

Neural Spatiotemporal Point Processes: Trends and Challenges

Anonymous authors

Paper under double-blind review

Abstract

Spatiotemporal point processes (STPPs) are probabilistic models for events occurring in continuous space and time. Real-world event data often exhibits intricate dependencies and heterogeneous dynamics. By incorporating modern deep learning techniques, STPPs can model these complexities more effectively than traditional approaches. Consequently, the fusion of neural methods with STPPs has become an active and rapidly evolving research area. In this review, we categorise existing approaches, unify key design choices, and explain the challenges of working with this data modality. We further highlight emerging trends and diverse application domains. Finally, we identify open challenges and gaps in the literature.

1 Introduction

Real-world events—such as urban crime incidents, epidemic spread, earthquakes, and environmental changes—can be represented as sequences of discrete events with both spatial and temporal components. Studying the spatiotemporal distribution of events and discovering the relationships among different types of events is an increasingly important area of research for understanding the dynamics and mechanisms of the occurrence of events. One such paradigm is the *spatiotemporal point process* (STPP) model, defined as a stochastic process that describes the spatial and temporal distribution of discrete events (Daley and Vere-Jones, 2007), which is well suited to capture the complex relationships between events, including self-excitation, and the interactions between events and spatial covariates across time and space.

Reviews on neural point processes have focused on modelling the temporal dynamics of events using neural networks (NNs) (Shchur et al., 2021; Lin et al., 2021; 2022). In contrast, reviews that address spatiotemporal event modelling have focused on traditional statistical methods (González et al., 2016; Reinhart, 2018). More recently, some reviews have explored specific aspects of the use of machine learning in STPPs (Wikle and Zammit-Mangion, 2023; Bernabeu et al., 2024), these typically focus on specific aspects rather than providing comprehensive coverage. Meanwhile, work on neural STPPs has achieved remarkable progress.

Our work bridges this gap by systematically exploring design choices, methodological innovations, and key challenges in neural STPPs. To our knowledge, no prior survey has comprehensively examined these aspects in this context.

1.1 Why We Need Neural STPPs

Structural differences between space and time make modelling challenging. Time is unidirectional, while spatial propagation is omnidirectional and is affected by environmental factors. Traditional methods rely on strong parametric assumptions and independence, limiting flexibility. They struggle with long-range dependencies, fail to capture heterogeneous dynamics across space and time, and cannot integrate multimodal data. Additionally, single-step predictions cause error accumulation. NN-based methods overcome these limitations by encoding space and time efficiently, handling dependencies, and learning heterogeneous patterns. They automate feature extraction, integrate diverse data, and scale effectively.

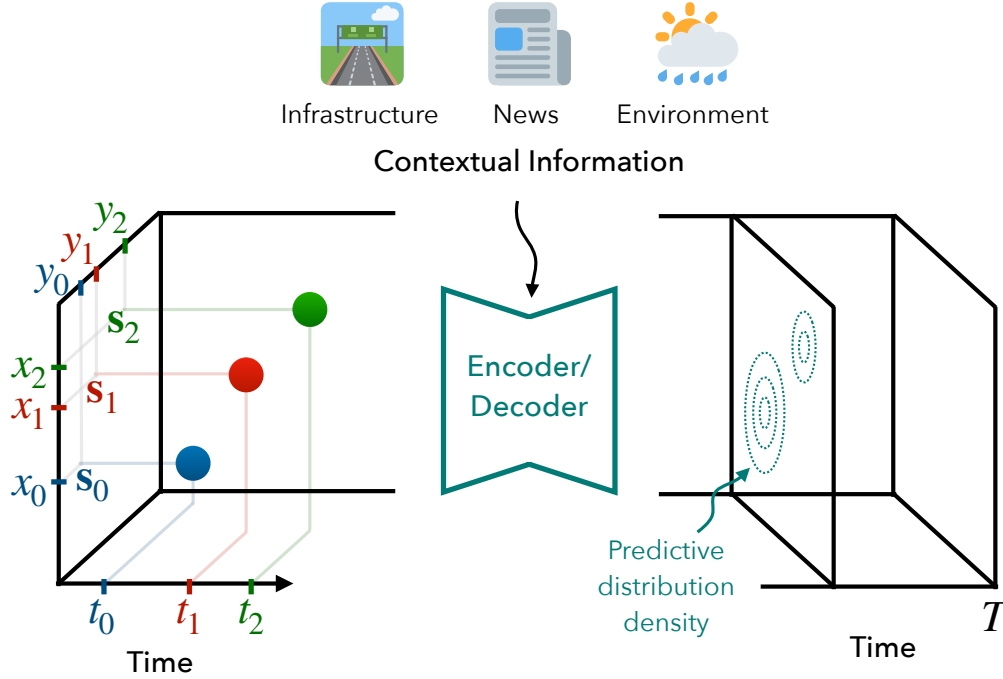


Figure 1: Schematic of the autoregressive construction of an STPP with two spatial dimensions. Three events are fed into an NN that predicts the likelihood of future event times and locations.

1.2 Scope and Structure of the Paper

This survey reviews neural STPPs, covering core models, applications, and key components in event modelling. We focus on studies using point processes with neural parameterisation to capture spatiotemporal dynamics. Our literature search included keyword-based queries, citation tracking, and seminal works in neural temporal point processes. By outlining fundamental principles and design choices, we provide a practical foundation for researchers. We first introduce necessary background and notation, then present available modelling choices, including architectures, training procedures, and metrics. The next section highlights notable applications of neural STPPs, and we conclude by discussing open challenges.

2 Background and Notation

2.1 Spatiotemporal Point Processes

STPPs are concerned with modelling sequences of random events in continuous space and time (Møller and Waagepetersen, 2003). A realisation or sample of an STPP is defined up to a *time horizon* $T \in \mathbb{R}_{\geq 0}$. It is a finite, ordered *event sequence* containing pairs $X = [(t_1, \mathbf{s}_1), (t_2, \mathbf{s}_2), \dots, (t_n, \mathbf{s}_n)]$, where $t_i \in [0, T]$ denotes the time and $\mathbf{s}_i \in \mathcal{S}$ the location of event i . Here, $\mathcal{S} \subseteq \mathbb{R}^d$ represents the spatial domain, which is typically a bounded region in d -dimensional Euclidean space and typically $d = 2$. For a given event sequence, $\mathcal{H}_t = \{(t_i, \mathbf{s}_i) \mid t_i < t\}$ denotes the history of the current realisation up to (but excluding) time point t (c.f. Figure 1).

2.2 Likelihood

An STPP can be specified by defining the likelihood of event sequences. The framework follows the *temporal priority principle* (Berzuini et al., 2012), which asserts that all causes must precede their effects. As a result,




Figure 2: A timeline of the reviewed methodological and application-focused works, along with an overview of neural architectures and application domains. Papers are grouped based on how they encode spatial information, though in some cases these categorisations are not strictly defined.

the likelihood of an event sequence X can be expressed in an auto-regressive form (Rasmussen, 2018):

$$f(X) = \underbrace{\left(\prod_{i=1}^n f^{\text{pred}}(t_i, \mathbf{s}_i \mid \mathcal{H}_{t_i}) \right)}_{\text{Likelihood of observed events}} \cdot \underbrace{(1 - F^{\text{pred}}(T \mid \mathcal{H}_{t_n}))}_{\text{Probability of no events after } t_n},$$

where $f^{\text{pred}}(t, \mathbf{s} \mid \mathcal{H}_t)$ is the *predictive distribution*, specifying the conditional probability density function (PDF) for the next event occurring at a given timestamp t and location \mathbf{s} , given the history of past events \mathcal{H}_t . The term $F^{\text{pred}}(T \mid \mathcal{H}_{t_n})$ represents the cumulative distribution function (CDF) of the predictive distribution, which gives the probability that an event occurs before or at time T , regardless of \mathbf{s} . Thus, $1 - F^{\text{pred}}(T \mid \mathcal{H}_{t_n})$ is the probability of no events occurring after the last observed event. Conveniently, this formulation also provides a principled approach for simulation (i.e., the generation of samples from the underlying stochastic model).

2.3 Intensity Function

In practice, an STPP is often described using a (*conditional*) *intensity function* (CIF) instead of a predictive distribution; they are mutually translatable (Chen et al., 2021). The CIF $\lambda(t, \mathbf{s} \mid \mathcal{H}_t)$, denoted by $\lambda^*(t, \mathbf{s})$ as a shorthand for its dependence on the \mathcal{H}_t , is defined as:

$$\lambda^*(t, \mathbf{s}) = \lim_{\Delta_s, \Delta_t \downarrow 0} \frac{\mathbb{P}(\text{Event} \in (B(\mathbf{s}, \Delta_s) \times [t, t + \Delta_t)) \mid \mathcal{H}_t)}{|B(\mathbf{s}, \Delta_s)| \Delta t}, \quad (1)$$

where $B(\mathbf{s}, \Delta_s)$ is a d -dimensional ball (a disk if $d = 2$) centered at $\mathbf{s} \in \mathcal{S}$ with radius Δ_s , and $|B(\mathbf{s}, \Delta_s)|$ denotes its volume.

Events can include additional information beyond location as a *mark* (Daley and Vere-Jones, 2007). While the core concept remains unchanged, this approach incorporates multimodal contextual data

3 Modelling Neural Spatiotemporal Point Processes

Neural STPPs model event evolution and probabilities using NNs, implicitly encoding a PDF over sequences. The chosen approach affects training efficiency, likelihood evaluation, and event generation. A common

method constructs latent representations from an event history to generate new events. Addressing the challenges of spatial encoding, we first review methods designed to address these challenges, then discuss multi-event generation. We do not focus on neural methods for temporal encoding as they are well-studied Shchur et al. (2021).

modelling choices depend on application needs. Crime models may focus on landmark influence, while earthquake prediction prioritizes accuracy. Whereas, epidemic models may emphasize county-level policy, reducing spatial granularity. Some applications require parametric assumptions for hypothesis testing, while safety-critical settings demand uncertainty quantification. These considerations shape spatial encoding strategies. While discretisation simplifies modelling, it reduces the ability to encode an inductive bias where nearby events exert stronger influence. Its performance is also sensitive to granularity, boundary effects, and local correlations. Graph-based methods like GNNs mitigate these issues by modelling inter-cell influences.

For continuous-space methods, using raw coordinates (e.g., latitude-longitude) ignores anisotropic propagation. Solutions include location-specific parametric kernels, contextual data (e.g., satellite images, geotagged text), and learned spatial embeddings. Attention models implicitly capture heterogeneous effects, while GNNs and heterogeneous kernels explicitly model spatial interactions.

3.1 History Event Encoder

Spatiotemporal event modelling requires a CIF that combines temporal evolution and spatial dependencies (c.f. Figure 1). A common approach to predicting an event (t_i, \mathbf{s}_i) is to first encode each historical event $(t_j, \mathbf{s}_j) \in \mathcal{H}_{t_i}$ as an embedding $\mathbf{e}_j = [\omega(t_j); \sigma(\mathbf{s}_j)]$. Some architectures directly use raw time t_j and space \mathbf{s}_j , while others transform them via $\omega(\cdot)$ (e.g., linear, trigonometric, or logarithmic mappings) and $\sigma(\cdot)$ (e.g., linear layers or one-hot encodings). A history encoder then processes embeddings into a latent state \mathbf{h}_i to parameterise the CIF.

RNN variants like GRU and LSTM are common history encoders, updating states as $\mathbf{h}_{i+1} = \text{RNN}(\mathbf{e}_i, \mathbf{h}_i)$ (Du et al., 2016; Omi et al., 2019; Shchur et al., 2020). Their sequential nature reduces storage needs but limits parallelisation. Additionally, they suffer from gradient vanishing and long-term memory loss (Le and Zuidema, 2016).

Attention mechanisms (Vaswani, 2017) overcome several limitations of recurrent encoders. They have demonstrated superior performance as history encoders for temporal point processes (Zhang et al., 2020; Zuo et al., 2020) and lately in STPPs (Zhou et al., 2022; Yuan et al., 2023). Nevertheless, the $\mathcal{O}(N^2)$ space complexity required to construct the attention matrix can pose practical challenges when dealing with very long event sequences.

3.2 Single Event Prediction

Neural STPPs predict the time and location of the next event. We review parametric, neural, mixture, diffusion, and continuous-time models, emphasising those that capture spatial heterogeneity, contextual influences, and graph-based interactions.

3.2.1 Kernel-Based Methods

Kernel-based methods model spatiotemporal dependencies by parameterising event influence through kernels. When combined with NNs, these kernels become more flexible while maintaining interpretability. The intensity function is typically defined as follows:

$$\lambda^*(t, \mathbf{s}) = \mu(t, \mathbf{s}) + \sum_{(t', \mathbf{s}') \in \mathcal{H}_t} K(t', t, \mathbf{s}', \mathbf{s}), \quad (2)$$

where $\mu(t, \mathbf{s})$ is the baseline rate (which can potentially vary over time and space), and $K(\cdot)$ models past event influence. Traditional models use stationary kernels, assuming time and space invariant relationships (e.g., ETAS (Musmeci and Vere-Jones, 1992), (Hawkes, 1971)), making strong parametric assumptions and missing heterogeneous effects. More expressive approaches learn spatial and temporal dependencies jointly or independently, improving adaptability.

Nonstationary kernels. Designing the influence kernel $K(\cdot)$ in Equation equation 2 is crucial for capturing how past events trigger future occurrences. Neural parameterisations of $K(\cdot)$ enable non-stationary dependencies that vary across space, time, and contextual marks, moving beyond fixed parametric forms.

A representative approach is the Gaussian mixture model of Zhu et al. (2021b), where $K(\cdot)$ is expressed as:

$$K(t', t, \mathbf{s}', \mathbf{s}) = \sum_{l=1}^L \phi_{\mathbf{s}'}^{(l)} \cdot g(t, t', \mathbf{s}, \mathbf{s}' \mid \Sigma_{\mathbf{s}'}^{(l)}, \mu_{\mathbf{s}'}^{(l)}). \quad (3)$$

A neural network embeds spatial coordinates \mathbf{s} to generate location-specific parameters $\mu_{\mathbf{s}}^{(l)}$, $\Sigma_{\mathbf{s}}^{(l)}$, and mixture weights $\phi_{\mathbf{s}}^{(l)}$ where $l \in 1, \dots, L$ with L denoting number of mixture components. Applying constraints ensures physical interpretability and reflects heterogeneous event diffusion. Visualising learned kernels reveals region-specific influence propagation, making $K(\cdot)$ a smooth, adaptive function.

Alternate formulations consider heterogeneous interactions between events modeled using Mercer’s theorem and neural basis functions(Zhu et al., 2022), leverage low-rank decomposition and a deep non-stationary influence kernel (Dong et al., 2022), or embed events in graphs for non-Euclidean interactions (Dong et al., 2023a). These can be generalised as:

$$K(t', t, \mathbf{s}', \mathbf{s}) = \sum_{r=1}^R \sum_{l=1}^L \alpha_{rl} \psi_l(t', t) \phi_r(\mathbf{s}', \mathbf{s}), \quad (4)$$

where $\psi_l(\cdot)$ and $\phi_r(\cdot)$ are neural feature maps modelling time displacement, spatial relations, or graph connectivity, with α_{rl} scaling their contributions. L and R refer to the number of basis kernels used to decompose the kernel in Equation equation 2.

For instance, Dong et al. (2023b) and Dong and Xie (2024) use learned Gaussian bases for anisotropic geography, while Dong et al. (2024) combine a stationary Gaussian kernel with a Graph Neural Network (GNN) mark kernel to capture network constraints and landmark effects.

Context. Event dynamics are influenced by contextual covariates like georeferenced images and text. Okawa et al. (2019) define event intensity as a spatially localised mixture of kernels:

$$\lambda^*(t, \mathbf{s} \mid \mathcal{D}) = \sum_{j=1}^M f(\mathbf{u}_j, \mathbf{z}_j; \theta) K(t, \mathbf{s}, \mathbf{u}_j), \quad (5)$$

where $K(\cdot)$ is a compactly supported Gaussian kernel, \mathcal{D} is a set of contextual features, and \mathbf{u}_j are spatiotemporal anchors uniformly distributed in time and space. Contextual features \mathbf{z}_j , extracted from \mathbf{u}_j , inform mixture weights $f(\mathbf{u}_j, \mathbf{z}_j; \theta)$ through a deep network combining image and text embeddings. These weights adapt to heterogeneous conditions (e.g., urban infrastructure, social events), enabling dynamic kernel weighting centered on \mathbf{u}_j while restricting contextual influence to local neighborhoods.

Zhang et al. (2023) replace the parametric kernel with a deep kernel $K(t, \mathbf{s}, \mathbf{u}) = k_\phi(g(t, \mathbf{s}), g(\mathbf{u}))$, where $g(\cdot)$ is a non-linear transformation by a deep NN, learning complex spatial correlations beyond standard Euclidean distance.

Okawa et al. (2022) proposes a method for incorporating high-dimensional contextual data into the Hawkes process. They introduce a weighting term in the excitation kernel, extracting relevant features using CNNs and employing continuous kernel convolution to transform discretised features into continuous space. This enables capturing spatial heterogeneity and external influences while ensuring tractable optimisation.

Semi-parametric and non-parametric kernels. (Zhou et al., 2022) introduce a non-parametric mixture-based intensity given by

$$\lambda^*(\mathbf{s}, t \mid \mathbf{z}) = \sum_{i=1}^{n+J} w_i K_{\mathbf{s}}(\mathbf{s}, \mathbf{s}_i; \gamma_i) K_t(t, t_i; \beta_i), \quad (6)$$

where each kernel $K_s(\cdot)$ and $K_t(\cdot)$ is a normalised radial basis function, and the parameters $\{w_i, \gamma_i, \beta_i\}$ are drawn from a latent process \mathbf{z} . The history is encoded by a Transformer 3.1 and decoded to parameterise the latent process. The samples of the latent process are further decoded via a feedforward network. This method captures uncertainty in the timing and spatial locations of events. By augmenting the observed n events with J randomly sampled representative points, this approach addresses global background intensity, thereby reducing the reliance on strong parametric assumptions.

Neural STPP models often rely on restrictive assumptions, such as conditional independence or unilateral dependence between the distributions of temporal and spatial events. These assumptions limit their ability to accurately predict events in real-world scenarios, where events exhibit complex interdependencies in both time and space. To address this, Yuan et al. (2023) proposed a framework that jointly models spatiotemporal event distributions via a diffusion-based approach without structural constraints. The model employs a spatiotemporal encoder that separately embeds time and space, fuses these into spatiotemporal representations using self-attention, and conditions a diffusion model on these hidden states. The diffusion process iteratively denoises event coordinates using a co-attention network that dynamically captures cross-dependent spatiotemporal interactions, enabling joint distribution learning without assuming independence or requiring integrable intensity functions. For spatial decoding, the model directly predicts continuous coordinates or can apply a rounding step for discrete locations. This method eliminates the need for approximation during sampling and supports continuous and discrete spatial domains.

Despite the inherent randomness in event times and locations, many STPP models offer only point predictions, lacking principled uncertainty quantification. This gap is especially problematic in marked STPPs, where reliable confidence scores for discrete event marks are essential. To address these challenges, Li et al. (2024) introduce a score-matching objective for estimating the pseudo-likelihood of marked STPPs, overcoming issues with intractable integral calculations while also providing uncertainty estimates. The score function represents the gradient of the logarithm of the conditional spatial distribution. They use the same CDN architecture as the backbone as Yuan et al. (2023) to predict the score function. Langevin dynamics is employed to sample event locations by iteratively refining draws according to the learned score. Thresholding the resulting sample density yields confidence regions for event locations, while a similar procedure provides confidence intervals for event times. The proposed framework not only predicts future events accurately, but also quantifies uncertainty, offering robust confidence bounds for both the event timing and discrete marks.

3.2.2 Continuous Time-Based Methods

The work by Chen et al. (2021) employs Continuous-time Normalising Flows (CNF), which is based on Neural Ordinary Differential Equations (Chen et al., 2018). This method separates the conditional CIF into two components: a temporal component and a spatial conditional distribution. Subsequent developments of CNFs further refine the spatial conditional distribution. For instance, the Time-Varying CNF introduces continuous flow transformations to handle time-varying observational data; however, it does not explicitly incorporate conditioning on the event history.

To address this limitation, Jump CNFs combine insights from Jia and Benson (2019) with CNFs, integrating discrete state updates at event times. This efficiently models abrupt dynamic changes with a computational cost of $\mathcal{O}(N)$ for N past events. To improve scalability for long event histories, Attentive CNFs use Transformer-based attention, enabling parallel trajectory computation while preserving non-trivial dependencies, achieving a balance between efficiency and representational power.

3.3 Multi-Event Prediction

Single-event prediction methods rely on stepwise CIF estimation, making multi-event forecasting computationally prohibitive in high-dimensional spaces due to repeated integration. Sequential history updates also propagate errors, leading to degraded accuracy over iterations. Multi-event prediction addresses these issues by jointly estimating event distributions, eliminating reliance on sequential updates. Transformer-based architectures enable parallelised predictions, bypassing error accumulation and improving efficiency.

Table 1: Common Evaluation Metrics for Neural Spatiotemporal Point Processes

Name	Definition	Advantages
NLL	Negative Log-Likelihood (likelihood of observed data given the predicted distribution).	Directly tied to MLE-based training; widely adopted for distributional fit.
HD	Hellinger Distance (a measure of distance between two probability distributions).	Fine-grained assessment of similarity; often used with known ground-truth distributions.
MMD	Maximum Mean Discrepancy (kernel-based comparison of two sample sets).	Distribution-agnostic; captures higher-order differences in generative quality.
MAE / MSE / RMSE	Mean/Absolute/Squared/Root Errors for point predictions (time or space).	Straightforward and interpretable; highlight large errors (MSE/RMSE) or typical errors (MAE).
MAPE	Mean Absolute Percentage Error (ratio-based prediction error).	Scale-independent; intuitive for relative errors across different magnitudes.
SL	Sequence Length (Wasserstien Distance between the categorical distribution of event sequence length).	Suitable for applications where event count within an interval outweighs spatial and temporal accuracy.
CD	Counting Distance (A generalisation of Wasserstein Distance for order TPPs to STPP using L1 distance)	Useful for evaluating the generative performance of model, especially for multi-step predictions
CS / ECE	Calibration Score / Expected Calibration Error (comparison of predicted vs. observed confidence).	Evaluates how well probability estimates reflect true event frequencies (calibration).

3.3.1 Multi-Step Decoders

Erfanian et al. (2022) proposes a new Transformer-based architecture augmented with normalising flows and probabilistic layers, which outputs a batch of L events based on the history of previous n events. In this approach, a separate spatial and temporal encoding is learned for events ranging from $n + 1$ to $n + L$, which are then injected into a probabilistic layer through a learned mapping between the parameters of an exponential distribution for time and a multivariate Gaussian distribution for space. Considering these as the base distributions they are passed into normalising flows, which convert them into more expressive distributions to model the joint distribution of batched events. The parameters of all distributions are learned independently for $l \in [n + 1, \dots, n + L]$. Even though separate distributions are used for space and time, the hidden state used for both as inputs considers spatiotemporal interdependencies.

3.3.2 Sampling Full Sequences

Another approach models the probability of the entire sequence or point set, rather than modelling the inter-event time and spatial distribution given history. This approach addresses a key problem with CIF parameterisation methods. In this modelling framework, the entire event sequence is embedded in the analysis. Moving away from the CIF parameterisation offers added benefits for modelling events, such as data imputation and multi-event prediction. Lüdke et al. (2023) model entire TPPs using diffusion models. Superposition and thinning properties of the CIF define the noising (forward) and denoising (backward) methods. In a recent study, Lüdke et al. (2025) extends this work to point processes defined on general metric spaces while generalising to order space, like STPPs. This work further develops the method, enabling flexible conditioning on various tasks without the need for explicit task-specific training. The approach separates the thinning and superposition operations into two independent processes. This separation allows the diffusion process to be defined as a stochastic interpolation between two point sets, entirely independent of the intensity function.

4 Parameter Estimation and Inference

The key objectives of event prediction include enhancing predictive performance, improving generalisation and robustness, understanding event dynamics through learned parameters, accounting for event behaviour heterogeneity, and capturing the influence of external factors. Predicting future events from a learned model requires sampling from its intensity function. Traditional statistical methods often rely on strong parametric assumptions for modelling event intensities, using techniques such as likelihood-based methods, partial likelihood, the EM algorithm, or Bayesian approaches.

In statistical inference, **Maximum Likelihood Estimation** (MLE) is most commonly used to fit classical and neural STPPs, typically by maximising the likelihood function or, equivalently, minimising the negative log-likelihood (NLL). For an observed sequence of N events, the NLL is given by (Daley et al., 2003):

$$\mathcal{L}_{NLL} = - \underbrace{\sum_{i=1}^N \log \lambda^*(t_i, \mathbf{s}_i)}_{\text{Log-likelihood of observed events}} + \underbrace{\int_{[0,T]} \int_{\mathcal{S}} \lambda^*(\tau, \mathbf{u}) d\tau d\mathbf{u}}_{\text{Expected number of events}}.$$

When using neural networks to parameterise the CIF, evaluating the integral term is typically intractable. This complexity often requires numerical methods (Chen et al., 2021) or Monte Carlo methods (Mei and Eisner, 2017) for likelihood evaluation. However, these strategies can be computationally expensive and prone to numerical errors, particularly in high-dimensional spatiotemporal domains. While certain simplifying assumptions, such as exponential decay (Du et al., 2016) and linear interpolation (Zuo et al., 2020), can lead to closed-form solutions or faster approximations, they often restrict the expressiveness of the model.

In the purely temporal settings, Zhou and Yu (2023) introduce a paradigm for efficient and non-parametric inference of TPPs. They approximate the influence function (c.f. Equation equation 2) via a NN, using **automatic integration** to compute its integral. A monotonically increasing integral network is trained, its partial derivative defines the CIF. This approach directly yields the CIF and its antiderivative from the network parameters, avoiding functional form restrictions. Building on this foundation, Zhou and Yu (2024) addresses the computational challenge of integrating the intensity function in 3D spatiotemporal domains by employing automatic integration. This approach learns an integral network, whose partial derivatives with respect to spatial and temporal inputs yield the intensity, ensuring an exact antiderivative without restricting the model’s functional form. Furthermore, a ProdNet factorisation of the influence function into 1D components enforces non-negativity while capturing spatiotemporal interactions. Maximising the log-likelihood of observed events learns a highly expressive CIF, improving spatiotemporal event prediction.

Zhou et al. (2022) integrates flexible non-parametric modelling with **amortised variational inference**. This model captures events in continuous time, capturing irregular sampling dynamics and unifying spatial and temporal dependencies. By employing a kernel-based intensity function, the approach allows for closed-form integration, addressing previously intractable likelihood computations. This design avoids computationally expensive numerical integration inherent in neural ODE-based approaches (Chen et al., 2021). Furthermore, this non-parametric approach avoids restrictive parametric assumptions. Training via amortised variational inference maximises the evidence lower bound (ELBO) of the likelihood while balancing reconstruction accuracy and posterior regularisation. The framework uses the kernel-based intensity for gradient computation, facilitating end-to-end optimisation of both encoder and decoder parameters.

Zhang et al. (2023) uses **score matching**, which minimises the Fisher divergence between the model’s log-density gradient and the data’s log-density gradient, thus bypassing the need to calculate the intractable integral. A denoising score matching (DSM) method is used, which improves stability by introducing a small amount of noise to the data, which avoids the computation of second derivatives. In contrast, Li et al. (2024) utilises a score-matching-based pseudolikelihood objective, which eliminates the need for explicit calculation of the normalising term that makes the likelihood integral intractable. The model decomposes the joint intensity function into a product of conditional distributions, allowing for the application of score-matching techniques for event times and locations, while using a conditional likelihood for event marks. This approach

is designed to overcome overconfidence and underconfidence by learning a posterior distribution that matches the actual data distribution, which also requires score-based sampling with Langevin dynamics.

Reinforcement learning (RL) frameworks provide a training approach that does not rely on likelihood calculations. Zhu et al. (2021b) employs an imitation learning framework to train their model. The learner policy is defined by a PDF associated with the CIF of the point process and parameterised by the model’s parameters. The goal is for this learner policy to replicate the expert policy reflected in the training data. The training process maximises the expected reward, determined by the Maximum Mean Discrepancy (MMD) between empirical distributions of the training data and data generated by the learner’s policy. This data-driven MMD reward function offers robustness to model mismatch as it compares data distributions instead of relying on a predefined likelihood. Additionally, the closed-form representation of the reward function enables computationally efficient optimisation via analytical gradient calculation, avoiding intensive inverse reinforcement learning.

5 Evaluation Metrics

Evaluating STPPs requires task-specific metrics, summarised in Table 1. Model fit is often assessed using **NLL**, which measures how well predicted distributions align with observed data. However, NLL prioritises overall distributional fit over individual event accuracy and can be biased due to computational approximations like Monte Carlo integration. Moreover, since NLL conditions on ground truth history, it is limited in evaluating true generative capacity.

Alternative distributional metrics address these shortcomings. **HD** (Zhou et al., 2022) directly compares learned and true distributions but requires ground truth intensity and spatial discretisation, limiting real-world use, but useful for testing the ability of models to recover known patterns. **MMD** (Zhu et al., 2021b) avoids strong distributional assumptions by comparing generated and observed sequences but is computationally expensive and sensitive to kernel selection. MMD can also be used with **CD** (Lüdke et al., 2025) as the distance measure. This improves distributional comparisons and is useful for evaluating the generative performance of STPPs. **SL**, measured via Wasserstein distance between two categorical distributions, provides additional insight into model performance by comparing event sequence lengths, especially useful when the task involves measuring case counts, such as in epidemiology or crime modelling.

For point prediction accuracy, **MAE**, **MSE**, and **RMSE** are commonly used. MAE is robust to uniform errors, while RMSE is more sensitive to large deviations. Aggregated event predictions rely on normalised MAE (NMAE) (Okawa et al., 2022) and **MAPE** (Okawa et al., 2019). While NMAE enables cross-scale comparisons, it depends on predefined spatiotemporal regions. MAPE, though intuitive, can be unstable near zero values. Mean relative error (MRE) (Dong et al., 2022) assesses intensity differences but also suffers from instability near zero. Prediction accuracy (ACC) is useful for event count estimation but, unlike metrics considering location and time, it assesses the accuracy of event counts only. Prediction accuracy (ACC) is useful for event count estimation, but doesn’t account for spatial and temporal precision.

Uncertainty quantification is crucial for robust evaluation. **CS** and **ECE** (Li et al., 2024) assess how well predicted confidence intervals align with observed distributions, though ECE’s binning requires adaptive methods for reliability.

Selecting metrics depends on the task and dataset. A comprehensive evaluation combines prediction accuracy, distributional fit, and uncertainty quantification. Recent work favours WD and MMD over NLL for assessing generative quality.

6 Applications

The existing literature on neural STPPs mainly underscores their applications in public safety and urban mobility.

6.1 Natural Disasters

Neural STPPs model earthquake and wildfire occurrences by capturing spatiotemporal dependencies. The ETAS model is widely used for earthquake forecasting, with neural extensions improving prediction. Nicolis et al. (2021) enhance ETAS using neural networks for seismic forecasting, while Zhang et al. (2024) integrate it with deep learning. For wildfires, Xu et al. (2023) incorporate environmental factors and remote sensing data to predict ignition probability and magnitude, demonstrating strong performance on California data.

6.2 Crime

STPPs are increasingly used for crime prediction. Dong and Xie (2024) model gun violence in Atlanta with a non-stationary Hawkes process, integrating socio-economic covariates to improve predictions. Dong et al. (2024) develop a spatiotemporal network point process for Valencia crime, mapping events onto streets to better capture urban dynamics.

6.3 Traffic

Zhu et al. (2021a) introduce an attention-based STPP integrating traffic sensor data and 911 call records to model congestion dynamics, capturing self-excitation and external influences. They use an NLP-inspired attention mechanism for temporal dependence and a ‘tail-up’ approach for spatial correlations on road networks. Jin et al. (2023) propose a congestion prediction model combining GCNs for spatial dependencies, Transformers for temporal patterns, and a continuous GRU with neural flow for instantaneous traffic behaviour.

6.4 Epidemiology

STPPs model infectious disease spread by capturing event transmission dynamics. Li et al. (2021) proposes an intensity-free STPP using generative adversarial imitation learning, while Dong et al. (2023b) develop a non-stationary STPP with an NN kernel to model heterogeneous COVID-19 case correlations and spatial variations.

7 Open Challenges

Neural STPPs have advanced event modelling but still face key obstacles—computational constraints, interpretability, reproducibility, and real-world applicability. Overcoming these will enable robust, scalable, and interpretable models.

7.1 Reproducibility

A significant barrier to advancing neural STPP research is the lack of standardised experimental setups and consistent baseline comparisons. In contrast to TPPs, which benefit from unified libraries like Xue et al. (2024), there is no comprehensive implementation for spatial and temporal methods in STPPs. A unified library incorporating diverse neural architectures, metrics, and training strategies would facilitate better ablation studies and benchmarking. While some perform well under specific conditions or on curated datasets, fair evaluations across architectures remain challenging without robust testing environments. Additionally, the limited availability of open-source tools further restricts wider adoption and reproducibility of these methods.

7.2 Benchmarking

A major gap in the literature is the absence of standardised benchmarking datasets. Existing datasets often suffer from selection bias, missing data, and varying granularity, complicating integration and analysis. Different datasets exhibit diverse phenomena, such as spatial heterogeneity, long-range dependencies, and

entangled spatiotemporal dynamics. Some datasets display self-exciting behaviour, while others show self-correcting patterns. The lack of a unified dataset with consistent testing conditions means benchmarking efforts are often limited to specific attributes or datasets where a method performs well. A comprehensive event database would aid in model development and evaluation. Additionally, the lack of contextual datasets hinders efforts to develop foundational models for event prediction, improving the understanding of event dynamics in data-scarce environments.

7.3 Architectures

Generative strategies and alternative training objectives represent promising yet underexplored avenues for STPP research. While GAN-based architectures and Wasserstein objectives have shown promise in temporal point processes (Xiao et al., 2017), their adaptation to spatiotemporal applications remains limited. Similarly, methods moving beyond MLE, particularly those leveraging Transformer-based models instead of RNNs, could enhance the flexibility and robustness of event modelling. Continuous-time neural TPPs, such as those proposed by (Biloš et al., 2021) and (Chen et al., 2024), have demonstrated potential for temporal modelling but have yet to be sufficiently applied to spatial event modelling. Many existing neural STPPs treat space and time independently, missing opportunities to improve temporal predictions by capturing spatial relationships. Integrating GNNs to enhance spatial modelling remains an open challenge.

7.4 Applicability

Despite improvements in predictive performance, neural STPPs are rarely utilised to inform policy decisions or design interventions. Their lack of interpretability limits their applicability in real-world scenarios, where generalizable, interpretable models with uncertainty quantification are essential. Policymakers often seek to understand the impact of interventions, necessitating models that incorporate contextual factors and causal effects in event propagation. This underscores the need for research in interpretable neural networks, neuro-symbolic methods, and counterfactual analysis for retrospective policy evaluation and deeper insights into event modelling.

7.5 Causality and Uncertainty

Policy-focused applications of neural STPPs require deeper insights into event propagation and the causal factors influencing event rates. However, causal inference and uncertainty quantification remain underexplored in this field. Bayesian methods, which could integrate expert knowledge through priors, have not been widely adopted in neural spatiotemporal event modelling. Similarly, causal representation learning must address spatiotemporal confounding, but current research has yet to fully integrate these methods into neural STPPs (Wang et al., 2024). Advancing these areas is critical for developing models that can inform decision-making and intervention design.

Ethical Statement

Point processes are used in ethically sensitive areas, such as police patrol allocation (Mohler et al., 2015), where biases may cause disparate harms (Alikhademi et al., 2022). Unintended consequences should be considered before using Neural STPPs for spatial interventions.

References

- Kiana Alikhademi, Emma Drobina, Diandra Prioleau, Brianna Richardson, Duncan Purves, and Juan E Gilbert. A review of predictive policing from the perspective of fairness. *Artificial Intelligence and Law*, pages 1–17, 2022.
- Alba Bernabeu, Jiancang Zhuang, and Jorge Mateu. Spatio-temporal hawkes point processes: A review. *Journal of Agricultural, Biological and Environmental Statistics*, pages 1–31, 2024.

- Carlo Berzuini, Philip Dawid, and Luisa Bernardinell. *Causality: Statistical perspectives and applications*. John Wiley & Sons, 2012.
- Marin Biloš, Johanna Sommer, Syama Sundar Rangapuram, Tim Januschowski, and Stephan Günnemann. Neural flows: Efficient alternative to neural odes. *Advances in neural information processing systems*, 34: 21325–21337, 2021.
- Ricky T. Q. Chen, Brandon Amos, and Maximilian Nickel. Neural spatio-temporal point processes. In *International Conference on Learning Representations*, 2021.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Yuqi Chen, Kan Ren, Yansen Wang, Yuchen Fang, Weiwei Sun, and Dongsheng Li. Contiformer: Continuous-time transformer for irregular time series modeling. *Advances in Neural Information Processing Systems*, 36, 2024.
- 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.
- Daryl J Daley, David Vere-Jones, et al. *An introduction to the theory of point processes: volume I: elementary theory and methods*. Springer, 2003.
- Zheng Dong and Yao Xie. Atlanta gun violence modeling via nonstationary spatio-temporal point processes. *arXiv preprint arXiv:2408.09258*, 2024.
- Zheng Dong, Xiuyuan Cheng, and Yao Xie. Spatio-temporal point processes with deep non-stationary kernels. *arXiv preprint arXiv:2211.11179*, 2022.
- Zheng Dong, Matthew Repasky, Xiuyuan Cheng, and Yao Xie. Deep graph kernel point processes. In *Temporal Graph Learning Workshop@ NeurIPS*, 2023a.
- Zheng Dong, Shixiang Zhu, Yao Xie, Jorge Mateu, and Francisco J Rodríguez-Cortés. Non-stationary spatio-temporal point process modeling for high-resolution covid-19 data. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 72(2):368–386, 2023b.
- Zheng Dong, Jorge Mateu, and Yao Xie. Spatio-temporal-network point processes for modeling crime events with landmarks. *arXiv preprint arXiv:2409.10882*, 2024.
- 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*, pages 1555–1564, 2016.
- Negar Erfanian, Santiago Segarra, and Maarten de Hoop. Beyond hawkes: Neural multi-event forecasting on spatio-temporal point processes. *arXiv preprint arXiv:2211.02922*, 2022.
- JA González, FJ Rodríguez-Cortés, O Cronie, and J Mateu. Spatio-temporal point process statistics: a review. *spat stat* 18: 505–544, 2016.
- Alan G Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1): 83–90, 1971.
- Junteng Jia and Austin R Benson. Neural jump stochastic differential equations. *Advances in Neural Information Processing Systems*, 32, 2019.
- Guangyin Jin, Lingbo Liu, Fuxian Li, and Jincui Huang. Spatio-temporal graph neural point process for traffic congestion event prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 14268–14276, 2023.
- Phong Le and Willem Zuidema. Quantifying the vanishing gradient and long distance dependency problem in recursive neural networks and recursive lstms. *arXiv preprint arXiv:1603.00423*, 2016.

- Shuang Li, Lu Wang, Xinyun Chen, Yixiang Fang, and Yan Song. Understanding the spread of covid-19 epidemic: A spatio-temporal point process view. *arXiv preprint arXiv:2106.13097*, 2021.
- Zichong Li, Qunzhi Xu, Zhenghao Xu, Yajun Mei, Tuo Zhao, and Hongyuan Zha. Beyond point prediction: Score matching-based pseudolikelihood estimation of neural marked spatio-temporal point process. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=CpI37NA7M0>.
- Haitao Lin, Cheng Tan, Lirong Wu, Zhangyang Gao, Stan Li, et al. An empirical study: extensive deep temporal point process. *arXiv preprint arXiv:2110.09823*, 2021.
- Haitao Lin, Lirong Wu, Guojiang Zhao, Pai Liu, and Stan Z Li. Exploring generative neural temporal point process. *arXiv preprint arXiv:2208.01874*, 2022.
- David Lüdke, Marin Biloš, Oleksandr Shchur, Marten Lienen, and Stephan Günnemann. Add and thin: Diffusion for temporal point processes. *Advances in Neural Information Processing Systems*, 36:56784–56801, 2023.
- David Lüdke, Enric Rabasseda Raventós, Marcel Kollovich, and Stephan Günnemann. Unlocking point processes through point set diffusion. In *The Thirteenth International Conference on Learning Representations*, 2025.
- 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.
- George O Mohler, Martin B Short, Sean Malinowski, Mark Johnson, George E Tita, Andrea L Bertozzi, and P Jeffrey Brantingham. Randomized controlled field trials of predictive policing. *Journal of the American statistical association*, 110(512):1399–1411, 2015.
- Jesper Moller and Rasmus Plenge Waagepetersen. *Statistical inference and simulation for spatial point processes*. CRC press, 2003.
- Fabio Musmeci and David Vere-Jones. A space-time clustering model for historical earthquakes. *Annals of the Institute of Statistical Mathematics*, 44:1–11, 1992.
- Orietta Nicolis, Francisco Plaza, and Rodrigo Salas. Prediction of intensity and location of seismic events using deep learning. *Spatial Statistics*, 42:100442, 2021.
- 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 *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 373–383, 2019.
- 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*, 111(8):2929–2950, 2022.
- Takahiro Omi, Kazuyuki Aihara, et al. Fully neural network based model for general temporal point processes. *Advances in neural information processing systems*, 32, 2019.
- Jakob Gulddahl Rasmussen. Lecture notes: Temporal point processes and the conditional intensity function. *arXiv preprint arXiv:1806.00221*, 2018.
- Alex Reinhart. A review of self-exciting spatio-temporal point processes and their applications. *Statistical Science*, 33(3):299–318, 2018.
- Oleksandr Shchur, Marin Biloš, and Stephan Günnemann. Intensity-free learning of temporal point processes. In *International Conference on Learning Representations*, 2020.

- Oleksandr Shchur, Ali Caner Türkmen, Tim Januschowski, and Stephan Günnemann. Neural temporal point processes: A review. *arXiv preprint arXiv:2104.03528*, 2021.
- A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- Kun Wang, Sumanth Varambally, Duncan Watson-Parris, Yi-An Ma, and Rose Yu. Discovering latent structural causal models from spatio-temporal data. *arXiv preprint arXiv:2411.05331*, 2024.
- Christopher K Wikle and Andrew Zammit-Mangion. Statistical deep learning for spatial and spatiotemporal data. *Annual Review of Statistics and Its Application*, 10(1):247–270, 2023.
- Shuai Xiao, Mehrdad Farajtabar, Xiaojing Ye, Junchi Yan, Le Song, and Hongyuan Zha. Wasserstein learning of deep generative point process models. *Advances in neural information processing systems*, 30, 2017.
- Chen Xu, Yao Xie, Daniel A Zuniga Vazquez, Rui Yao, and Feng Qiu. Spatio-temporal wildfire prediction using multi-modal data. *IEEE Journal on Selected Areas in Information Theory*, 4:302–313, 2023.
- Siqiao Xue, Xiaoming Shi, Zhixuan Chu, Yan Wang, Hongyan Hao, Fan Zhou, Caigao JIANG, Chen Pan, James Y. Zhang, Qingsong Wen, JUN ZHOU, and Hongyuan Mei. EasyTPP: Towards open benchmarking temporal point processes. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=PJwAkg0z7h>.
- Yuan Yuan, Jingtao Ding, Chenyang Shao, Depeng Jin, and Yong Li. Spatio-temporal diffusion point processes. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3173–3184, 2023.
- Haoyuan Zhang, Shuya Ke, Wenqi Liu, and Yongwen Zhang. A combining earthquake forecasting model between deep learning and epidemic-type aftershock sequence (etas) model. *Geophysical Journal International*, 239(3):1545–1556, 2024.
- Qiang Zhang, Aldo Lipani, Omer Kirnap, and Emine Yilmaz. Self-attentive hawkes process. In *International conference on machine learning*, pages 11183–11193. PMLR, 2020.
- Yixuan Zhang, Quyu Kong, and Feng Zhou. Integration-free training for spatio-temporal multimodal co-variate deep kernel point processes. *Advances in Neural Information Processing Systems*, 36:25031–25049, 2023.
- Zihao Zhou and Rose Yu. Automatic integration for fast and interpretable neural point processes. In *Learning for Dynamics and Control Conference*, pages 573–585. PMLR, 2023.
- Zihao Zhou and Rose Yu. Automatic integration for spatiotemporal neural point processes. *Advances in Neural Information Processing Systems*, 36, 2024.
- 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*, pages 777–789. PMLR, 2022.
- Shixiang Zhu, Ruyi Ding, Minghe Zhang, Pascal Van Hentenryck, and Yao Xie. Spatio-temporal point processes with attention for traffic congestion event modeling. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):7298–7309, 2021a.
- Shixiang Zhu, Shuang Li, Zhigang Peng, and Yao Xie. Imitation learning of neural spatio-temporal point processes. *IEEE Transactions on Knowledge and Data Engineering*, 34(11):5391–5402, 2021b.
- Shixiang Zhu, Haoyun Wang, Xiuyuan Cheng, and Yao Xie. Neural spectral marked point processes. In *International Conference on Learning Representations*, 2022.
- Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. Transformer hawkes process. In *International conference on machine learning*, pages 11692–11702. PMLR, 2020.