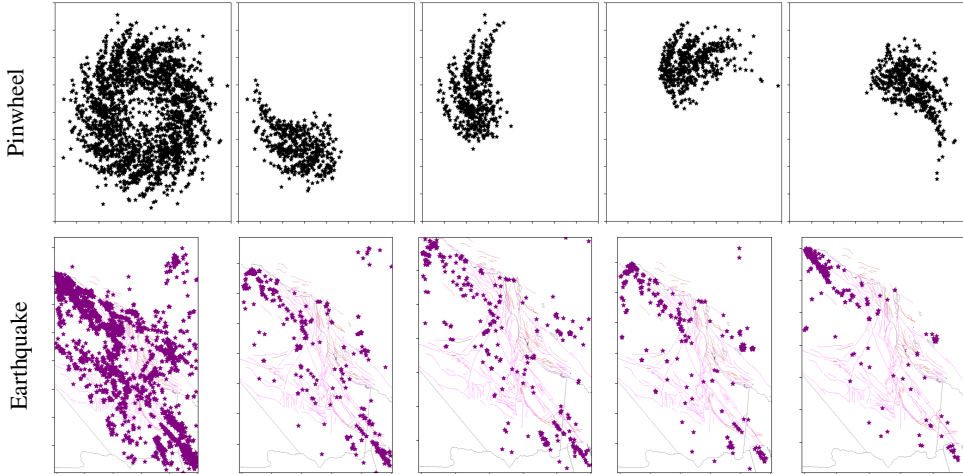


Figure 11: **Top:** Pinwheel synthetic dataset where simulated Hawkes times are assigned to the events of clusters in a clockwise way. These timings are simulated using thinning algorithm Ogata (1981). **Bottom:** South California earthquakes distributed over the main faults.

A.7 HYPER-PARAMETERS

We have set the number of attention layers in both the encoder and the decoder to 6, whereas each layer contains 6 heads individually. The reason behind these choices are what proposed by the main transformer network proposed in NLP Vaswani et al. (2017), where they suggest to choose between 6 to 8 attention heads and layers. Before inputting any of the sequences to the encoder and the decoder, we first used Multi-layer perceptron (MLP) layers with ELU activation functions, to embed to a 64 dimensional space. The latent representations \mathbf{h}_{t_i} , $\mathbf{h}_{\mathbf{x}_i}$, depicted in Figure 1 are in a 64 dimensional space, whereas we used another set of MLP layers with ELU activation functions to transfer back h_{t_i} , and $\mathbf{h}_{\mathbf{x}_i}$ to the same dimension as t_i , and \mathbf{x}_i before injecting them into the probabilistic layers. This was due to the use of bijective layers (normalizing flows) following the probabilistic layers only being capable of mapping to the same dimensional space.

Table 3: Hyper-parameters and layer-types used in our model for both training and test.

parameter	value	parameter	value
Epochs	1k	Batch-size	32
Sequence-length	500	learning-rate (starting)	1e-3
Input-length	497	Output-length	3
Attention layers	6	Attention heads	6
Marker-embedding dim	64	Dropout-rate	0.1~0.2
λ_1 (for L_1 norm)	0.1	λ_2 (for L_2 norm)	0.1

probabilistic layer	type	bijective layer	type
time	exponential	time	softsign
space	multi-variate Gaussian	space	RealNVP




Figure 12: Histograms of different markers associated with different datasets used in this work. The "extra marker"’s associated with Earthquake, Citibike, Covid-19, and pinwheel datasets are the magnitude, birth year, number of cases, and all-one intensity, respectively.

As mentioned earlier in Section 3.1, we used multi-variate Gaussian and exponential probabilistic layers for space and time, respectively as considered by the Hawkes model. The markers histograms shown by Figure 12 as well indicate the usefulness of these choices for our probabilistic layers. We also tried modeling the time distribution using Gaussian probabilistic layers which led to bad results.

For the bijective flow layers we tried using RealNVP Dinh et al. (2016) and the Masked Autoregressive flows Papamakarios et al. (2017) for the space distributions, and we tried softplus and softsign flows for time. Our results indicate that the RealNVP and softplus are capable of better modeling the expected space and time batched distributions. Note that both the RealNVP and Masked Autoregressive flows are only applicable to data with dimension of higher than 1, and therefore are not good candidates for transforming the time batched distributions.

To overcome the overfitting when training our model, we used dropout layers with rates listed as $\{0.1, 0.15, 0.2\}$. For the earthquake and Hawkes pinwheel synthetic data, both 0.1 and 0.15 were

good options whereas for the Citibike and Covid-19 we had to increase the rate to 0.2 to reduce the overfitting.

All the mentioned hyperparameters can be set differently in our network using `argparsing`⁴, once the code is publicly available. Our network can be applied to any discrete form data with variable dimension assigned markers for future use.

⁴Code is available as supplementary material and will be publicly released in github if the paper is accepted.