

Neural Point Process for Learning Spatiotemporal Event Dynamics

Zihao Zhou

UC San Diego

ziz244@ucsd.edu

Xingyi Yang

National University of Singapore

xyang@u.nus.edu

Ryan A. Rossi

Adobe Research

ryrossi@adobe.com

Handong Zhao

Adobe Research

hazhao@adobe.com

Rose Yu

UC San Diego

roseyu@ucsd.edu

Editors: R. Firooz, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

Learning the dynamics of spatiotemporal events is a fundamental problem. Neural point processes enhance the expressivity of point process models with deep neural networks. However, most existing methods only consider temporal dynamics without spatial modeling. We propose Deep Spatiotemporal Point Process (DeepSTPP), a deep dynamics model that integrates spatiotemporal point processes. Our method is flexible, efficient, and can accurately forecast irregularly sampled events over space and time. The key construction of our approach is the nonparametric space-time intensity function, governed by a latent process. The intensity function enjoys closed form integration for the density. The latent process captures the uncertainty of the event sequence. We use amortized variational inference to infer the latent process with deep networks. Using synthetic datasets, we validate our model can accurately learn the true intensity function. On real-world benchmark datasets, our model demonstrates superior performance over state-of-the-art baselines. Our code and data can be found at the [link](#).

Keywords: spatiotemporal dynamics, neural point processes, kernel density estimation

1. Introduction

Accurate modeling of spatiotemporal event dynamics is fundamentally important for disaster response (Veen and Schoenberg, 2008), logistic optimization (Safikhani et al., 2018) and social media analysis (Liang et al., 2019). Compared to other sequence data such as texts or time series, spatiotemporal events occur irregularly with uneven time and space intervals.

Discrete-time deep dynamics models such as recurrent neural networks (RNNs) (Hochreiter and Schmidhuber, 1997; Chung et al., 2014) assume events to be evenly sampled. Interpolating an irregular sampled sequence into a regular sequence can introduce significant biases (Rehfeld et al., 2011). Furthermore, event sequences contain strong spatiotemporal dependencies. The rate of an event depends on the preceding events, as well as the events geographically correlated to it.

Spatiotemporal point processes (STPP) (Daley and Vere-Jones, 2007; Reinhart et al., 2018) provides the statistical framework for modeling continuous-time event dynamics. As shown in Fig-

ure 1, given the history of events sequence, STPP estimates the intensity function that is evolving in space and time. However, traditional statistical methods for estimating STPPs often require strong modeling assumptions, feature engineering, and can be computationally expensive.

Machine learning community is observing a growing interest in continuous-time deep dynamics models that can handle irregular time intervals. For example, Neural ODE (Chen et al., 2018) parametrizes the hidden states in an RNN with an ODE. Shukla and Marlin (2018) uses a separate network to interpolate between reference time points. Neural temporal point process (TPP) (Mei and Eisner, 2017; Zhang et al., 2020; Zuo et al., 2020) is an exciting area that combines fundamental concepts from temporal point processes with deep learning to model continuous-time event sequences, see a recent review on neural TPP (Shchur et al., 2021). However, most of the existing models only focus on *temporal* dynamics without considering *spatial* modeling.

In the real world, while time is a unidirectional process (arrow of time), space extends in multiple directions. This fundamental difference from TPP makes it nontrivial to design a unified STPP model. The naive approach to approximate the intensity function by a deep neural network would lead to intractable integral computation for likelihood. Prior research such as Du et al. (2016) discretizes the space as ‘‘markers’’ and uses marked TPP to classify the events. This approach cannot produce the space-time intensity function. Okawa et al. (2019) models the spatiotemporal density using a mixture of symmetric kernels, which ignores the unidirectional property of time. Chen et al. (2021) proposes to model temporal intensity and spatial density separately with neural ODE, which is computationally expensive.

We propose a simple yet computationally efficient approach to learning STPP. Our model, *Deep Spatiotemporal Point Process* (DeepSTPP) marries the principles of spatiotemporal point processes with deep learning. We take a non-parametric approach and model the space-time intensity function as a mixture of kernels. The parameters of the intensity function are governed by a latent stochastic process no sampling which captures the uncertainty of the event sequence. The latent process is then inferred via amortized variational inference. That is, we draw a sample from the variational distribution for every event. We use a Transformer network to parametrize the variational distribution conditioned on the previous events.

Compared with existing approaches, our model is non-parametric, hence does not make assumptions on the parametric form of the distribution. Our approach learns the space-time intensity function jointly without requiring separate models for temporal intensity function and spatial density as in Chen et al. (2021). Our model is probabilistic by nature and can describe various uncertainties in the data. More importantly, our model enjoys closed form integration, making it feasible for processing large-scale event datasets. To summarize, our work makes the following key contributions:

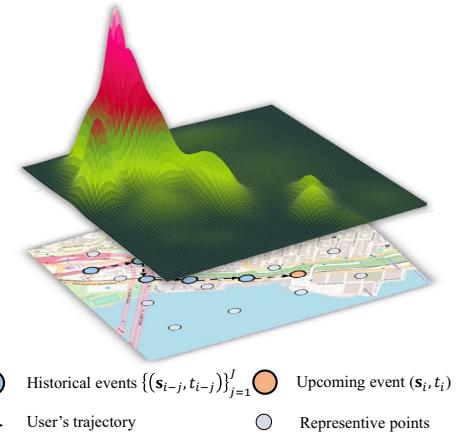


Figure 1: Illustration of learning spatiotemporal point process. We aim to learn the space-time intensity function given the historical event sequence and representative points as background.

- **Deep Spatiotemporal Point Process.** We propose a novel Deep Point Process model for forecasting unevenly sampled spatiotemporal events. It integrates deep learning with spatiotemporal point processes to learn continuous space-time dynamics.
- **Neural Latent Process.** We model the space-time intensity function using a non-parametric approach, governed by a latent stochastic process. We use amortized variational inference to perform inference on the latent process conditioned on the previous events.
- **Effectiveness.** We demonstrate our model using many synthetic and real-world spatiotemporal event forecasting tasks, where it achieves superior performance in accuracy and efficiency. We also derive and implement efficient algorithms for simulating STPPs.

2. Methodology

We first introduce the background of the spatiotemporal point process and then describe our approach to learning the underlying spatiotemporal event dynamics.

2.1. Background on Spatiotemporal Point Process

Spatiotemporal Point Process. Spatiotemporal point process (STPP) models the number of events $N(\mathcal{S} \times (a, b))$ that occurred in the Cartesian product of the spatial domain $\mathcal{S} \subseteq \mathbb{R}^2$ and the time interval $(a, b]$. It is characterized by a non-negative *space-time intensity function* given the history $\mathcal{H}_t := \{(\mathbf{s}_1, t_1), \dots, (\mathbf{s}_n, t_n)\}_{t_n \leq t}$:

$$\lambda^*(\mathbf{s}, t) := \lim_{\Delta s \rightarrow 0, \Delta t \rightarrow 0} \frac{\mathbb{E}[N(B(\mathbf{s}, \Delta s) \times (t, t + \Delta t)) | \mathcal{H}_t]}{B(\mathbf{s}, \Delta s) \Delta t} \quad (1)$$

which is the probability of finding an event in an infinitesimal time interval $(t, t + \Delta t]$ and an infinitesimal spatial ball $\mathcal{S} = B(\mathbf{s}, \Delta s)$ centered at location \mathbf{s} .

Example 1: Spatiotemporal Hawkes process (STH). Spatiotemporal Hawkes (or self-exciting) process assumes every past event has an additive, positive, decaying, and spatially local influence over future events. Such a pattern resembles neuronal firing and earthquakes. It is characterized by the following intensity function (Reinhart et al., 2018):

$$\lambda^*(\mathbf{s}, t) := \mu g_0(\mathbf{s}) + \sum_{i: t_i < t} g_1(t, t_i) g_2(\mathbf{s}, \mathbf{s}_i) : \mu > 0 \quad (2)$$

where $g_0(\mathbf{s})$ is the probability density of a distribution over \mathcal{S} , g_1 is the triggering kernel and is often implemented as the exponential decay function, $g_1(\Delta t) := \alpha \exp(-\beta \Delta t) : \alpha, \beta > 0$, and $g_2(\mathbf{s}, \mathbf{s}_i)$ is the density of an unimodal distribution over \mathcal{S} centered at \mathbf{s}_i .

Example 2: Spatiotemporal Self-Correcting process (STSC). Self-correcting spatiotemporal point process (Isham and Westcott, 1979) assumes that the background intensity increases with a varying speed at different locations, and the arrival of each event reduces the intensity nearby. STSC can model certain regular event sequences, such as an alternating home-to-work travel sequence. It has the following intensity function:

$$\lambda^*(\mathbf{s}, t) = \mu \exp \left(g_0(\mathbf{s}) \beta t - \sum_{i: t_i < t} \alpha g_2(\mathbf{s}, \mathbf{s}_i) \right) : \alpha, \beta, \mu > 0 \quad (3)$$

Here $g_0(\mathbf{s})$ is the density of a distribution over \mathcal{S} , and $g_2(\mathbf{s}, \mathbf{s}_i)$ is the density of an unimodal distribution over \mathcal{S} centered at location \mathbf{s}_i .

Maximum likelihood Estimation. Given a history of n events \mathcal{H}_t , the joint log-likelihood function of the observed events for STPP is as follows:

$$\log p(\mathcal{H}_t) = \sum_{i=1}^n \log \lambda^*(\mathbf{s}_i, t_i) - \int_{\mathcal{S}} \int_0^t \lambda^*(\mathbf{u}, \tau) d\mathbf{u} d\tau \quad (4)$$

Here, the space-time intensity function $\lambda^*(\mathbf{s}, t)$ plays a central role. Maximum likelihood estimation seeks the optimal $\lambda^*(\mathbf{s}, t)$ from data that optimizes Eqn. 4.

Predictive distribution. Denote the probability density function (PDF) for STPP as $f(\mathbf{s}, t | \mathcal{H}_t)$ which represents the conditional probability that next event will occur at location \mathbf{s} and time t , given the history. The PDF is closely related to the intensity function:

$$f(\mathbf{s}, t | \mathcal{H}_t) = \frac{\lambda^*(\mathbf{s}, t)}{1 - F^*(\mathbf{s}, t | \mathcal{H}_t)} = \lambda^*(\mathbf{s}, t) \exp \left(- \int_{\mathcal{S}} \int_{t_n}^t \lambda^*(\mathbf{u}, \tau) d\tau d\mathbf{u} \right) \quad (5)$$

where F is the cumulative distribution function (CDF), see derivations in Appendix ???. This means the intensity function specifies the expected number of events in a region conditional on the past.

The predicted time of the next event is the expected value of the predictive distribution for time $f^*(t)$ in the entire spatial domain:

$$\mathbb{E}[t_{n+1} | \mathcal{H}_t] = \int_{t_n}^{\infty} t \int_{\mathcal{S}} f^*(\mathbf{s}, t) d\mathbf{s} dt = \int_{t_n}^{\infty} t \exp \left(- \int_{t_n}^t \lambda^*(\tau) d\tau \right) \lambda^*(t) d\mathbf{s} dt$$

Similarly, the predicted location of the next event evaluates to:

$$\mathbb{E}[\mathbf{s}_{n+1} | \mathcal{H}_t] = \int_{\mathcal{S}} \mathbf{s} \int_{t_n}^{\infty} f^*(\mathbf{s}, t) dt d\mathbf{s} = \int_{t_n}^{\infty} \exp \left(- \int_{t_n}^t \lambda^*(\tau) d\tau \right) \int_{\mathcal{S}} \mathbf{s} \lambda^*(\mathbf{s}, t) d\mathbf{s} dt$$

Unfortunately, Eqn. (4) is generally intractable. It requires either strong modeling assumptions or expensive Monte Carlo sampling. We propose the Deep STPP model to simplify the learning.

2.2. Deep Spatiotemporal Point Process (DSTPP)

We propose DeepSTPP, a simple and efficient approach for learning the space-time event dynamics. Our model (1) introduces a latent process to capture the uncertainty, (2) parametrizes the latent process with deep neural networks to increase model expressivity, and (3) approximates the intensity function with a set of spatial and temporal kernel functions.

Neural latent process. Given a sequence of n events, we wish to model the conditional density of observing the next event given the history $f(\mathbf{s}, t | \mathcal{H}_t)$. We introduce a latent process to capture the uncertainty of the event history and infer the latent process with amortized variational inference. The latent process dictates the parameters in the space-time intensity function. We sample from the latent process using the re-parameterization trick (Kingma and Welling, 2013).

As shown in Figure 2, given the sequence with n events $\mathcal{H}_t = \{(\mathbf{s}_1, t_1), \dots, (\mathbf{s}_n, t_n)\}_{t_n \leq t}$, we encode the entire sequence into the high-dimensional embedding. We use positional encoding to encode the sequence order. To capture the stochasticity in the temporal dynamics, we introduce a

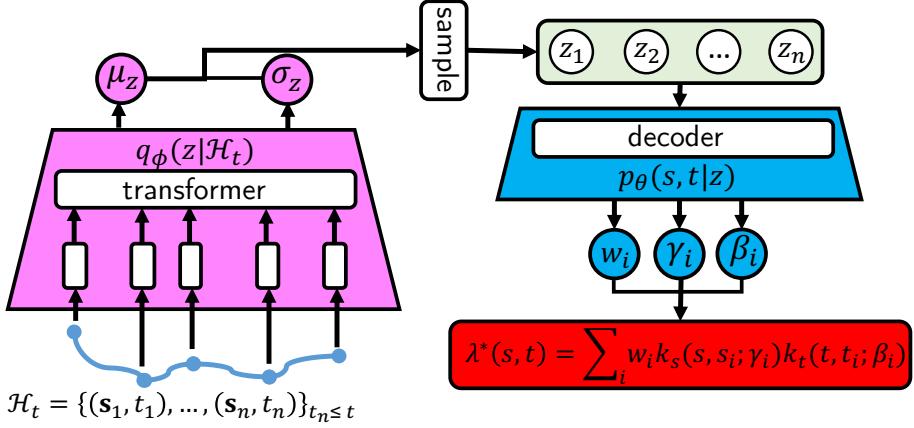


Figure 2: Design of our DeepSTPP model. For a historical event sequence, we encode it with a transformer network and map to the latent process (z_1, \dots, z_n) . We use a decoder to generate the parameters (w_i, γ_i, β_i) for each event i given the latent process. The estimate intensity is calculated using kernel functions k_s and k_t and the decoded parameters.

latent process $z = (z_1, \dots, z_n)$ for the entire sequence. We assume the latent process follows a multivariate Gaussian at each time step:

$$z_i \sim q_\phi(z_i | \mathcal{H}_t) = \mathcal{N}(\mu, \text{Diag}(\sigma)) \quad (6)$$

where the mean μ and covariance $\text{Diag}(\sigma)$ are the outputs of the embedding neural network. In our implementation, we found using a Transformer (Vaswani et al., 2017) with sinusoidal positional encoding to be beneficial. The positions to be encoded are the normalized event time instead of the index number, to account for the unequal time interval. Recently, Zuo et al. (2020) also demonstrated that Transformer enjoys better performance for learning the intensity in temporal point processes.

Non-parametric model. We take a non-parametric approach to model the space-time intensity function $\lambda^*(\mathbf{s}, t)$ as:

$$\lambda^*(\mathbf{s}, t | z) = \sum_{i=1}^{n+J} w_i k_s(\mathbf{s}, \mathbf{s}_i; \gamma_i) k_t(t, t_i; \beta_i) \quad (7)$$

Here $w_i(z), \gamma_i(z), \beta_i(z)$ are the parameters for each event that is conditioned on the latent process. Specifically, w_i represents the non-negative intensity magnitude, implemented with a soft-plus activation function. J is the number of representative points that we will introduce later. $k_s(\cdot, \cdot)$ and $k_t(\cdot, \cdot)$ are the spatial and temporal kernel functions, respectively. For both kernel functions, we parametrize them as a normalized RBF kernel:

$$k_s(\mathbf{s}, \mathbf{s}_i) = \alpha^{-1} \exp(-\gamma_i \|\mathbf{s} - \mathbf{s}_i\|), \quad k_t(t, t_i) = \exp(-\beta_i \|t - t_i\|) \quad (8)$$

where the bandwidth parameter γ_i controls an event's influence over the spatial domain. The parameter β_i is the decay rate that represents the event's influence over time. $\alpha = \int_S \exp(-\gamma_i \|\mathbf{s} - \mathbf{s}_i\|) d\mathbf{s}$ is the normalization constant.

We use a decoder network to generate the parameters $\{w_i, \gamma_i, \beta_i\}$ given z separately, shown in Figure 2. Each decoder is a 4-layer feed-forward network. We use a softplus activation function to

ensure w_i and γ_i are positive. The decay rate β_i can be any real number, such that an event can also have constant or increasing triggering intensity over time.

Representative Points. In addition to n historical events, we also randomly sample J representative points from the spatial domain to approximate the background intensity. This is to account for the influence from unobserved events in the background, with varying rates at different locations. The model design in (7) enjoys a closed form integration, which gives the conditional PDF as:

$$f(\mathbf{s}, t | \mathcal{H}_t, z) = \lambda^*(\mathbf{s}, t | z) \exp \left(- \sum_{i=1}^{n+J} \frac{w_i}{\beta_i} [k_t(t_n, t_i) - k_t(t, t_i)] \right) \quad (9)$$

See the derivation details in Appendix ???. DeepSTPP circumvents the integration of the intensity function and enjoys fast inference in forecasting future events. In contrast, NSTPP [Chen et al. \(2021\)](#) is relatively inefficient as its ODE solver also requires additional numerical integration.

Parameter learning. Due to the latent process, the posterior becomes intractable. Instead, we use amortized inference by optimizing the evidence lower bound (ELBO) of the likelihood. In particular, given event history \mathcal{H}_t , the conditional log-likelihood of the next event is:

$$\log p(\mathbf{s}, t | \mathcal{H}_t) \geq \log p_\theta(\mathbf{s}, t | \mathcal{H}_t, z) + \text{KL}(q_\phi(z | \mathcal{H}_t) || p(z)) \quad (10)$$

$$= \log \lambda^*(\mathbf{s}, t | z) - \int_{t_n}^t \lambda^*(\tau) d\tau + \text{KL}(q || p) \quad (11)$$

where ϕ represents the parameters of the encoder network, and θ are the parameters of the decoder network. $p(z)$ is the prior distribution, which we assume to be Gaussian. $\text{KL}(\cdot || \cdot)$ is the Kullback–Leibler divergence between two distributions. We can optimize the objective function in Eqn. (11) w.r.t. the parameters ϕ and θ using back-propagation.

3. Related Work

Spatiotemporal Dynamics Learning. Modeling the spatiotemporal dynamics of a system in order to forecast the future is a fundamental task in many fields. Most work on spatiotemporal dynamics has been focused on spatiotemporal data measured at regular space-time intervals, e.g., ([Xingjian et al., 2015](#); [Li et al., 2018](#); [Yao et al., 2019](#); [Fang et al., 2019](#); [Geng et al., 2019](#)). For discrete spatiotemporal events, statistical methods include space-time point process, see ([Moller and Waagepetersen, 2003](#); [Mohler et al., 2011](#)). [Zhao et al. \(2015\)](#) proposes multi-task feature learning whereas [Yang et al. \(2018\)](#) proposes RNN-based model to predict spatiotemporal check-in events. These discrete-time models assume data are sampled evenly, thus are unsuitable for our task.

Continuous Time Sequence Models. Continuous time sequence models provide an elegant approach for describing irregular sampled time series. For example, [Chen et al. \(2018\)](#); [Jia and Benson \(2019\)](#); [Dupont et al. \(2019\)](#); [Gholami et al. \(2019\)](#); [Finlay et al. \(2020\)](#); [Kidger et al. \(2020\)](#); [Norcliffe et al. \(2021\)](#) assume the latent dynamics are continuous and can be modeled by an ODE. But for high-dimensional spatiotemporal processes, this approach can be computationally expensive. [Che et al. \(2018\)](#); [Shukla and Marlin \(2018\)](#) modify the hidden states with exponential decay. GRU-ODE-Bayes proposed by [De Brouwer et al. \(2019\)](#) introduces a continuous-time version of GRU and a Bayesian update network capable of handling sporadic observations. However, [Mozer et al. \(2017\)](#) shows that there is no significant benefit of using continuous-time RNN for discrete event data. Special treatment is still needed for modeling unevenly sampled events.

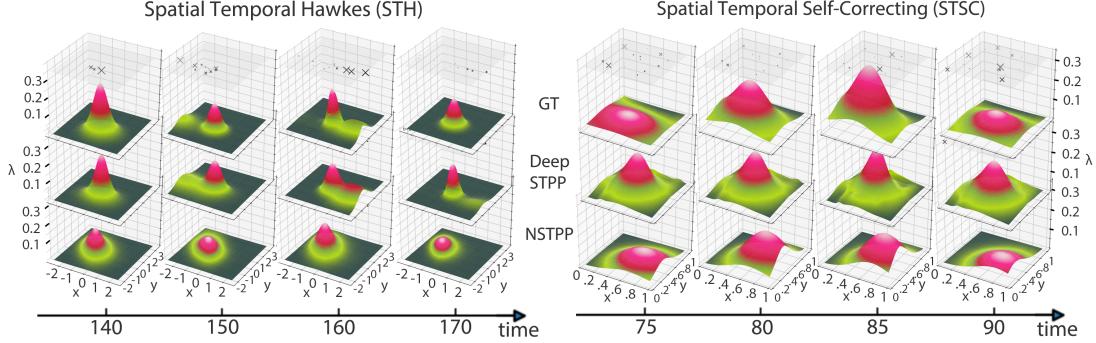


Figure 3: Ground-truth and learned intensity on two synthetic data. **Top:** ground-truth; **Middle:** learned intensity by our DeepSTPP model. **Bottom:** learned conditional intensity by NSTPP. The crosses on the top represent the event history, larger crosses are more recent events.

Deep Point Process. Point process is well-studied in statistics (Moller and Waagepetersen, 2003; Daley and Vere-Jones, 2007; Reinhart et al., 2018). Early work such as Linderman and Adams (2014) applies the graph model and Bayesian approach to infer the latent dynamic in point processes. Deep point process couples deep learning with point process and has received considerable attention. For example, neural Hawkes process applies RNNs to approximate the temporal intensity function (Du et al., 2016; Mei and Eisner, 2017; Xiao et al., 2017; Zhang et al., 2020), and Zuo et al. (2020) employs Transformers. Shang and Sun (2019) integrates graph convolution structure. However, all existing works focus on temporal point processes without spatial modeling. For datasets with spatial information, they discretize the space and treat them as discrete ‘‘markers’’. Okawa et al. (2019) extends Du et al. (2016) for spatiotemporal event prediction but they only predict the density instead of the next location and time of the event. Zhu et al. (2019) parameterizes the spatial kernel with a neural network embedding without consider the temporal sequence. Recently, Chen et al. (2021) proposes neural spatiotemporal point process (NSTPP) which combines continuous-time neural networks with continuous-time normalizing flows to parameterize spatiotemporal point processes. However, this approach is quite computationally expensive, which requires evaluating the ODE solver for multiple time steps.

4. Experiments

We evaluate DeepSTPP for spatiotemporal prediction using both synthetic and real-world data.

Baselines We compare DeepSTPP with the state-of-the-art models, including

- Spatiotemporal Hawkes Process (MLE) (Reinhart et al., 2018): it learns a spatiotemporal parametric intensity function using maximum likelihood estimation, see derivation in Appendix ??.
- Recurrent Marked Temporal Point Process (RMTPP) (Du et al., 2016): it uses GRU to model the temporal intensity function. We modify this model to take spatial location as marks.
- Neural Spatiotemporal Point Process (NSTPP) (Chen et al., 2021): state of the art neural point process model that parameterizes the spatial PDF and temporal intensity with continuous-time normalizing flows. Specifically, we use Jump CNF as it is a better fit for Hawkes processes.

All models are implemented in PyTorch, trained using the Adam optimizer. We set the number of representative points to be 100. The details of the implementation are deferred to the Appendix ???. For the baselines, we use the authors' original repositories whenever possible.

Datasets. We simulated two types of STPPs: spatiotemporal Hawkes process (STH) and spatiotemporal self-correcting process (STSC) . For both STPPs, we generate three synthetic datasets, each with a different parameter setting, denoted as DS1, DS2, and DS3 in the tables. We also derive and implement efficient algorithms for simulating STPPs based on Ogata's thinning algorithm ([Ogata, 1981](#)). We view the simulator construction as an independent contribution from this work. The details of the simulation can be found in Appendix ???. We use two real-world spatiotemporal event datasets from NSTPP ([Chen et al., 2021](#)) to benchmark the performance.

- **Earthquakes Japan:** catalog earthquakes data including the location and time of all earthquakes in Japan from 1990 to 2020 with magnitude of at least 2.5 from the U.S. Geological Survey. There are in total 1,050 sequences. The number of events per sequences ranges between 19 to 545 ¹.
- **COVID-19:** daily county level COVID-19 cases data in New Jersey state published by The New York Times. There are 1,650 sequences and the number of events per sequences ranges between 7 to 305.

For both synthetic data and real-world data, we partition long event sequences into non-overlapping subsequences according to a fixed time range T . The targets are the last event, and the input is the rest of the events. The number of input events varies across subsequences. For each dataset, we split each into train/val/test sets with the ratio of 8:1:1. All results are the average of 3 runs.

Table 1: Test log likelihood (LL) and Hellinger distance of distribution (HD) on synthetic data (LL higher is better, HD lower is better). Comparison between ours and NSTPP on synthetic datasets from two type of spatiotemporal point processes.

		Spatiotemporal Hawkes process						Spatiotemporal Self Correcting process					
		DS1		DS2		DS3		DS1		DS2		DS3	
		LL	HD	LL	HD	LL	HD	LL	HD	LL	HD	LL	HD
DeepSTPP (ours)		-3.8420	0.0033	-3.1142	0.4920	-3.6327	0.0908	-1.2248	0.2348	-1.4915	0.1813	-1.3927	0.2075
NSTPP		-5.3110	0.5341	-4.8564	0.5849	-3.7366	0.1498	-2.0759	0.5426	-2.3612	0.3933	-3.0599	0.3097

4.1. Synthetic Experiment Results

For synthetic data, we know the ground truth intensity function. We compare our method with the best possible estimator: maximum likelihood estimator (MLE), as well as the NSTPP model. The MLE is learned by optimizing the log-likelihood using the BFGS algorithm. RMTPP can only learn the temporal intensity thus is not included in this comparison.

Table 2: Estimated $\lambda^*(t)$ MAPE on synthetic data

	STH			STSC		
	DS1	DS2	DS3	DS1	DS2	DS3
DeepSTPP	3.33	369.44	11.30	7.84	3.22	20.98
NSTPP	53.41	17.69	3.85	99.99	39.33	37.39
RMTPP	263.83	729.78	0.62	45.55	21.26	37.46
MLE	2.98	11.30	4.38	27.38	18.20	20.01

1. The statistics differ slightly from the original paper due to updates in the data source.

Predictive log-likelihood. Table 1 shows the comparison of the predictive distribution for space and time. We report Log Likelihood (LL) of $f(\mathbf{s}, t | \mathcal{H}_t)$ and the Hellinger Distance (HD) between the predictive distributions and the ground truth averaged over time.

On both the STH and STSC datasets with different parameter settings, DeepSTPP outperform the baseline NSTPP in terms of LL and HD. It shows that DeepSTPP can estimate the spatiotemporal intensity more accurately for point processes with unknown parameters.

Temporal intensity estimate. Table 2 shows the mean absolute percentage error (MAPE) between the models’ estimated temporal intensity and the ground truth $\lambda^*(t)$ over a short sampled range. On the STH datasets, since MLE has the correct parametric form, it is the theoretical optimum. Compared to baselines, DeepSTPP generally obtained the same or lower MAPE. It shows that joint spatiotemporal modeling also improve the performance of temporal prediction.

Intensity visualization. Figure 3 visualizes the learned space-time intensity and the ground truth for STH and STSC, providing strong evidence that DeepSTPP can correctly learn the underlying dynamics of the spatiotemporal events. Especially, NSTPP has difficulty in modeling the complex dynamics of the multimodal distribution such as the spatiotemporal Hawkes process. NSTPP sometimes produces overly smooth intensity surfaces, and lost most of the details at the peak. In contrast, our DeepSTPP can better fit the multimodal distribution through the form of kernel summation and obtain more accurate intensity functions.

Computational efficiency. Figure 4 provides the run time comparison for the training between DeepSTPP and NSTPP for 100 epochs. To ensure a fair comparison, all experiments are conducted on 1 GTX 1080 Ti with Intel Core i7-4770 and 64 GB RAM. Our method is 100 times faster than NSTPP in training. It is mainly because our spatiotemporal kernel formulation has a close form of integration, which bypasses the complex and cumbersome numerical integration.

4.2. Real-World Experiment Results

For real-world data evaluation, we report the conditional spatial and temporal log-likelihoods, i.e., $\log f^*(\mathbf{s}|t)$ and $\log f^*(t)$, of the final event given the input events, respectively. The total log-likelihood, $\log f^*(s, t)$, is the summation of the two values.

Predictive performances. As our model is probabilistic, we compare against baselines models on the test predictive LL for space and time separately in Table 3. RMTPP

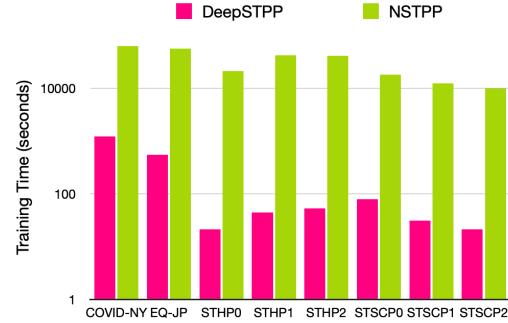


Figure 4: Log train time comparison on all datasets

Table 3: Test log likelihood (LL) comparison for space and time on real-world data over 3 runs.

LL	COVID-19 NY		Earthquake JP	
	Space	Time	Space	Time
DeepSTPP	-0.1150 ± 0.0109	2.4583 ± 0.0008	-4.4025 ± 0.0128	0.4173 ± 0.0014
NSTPP	-0.0798 ± 0.0433	2.6364 ± 0.0111	-4.8141 ± 0.1165	0.3192 ± 0.0124
RMTPP	-	2.4476 ± 0.0039	-	0.3716 ± 0.0077

can only produce temporal intensity thus we only include the time likelihood. We observe that DeepSTPP outperforms NSTPP most of the time in terms of accuracy. It takes only half of the time to train, as shown in Figure 4. Furthermore, we see that STPP models (first three rows) achieve higher LL compared with only modeling the time (RMTPP). It suggests the additional benefit of joint spatiotemporal modeling to increases the time prediction ability.

Ablation study We conduct ablation studies on the model design. Our model assumes a global latent process z that governs the parameters $\{w_i, \beta_i, \gamma_i\}$ with separate decoders. We examine other alternative designs experimentally. (1) *Shared decoders*: We use one shared decoder to generate model parameters. Shared decoders input the sampled z to one decoder and partition its output to generate model parameters. (2) *Separate process*: We assume that each of the $\{w_i, \beta_i, \gamma_i\}$ follows a separate latent process and we sample them separately. Separate processes use three sets of means and variances to sample $\{w_i, \beta_i, \gamma_i\}$ separately. (3) *LSTM encoder*: We replace the Transformer encoder with a LSTM module.

As shown in Table 4, we see that (1) *Shared decoders* decreases the number of parameters but reduces the performance. (2) *Separate process* largely increases the number of parameters but has negligible influences in test log-likelihood. (3) *LSTM encoder*: changing the encoder from Transformer to LSTM also results in slightly worse performance. Therefore, we validate the design of DeepNSTPP: we assume all distribution parameters are governed by one single hidden stochastic process with separate decoders and a Transformer as encoder.

5. Conclusion

We propose a family of deep dynamics models for irregularly sampled spatiotemporal events. Our model, Deep Spatiotemporal Point Process (DeepSTPP), integrates a principled spatiotemporal point process with deep neural networks. We derive a tractable inference procedure by modeling the space-time intensity function as a composition of kernel functions and a latent stochastic process. We infer the latent process with neural networks following the variational inference procedure. Using synthetic data from the spatiotemporal Hawkes process and self-correcting process, we show that our model can learn the spatiotemporal intensity accurately and efficiently. We demonstrate superior forecasting performance on many real-world benchmark spatiotemporal event datasets. Future work include further considering the mutual-exciting structure in the intensity function, as well as modeling multiple heterogeneous spatiotemporal processes simultaneously.

Acknowledgments

This work was supported in part by Adobe Data Science Research Award, U.S. Department Of Energy, Office of Science, U. S. Army Research Office under Grant W911NF-20-1-0334, and NSF Grant 2134274.

References

Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific reports*, 8(1):1–12, 2018.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in neural information processing systems*, pages 6571–6583, 2018.

Ricky TQ Chen, Brandon Amos, and Maximilian Nickel. Neural spatio-temporal point processes. *ICLR*, 2021.

Junyoung Chung, Caglar Gülcöhre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR*, abs/1412.3555, 2014.

Daryl J Daley and David Vere-Jones. *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media, 2007.

Edward De Brouwer, Jaak Simm, Adam Arany, and Yves Moreau. Gru-ode-bayes: Continuous modeling of sporadically-observed time series. In *Advances in Neural Information Processing Systems*, pages 7379–7390, 2019.

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 *KDD*, pages 1555–1564, 2016.

Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. In *Advances in Neural Information Processing Systems*, pages 3140–3150, 2019.

Shen Fang, Qi Zhang, Gaofeng Meng, Shiming Xiang, and Chunhong Pan. Gstnet: Global spatial-temporal network for traffic flow prediction. In *IJCAI*, pages 2286–2293, 2019.

Chris Finlay, Jörn-Henrik Jacobsen, Levon Nurbekyan, and Adam M Oberman. How to train your neural ode. *arXiv preprint arXiv:2002.02798*, 2020.

Xu Geng, Yaguang Li, Leye Wang, Lingyu Zhang, Qiang Yang, Jieping Ye, and Yan Liu. Spatiotemporal multi-graph convolution network for ride-hailing demand forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3656–3663, 2019.

Amir Gholami, Kurt Keutzer, and George Biros. Anode: Unconditionally accurate memory-efficient gradients for neural odes. *arXiv preprint arXiv:1902.10298*, 2019.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

Valerie Isham and Mark Westcott. A self-correcting point process. *Stochastic processes and their applications*, 8(3):335–347, 1979.

Junteng Jia and Austin R Benson. Neural jump stochastic differential equations. In *NeurIPS*, pages 9847–9858, 2019.

Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *NeurIPS*, 2020.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *ICLR*, 2018.

Wenwei Liang, Wei Zhang, and Xiaoling Wang. Deep sequential multi-task modeling for next check-in time and location prediction. In *International Conference on Database Systems for Advanced Applications*, pages 353–357. Springer, 2019.

Scott Linderman and Ryan Adams. Discovering latent network structure in point process data. In *International conference on machine learning*, pages 1413–1421. PMLR, 2014.

Hongyuan Mei and Jason Eisner. The neural hawkes process: A neurally self-modulating multivariate point process. In *NeurIPS*, 2017.

George O Mohler, Martin B Short, P Jeffrey Brantingham, Frederic Paik Schoenberg, and George E Tita. Self-exciting point process modeling of crime. *Journal of the American Statistical Association*, 106(493):100–108, 2011.

Jesper Moller and Rasmus Plenge Waagepetersen. *Statistical inference and simulation for spatial point processes*. CRC Press, 2003.

Michael C Mozer, Denis Kazakov, and Robert V Lindsey. Discrete event, continuous time rnns. *arXiv:1710.04110*, 2017.

Alexander Norcliffe, Cristian Bodnar, Ben Day, Jacob Moss, and Pietro Lio. Neural ode processes. *ICLR*, 2021.

Yosihiko Ogata. On lewis’ simulation method for point processes. *IEEE transactions on information theory*, 27(1):23–31, 1981.

Maya Okawa, Tomoharu Iwata, Takeshi Kurashima, Yusuke Tanaka, Hiroyuki Toda, and Naonori Ueda. Deep mixture point processes: Spatio-temporal event prediction with rich contextual information. In *KDD*, pages 373–383, 2019.

Kira Rehfeld, Norbert Marwan, Jobst Heitzig, and Jürgen Kurths. Comparison of correlation analysis techniques for irregularly sampled time series. *Nonlinear Processes in Geophysics*, 18(3):389–404, 2011.

Alex Reinhart et al. A review of self-exciting spatio-temporal point processes and their applications. *Statistical Science*, 33(3):299–318, 2018.

Abolfazl Safikhani, Camille Kamga, Sandeep Mudigonda, Sabiheh Sadat Faghih, and Bahman Moghimi. Spatio-temporal modeling of yellow taxi demands in new york city using generalized star models. *International Journal of Forecasting*, 2018.

Jin Shang and Mingxuan Sun. Geometric hawkes processes with graph convolutional recurrent neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4878–4885, 2019.

Oleksandr Shchur, Ali Caner Türkmen, Tim Januschowski, and Stephan Günnemann. Neural temporal point processes: A review. *arXiv preprint arXiv:2104.03528*, 2021.

Satya Narayan Shukla and Benjamin Marlin. Interpolation-prediction networks for irregularly sampled time series. In *ICLR*, 2018.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017.

Alejandro Veen and Frederic P Schoenberg. Estimation of space–time branching process models in seismology using an em–type algorithm. *Journal of the American Statistical Association*, 103(482):614–624, 2008.

Shuai Xiao, Mehrdad Farajtabar, Xiaojing Ye, Junchi Yan, Le Song, and Hongyuan Zha. Wasserstein learning of deep generative point process models. In *NeurIPS*, pages 3247–3257, 2017.

Shi Xingjian, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *NeurIPS*, pages 802–810, 2015.

Guolei Yang, Ying Cai, and Chandan K. Reddy. Recurrent spatio-temporal point process for check-in time prediction. In *CIKM*, pages 2203–2211. ACM, 2018.

Huaxiu Yao, Xianfeng Tang, Hua Wei, Guanjie Zheng, and Zhenhui Li. Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5668–5675, 2019.

Qiang Zhang, Aldo Lipani, Omer Kirnap, and Emine Yilmaz. Self-attentive hawkes processes. In *International Conference on Machine Learning (ICML)*, 2020.

Liang Zhao, Qian Sun, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. Multi-task learning for spatio-temporal event forecasting. In *KDD*, pages 1503–1512, 2015.

Shixiang Zhu, Shuang Li, Zhigang Peng, and Yao Xie. Imitation learning of neural spatio-temporal point processes. *arXiv preprint arXiv:1906.05467*, 2019.

Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. Transformer hawkes process. *International Conference on Machine Learning (ICML)*, 2020.