

Imitation Learning of Neural Spatio-Temporal Point Processes

Shixiang Zhu, Shuang Li, Zhigang Peng, Yao Xie

Abstract—We present a novel Neural Embedding Spatio-Temporal (NEST) point process model for spatio-temporal discrete event data and develop an efficient imitation learning (a type of reinforcement learning) based approach for model fitting. Despite the rapid development of one-dimensional temporal point processes for discrete event data, the study of spatial-temporal aspects of such data is relatively scarce. Our model captures complex spatio-temporal dependence between discrete events by carefully design a mixture of heterogeneous Gaussian diffusion kernels, whose parameters are parameterized by neural networks. This new kernel is the key that our model can capture intricate spatial dependence patterns and yet still lead to interpretable results as we examine maps of Gaussian diffusion kernel parameters. The imitation learning model fitting for the NEST is more robust than the maximum likelihood estimate. It directly measures the divergence between the empirical distributions between the training data and the model-generated data. Moreover, our imitation learning-based approach enjoys computational efficiency due to the explicit characterization of the reward function related to the likelihood function; furthermore, the likelihood function under our model enjoys tractable expression due to Gaussian kernel parameterization. Experiments based on real data show our method's good performance relative to the state-of-the-art and the good interpretability of NEST's result.

Index Terms—Spatio-temporal point processes, Generative model, Imitation learning

1 INTRODUCTION

Spatio-temporal event data has become ubiquitous, emerging from various applications such as social media data, crime events data, social mobility data, and electronic health records. Such data consist of a sequence of times and locations that indicate when and where the events occurred. Studying generative models for discrete events data has become a hot area in machine learning and statistics: it reveals the pattern of the data, helps us to understand the data dynamic and information diffusion, as well as serves as an essential step to enable subsequent machine learning tasks.

Point process models (see [1] for an overview) have become a common choice for generative models of discrete event data. In particular, the self and mutual exciting processes, also known as the Hawkes processes, are popular since they can capture past events' influence on future events over time, space, and networks.

Despite the rapid development of one-dimensional temporal point processes models for discrete event data, studies focusing on *spatial-temporal* aspects of such data are relatively scarce. The original works of [2, 3] develop the so-called epidemic-type aftershock sequence (ETAS) model, which is still widely used, suggesting an exponential decaying diffusion kernel function. This model captures the seismic activities' mechanism and is convenient to fit, as the kernel function is homogeneous at all locations with the same

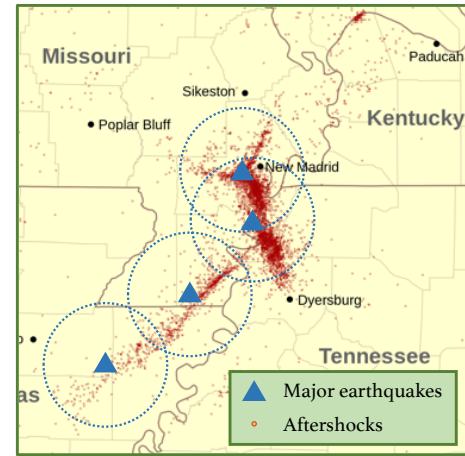


Fig. 1: A motivating example of seismic activities: four earthquakes and their aftershocks occurred in New Madrid, MO., in the United States since 1811. The blue triangles represent the major earthquakes, and the dotted circles represent the estimated aftershock regions suggested by the ETAS. The small red dots represent the actual aftershocks caused by the major earthquakes. We can observe that the locations of actual aftershocks are related to the geologic structure of faults, and the vanilla ETAS model may not sufficiently capture such complex spatial dependence.

oval shape (Fig. 2). However, these classical models for spatio-temporal event data (usually statistical models in nature) tend to make strong parametric assumptions on the conditional intensity.

However, in specific scenarios, the simplified spatio-temporal models based on the ETAS may lack flexibility.

• Shixiang Zhu, Shuang Li, and Yao Xie were with H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, 30332. Zhigang Peng was with School of Earth and Atmospheric Sciences, Georgia Institute of Technology, Atlanta, GA, 30332

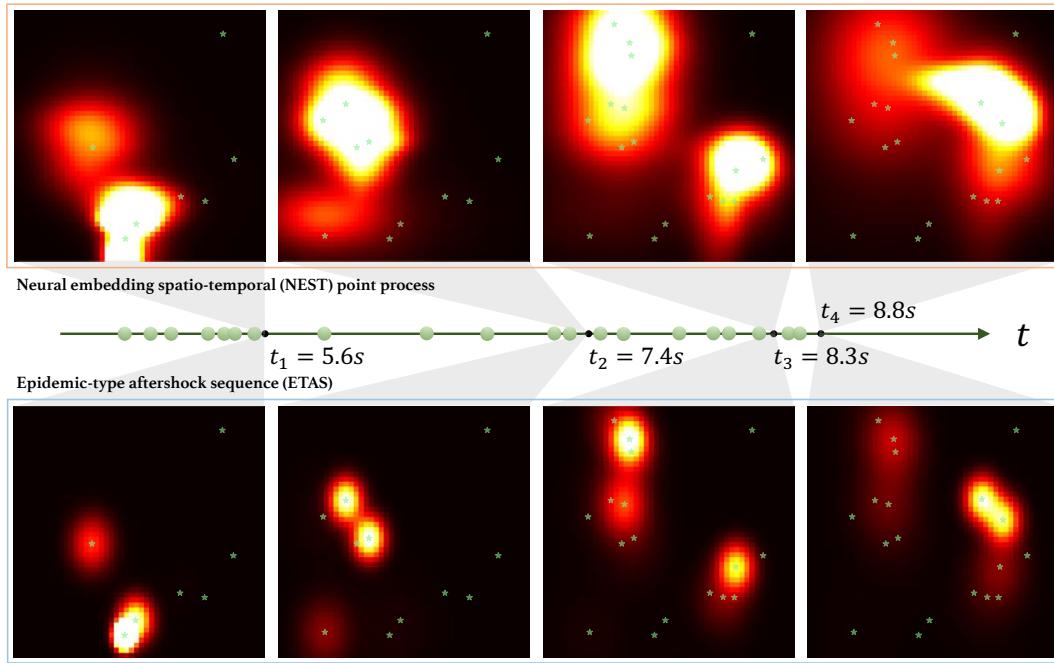


Fig. 2: Comparison between the standard epidemic-type aftershock sequence (ETAS) and the proposed neural embedding spatio-temporal (NEST) point process models, fitted to the same set of police 911 calls-for-service data. We visualize snapshots of the conditional intensities $\lambda^*(t, s), s \in \mathcal{S}$ at four times for the same sequence of events that occurred in a square region $\mathcal{S} = [-1, +1] \times [-1, +1]$; brighter areas are higher intensities where the next event is more likely to occur. The green dots on the timeline represent when the events occur, and the green stars in the square region represent the events' locations. Snapshots are taken at t_1, t_2, t_3 , and t_4 , which are indicated by the black dots on the timeline. As self-exciting point processes, the occurrence of a new event will abruptly increase the local intensity, and the influence will decay and diffuse to the surrounding regions over time. As we can see from the comparison, the NEST can capture more intricate spatial dynamics than the ETAS.

It does not capture the anisotropic spatial influence or other intricate spatial dependence structure. Take earthquake event data as an example, consisting of a sequence of records of seismic activities: their times and locations. The aftershocks are minor seismic activities that are triggered by the major earthquakes. According to the study [4], it has been shown that most of the recent seismic events that occurred in New Madrid, MO, are aftershocks of four earthquakes of magnitude 7.5 in 1811. As shown in Fig. 1, the distribution of the minor seismic activities is in a complex shape (clearly not “circles” or isotropic), and the spatial correlation between seismic activities is related to the geologic structure of faults through the complex physical mechanism and usually exhibits a heterogeneous conditional intensity. For instance, most aftershocks occur along the fault plane or other faults within the volume affected by the mainshock’s strain. This creates a spatial profile of the intensity function that we would like to capture through the model, such as the direction and shape of the intensity function at different locations, to provide useful information to geophysicists’ scientific study.

On the other hand, when developing spatio-temporal models, we typically want to generate some statistical interpretations (e.g., temporal correlation, spatial correlation), which may not be easily derived from a complete neural network model. Thus, a generative model based on specifying

the conditional intensity of point process models is a popular approach. For example, recent works [5, 6, 7, 8, 9, 10, 11] have achieved many successes in modeling temporal event data (some with marks) which are correlated in time. It remains an open question on extending this type of approach to include the spatio-temporal point processes. One challenge is how to address the computational challenge associated with evaluating the log-likelihood function (see expression (3)). This can be intractable for the general models without a carefully crafted structure since it requires integrating the conditional intensity function in a continuous spatial and temporal space. Another challenge is regarding how to develop robust model-fitting methods that do not rely too much on the modeling assumption.

This paper first presents a novel neural embedding spatio-temporal (NEST) point process for spatio-temporal discrete event data. Our proposed NEST model tackles flexible representation for complex spatial dependence, interpretability, and computational efficiency, through meticulously designed neural networks capturing spatial information. We generalize the idea of using a Gaussian diffusion kernel to model spatial correlation by introducing the more flexible heterogeneous mixture of Gaussian diffusion kernels with shifts, rotations, and anisotropic shapes. Such a model can still be efficiently represented using a handful of parameters (compared with a full neural network model such as convolutional neural

networks (CNN) over space). The Gaussian diffusion kernels are parameterized by neural networks, which allows the kernels to vary continuously over locations. This is the key that our model can capture intricate spatial dependence patterns and yet still lead to interpretable results as we examine maps of Gaussian diffusion kernel parameters. As shown in Fig. 2, our model is able to represent arbitrary diffusion shape at different locations in contrast to the ETAS developed by [12, 2, 3].

Second, we develop computationally efficient approaches for fitting the NEST model based on imitation learning (IL) [13, 7], and compare it with the maximum likelihood estimation (MLE) method. The imitation learning model fitting for the NEST is more flexible and robust. It directly measures the divergence between the empirical distributions between the training data and the model-generated data. Moreover, our imitation learning-based approach enjoys computational efficiency due to the explicit characterization of the reward function related to the likelihood function; furthermore, the likelihood function under our model enjoys tractable expression due to Gaussian kernel parameterization. Experiments based on synthetic and real data show our method's superior performance in event prediction and model interpretability compared to the state-of-the-art.

The rest of the paper is organized as follows. We start by reviewing related literature. Then Section 2 introduces the background of spatio-temporal point processes and related classic models. We present our proposed NEST model in Section 3 and describe our imitation learning framework for model fitting in Section 4. Section 5 contains experimental results based on synthetic and real data. Finally, Section 6 concludes the paper with discussions.

1.1 Related work

Existing literature on spatial-temporal point process modeling usually makes simplified parametric assumptions on the conditional intensity based on the prior knowledge of the processes. Landmark works [3, 2] suggest an exponential decaying kernel function. This model captures seismic activities' mechanism to a certain degree and is easy to learn, as the kernel function is homogeneous at all locations. However, in applications to other scenarios, such a parametric model may lack flexibility.

Another approach to obtaining generative models for temporal point processes is based on the idea of imitation learning and reinforcement learning. Good performance has been achieved for modeling temporal point processes [7] and marked temporal point processes [8]. The premise is to formulate the generative model as a policy for reinforcement learning and extract policy from sequential data as if it were obtained by reinforcement learning [14] followed by inverse reinforcement learning [15, 16]. In this way, the generative model is parameterized by neural networks [5, 17, 6, 18]. However, representing the conditional intensity entirely using neural networks may lack certain interpretability. Compressing all the information by neural networks may also miss the opportunity to consider prior knowledge about the point processes. Also, it remains an open question on how to extend this approach to include the spatial component. The spatial-temporal point processes are significantly more

challenging to model than the temporal point processes since the spatio-temporal dependence is much more intricate than the one-dimensional temporal dependence.

Other works such as [19, 20] have achieved some success in modeling complicated spatial patterns of crime by considering the spatial influence as *hotspots* with full parametric models. Some works [21, 22, 23, 10] consider the event sequences as temporal point processes without incorporating spatial information leverages non-parametric approaches. As a compromise, [24, 25, 6] seek to learn the spatial and temporal pattern jointly by multivariate point processes with discretized the spatial structure.

2 BACKGROUND

This section revisits the definitions of the spatio-temporal point processes (STPP) and one of the most commonly used STPP model – epidemic-type aftershock sequence (ETAS).

2.1 Spatio-temporal point processes (STPP)

STPP consists of an ordered sequence of events in time and location space. Assume the time horizon is $[0, T]$ and the data is given by $\{a_1, a_2, \dots, a_{N(T)}\}$, which are a set of sequences of events ordered in time. Each a_i is a spatio-temporal tuple $a_i = (t_i, s_i)$, where $t_i \in [0, T]$ is the time of the event and $s_i \in \mathcal{S} \subseteq \mathbb{R}^2$ is the associated location of the i -th event. We denote by $N(T)$ the number of the events in the sequence between time $[0, T]$ and in the region \mathcal{S} .

The joint distribution of a STPP is completely characterized by its conditional intensity function $\lambda(t, s|\mathcal{H}_t)$. Given the event history $\mathcal{H}_t = \{(t_i, s_i) | t_i < t\}$, the intensity corresponds to the probability of observing an event in an infinitesimal around (t, s) :

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

where $N(A)$ is the counting measure of events over the set $A \subseteq [0, T] \times \mathcal{S}$, $B(s, \Delta s)$ denotes a Euclidean ball centered at s with radius Δs , and $|\cdot|$ is the Lebesgue measure. Below, for the notational simplicity, we denote the conditional intensity function $\lambda(t, s|\mathcal{H}_t)$ as $\lambda^*(t, s)$.

For instance, a type of self-exciting point processes, Hawkes processes [26] has been widely used to capture the mutual excitation among temporal events. Assuming that influence from past events are linearly additive for the current event, the conditional intensity function of a Hawkes process is defined as

$$\lambda(t|\mathcal{H}_t) = \lambda_0 + \sum_{t_i < t} \nu(t - t_i),$$

where $\lambda_0 \geq 0$ is the background intensity of events, $\nu(\cdot) \geq 0$ is the *triggering function* that captures temporal dependencies of the past events. The triggering function can be chosen in advance, for instance, in the one-dimensional case $\nu(t - t_i) = \alpha \exp\{-\beta(t - t_i)\}$.

2.2 ETAS model

The most commonly used kernel function for spatio-temporal point processes is the standard *diffusion kernel* function proposed by epidemic type aftershock-sequences (ETAS) [27], which was initially introduced to model the earthquake events, but now widely used in many other applications [2, 3, 24, 21, 23, 22]. ETAS model assumes that the influence over time and space decouples, and the influence decays exponentially over time. Over space decay only depends on distance (thus, it is a spherical model). Therefore, *The ETAS model does not capture the anisotropic shape of the kernel*. This is a simplification and may not capture complex spatial dependence. The ETAS model can also deal with scalar-valued marks (e.g., the magnitude of earthquakes), which we will not discuss here while only focusing on spatio-temporal interactions between events. We also note that one of the reasons that the ETAS is a popular model is due to its interpretability.

3 PROPOSED MODEL

To capture the complex and heterogenous spatial dependence in discrete events, we present a novel *continuous-time and continuous-space* point process model, called neural embedding spatio-temporal (NEST) model. The NEST uses flexible neural networks to represent the conditional intensity's spatial heterogeneity while retaining interpretability as a semi-parametric statistical model.

3.1 Spatially heterogeneous Gaussian diffusion kernel

We start by specifying the conditional probability of the point process model, as it will uniquely specify the joint distribution of a sequence of events. First, to obtain a similar interpretation as the ETAS model [2], we define a similar parametric form for the conditional intensity function

$$\lambda^*(t, s) = \lambda_0 + \sum_{j: t_j < t} \nu(t, t_j, s, s_j), \quad (1)$$

where $\lambda_0 > 0$ is a constant background rate, ν is the kernel function that captures the influence of the past events \mathcal{H}_t . The form of the kernel function ν determines the profile of the spatio-temporal dependence of events.

We assume the kernel function takes a standard Gaussian diffusion kernel over space and decays exponentially over time. We adopt a mixture of generalized Gaussian diffusion kernels to enhance the spatial expressiveness, which is location dependent. Thus, it can capture a more complicated spatial-nonhomogeneous structure. Given all past events \mathcal{H}_t , we define

$$\nu(t, t', s, s') = \sum_{k=1}^K \phi_{s'}^{(k)} \cdot g(t, t', s, s' | \Sigma_{s'}^{(k)}, \mu_{s'}^{(k)}), \quad (2)$$

$$\forall t' < t, s \in \mathcal{S},$$

where K is a hyper-parameter that defines the number of components of the Gaussian mixture; $\mu_{s'}^{(k)} \in \mathbb{R}^2$ and $\Sigma_{s'}^{(k)} \in \mathbb{R}^{2 \times 2}$ are the mean and covariance matrix parameters for the k th Gaussian component at location s' ; $\phi_{s'}^{(k)} : \mathcal{S} \rightarrow \mathbb{R}$ is the corresponding weight of the component that satisfies $\sum_{k=1}^K \phi_{s'}^{(k)} = 1, \forall s' \in \mathcal{S}$. The exact forms of $\mu_{s'}^{(k)}, \Sigma_{s'}^{(k)}$, and

$\phi_{s'}^{(k)}$ will be specified later. In the following discussions, we focus on describing a single Gaussian component and omit the superscript k for the notational simplicity.

Now each Gaussian diffusion kernel is defined as

$$g(t, t', s, s' | \Sigma_{s'}, \mu_{s'}) = \frac{Ce^{-\beta(t-t')}}{2\pi\sqrt{|\Sigma_{s'}|}(t-t')} \cdot \exp \left\{ -\frac{(s-s' - \mu_{s'})^T \Sigma_{s'}^{-1} (s-s' - \mu_{s'})}{2(t-t')} \right\},$$

where $\beta > 0$ controls the temporal decay rate; $C > 0$ is a constant that decides the magnitude; $\mu_s = [\mu_x(s), \mu_y(s)]^T$ and Σ_s denote the mean and covariance parameters of the diffusion kernel; $|\cdot|$ denotes the determinant of a covariance matrix. Note that the structure of the kernel function $g(\cdot | \Sigma_s, \mu_s)$ may vary over "source" locations $s \in \mathcal{S}$. To be specific, Σ_s is defined as a positive semi-definite matrix

$$\Sigma_s = \begin{pmatrix} \sigma_x^2(s) & \rho(s)\sigma_x(s)\sigma_y(s) \\ \rho(s)\sigma_x(s)\sigma_y(s) & \sigma_y^2(s) \end{pmatrix}.$$

The parameters μ_s and Σ_s control the shape (shift, rotation, et cetera) of each Gaussian component. As shown in Fig. 3, parameters $\sigma_x(s), \sigma_y(s), \rho(s)$ may be different at different location s and jointly control the spatial structure of the diffusion. Parameters $\mu_x(s), \mu_y(s)$ define the offset of the center of the diffusion from the location s . Let $\mu_x : \mathcal{S} \rightarrow \mathbb{R}$, $\mu_y : \mathcal{S} \rightarrow \mathbb{R}$, $\sigma_x : \mathcal{S} \rightarrow \mathbb{R}^+$, $\sigma_y : \mathcal{S} \rightarrow \mathbb{R}^+$, and $\rho : \mathcal{S} \rightarrow (-1, 1)$ be non-linear mappings from location space \mathcal{S} to the corresponding parameter space. To capture intricate spatial dependence, we represent such non-linear mappings using neural networks.

3.2 Comparison with ETAS model

In the standard ETAS, the kernel function can be thought of as a special instance of the proposed heterogeneous Gaussian diffusion kernel in the NEST with *a single component whose parameters do not vary over space and time*: the kernel function defined in (2) can be simplified to $\nu(t, t', s, s') = g(t, t', s, s' | \Sigma, \mu)$, where the spatial and temporal parameters are location invariant, i.e., $\Sigma \equiv \text{diag}\{\sigma_x^2, \sigma_y^2\}$ and $\mu \equiv 0$. Compared with the standard Gaussian diffusion kernel used in ETAS, we introduce additional parameters ρ, μ_x, μ_y that allows the diffusion to shift, rotate, or stretch in the space. A comparison between the spatio-temporal kernels used in the ETAS and the NEST is presented in Fig. 2.

3.3 Deep neural network representation

Recall that parameters in each Gaussian component are determined by a set of non-linear mappings $\{\rho(s), \sigma_x(s), \sigma_y(s), \mu_x(s), \mu_y(s)\}$. We capture these non-linear spatial dependencies using a deep neural network through a latent embedding.

Assume the spatial structure at location s can be summarized by a latent embedding vector $\mathbf{h}(s) \in \mathbb{R}^d$, where d is the dimension of the embedding. The parameters of a Gaussian component at location s can be represented by an output layer of the neural network. The input of the layer is the

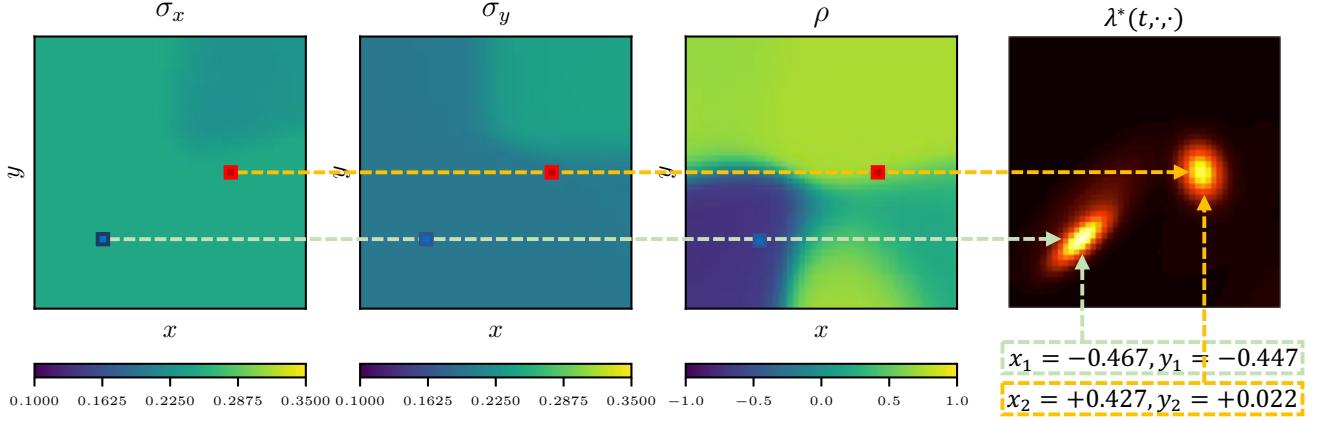


Fig. 3: An example of kernel used in the NEST model: σ_x, σ_y, ρ defines a Gaussian component in the heterogeneous Gaussian diffusion kernel. The rightmost subfigure shows the conditional intensity at time t , where two points occurred at location (x_1, y_1) and (x_2, y_2) have triggered the two diffusions (the bright spots) with different shapes. This Gaussian component is specified by parameters (mean, covariance) that vary continuously over space, which are themselves represented by neural networks.

latent embedding $\mathbf{h}(s)$. This layer is specified as follows. The mean parameters are specified by

$$\begin{aligned}\mu_x(s) &= C_x \cdot (\text{sigm}(\mathbf{h}(s)^T W_{\mu_x} + b_{\mu_x}) - 1/2), \\ \mu_y(s) &= C_y \cdot (\text{sigm}(\mathbf{h}(s)^T W_{\mu_y} + b_{\mu_y}) - 1/2),\end{aligned}$$

where C_x, C_y are preset constants that control the shift of the center of Gaussian components from location s , $\text{sigm}(x) = 1/(1 + e^{-x})$ is the sigmoid function which gives an output in the range $[0, 1]$. The variance and the correlation parameters are specified by

$$\begin{aligned}\sigma_x(s) &= \text{softplus}(\mathbf{h}(s)^T W_{\sigma_x} + b_{\sigma_x}), \\ \sigma_y(s) &= \text{softplus}(\mathbf{h}(s)^T W_{\sigma_y} + b_{\sigma_y}), \\ \rho(s) &= 2 \cdot \text{sigm}(\mathbf{h}(s)^T W_{\rho} + b_{\rho}) - 1,\end{aligned}$$

where $\text{softplus} = \log(1 + e^x)$ is a smooth approximation of the ReLU function. The parameters in the network $\theta_w = \{W_{\sigma_x}, W_{\sigma_y}, W_{\mu_x}, W_{\mu_y}, W_{\rho}\}$ and $\theta_b = \{b_{\sigma_x}, b_{\sigma_y}, b_{\mu_x}, b_{\mu_y}, b_{\rho}\}$ are weight-vectors and biases in the output layer of the Gaussian component. Note that we omitted the superscript k of each parameter in the discussion above for notational simplicity. However, it should be understood that each Gaussian component will have its own set of parameters. Finally, the weight of each component is given by $\phi_s^{(k)}$, which is defined through the soft-max function

$$\phi_s^{(k)} = e^{\mathbf{h}(s)^T W_{\phi}^{(k)}} / \sum_{\kappa=1}^K e^{\mathbf{h}(s)^T W_{\phi}^{(\kappa)}}.$$

where $W_{\phi}^{(k)} \in \mathbb{R}^d$ is a weight vector to be learned. The latent embedding $\mathbf{h}(s)$ is characterized by another neural network defined as $\mathbf{h}(s) = \psi(s|\theta_h)$, where $\psi(\cdot) : \mathbb{R}^2 \rightarrow \mathbb{R}^d$ is a fully-connected multi-layer neural network function taking spatial location s as input; θ_h contains the parameters in this neural network. In our experiments later, we typically use three-layer neural networks where each layer's width is 64.

In summary, the NEST with heterogeneous Gaussian mixture diffusion kernel is jointly parameterized by $\theta =$

$\{\beta, \theta_h, \{W_{\phi}^{(k)}, \theta_w^{(k)}, \theta_b^{(k)}\}_{k=1, \dots, K}\}$. The architecture is summarized in Fig. 4. In the following, we denote the conditional intensity as $\lambda_{\theta}^*(s, t)$ defined in (1), to make the dependence on the parameters more explicit. We emphasize that the Gaussian diffusion kernels' parameters vary *continuously* over location, and are represented by flexible neural networks; this is the key that our model can capture intricate spatial dependence in practice.

4 COMPUTATIONALLY EFFICIENT LEARNING

In this section, we define two approaches to learn parameters for the NEST model: (1) the maximum likelihood-based approach, and (2) the imitation learning-based approach, using a policy parameterized by the conditional intensity and a non-parametric reward function based on the maximum mean discrepancy (MMD) metric [13].

4.1 Maximum likelihood approach

The model parameters can be estimated via maximum likelihood estimate (MLE) since we have the explicit form of the conditional intensity function. Given a sequence of events $\mathbf{a} = \{a_0, a_1, \dots, a_n\}$ occurred on $(0, T] \times \mathcal{S}$ with length n , where $a_i = (t_i, s_i)$, the log-likelihood is given by

$$\ell(\theta) = \left(\sum_{i=1}^n \log \lambda_{\theta}^*(t_i, s_i) \right) - \int_0^T \int_{\mathcal{S}} \lambda_{\theta}^*(\tau, r) dr d\tau. \quad (3)$$

A crucial step to tackle the computational challenge is to evaluate the integral in (3). Here, we can obtain a closed-form expression for the likelihood function using the following proposition. This can reduce the integral to an analytical form, which can be evaluated directly without numerical integration (see the proof in the appendix).

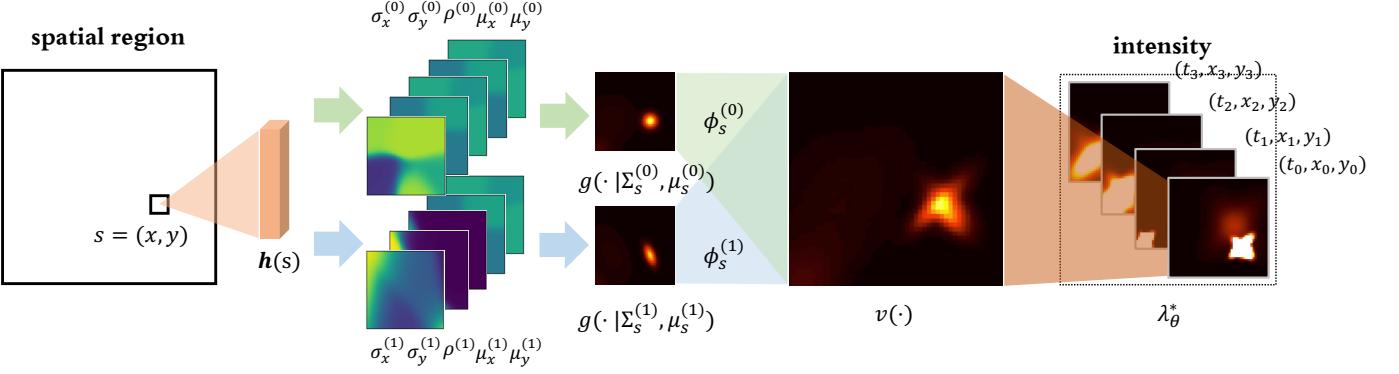


Fig. 4: An illustration for NEST’s neural network architecture based on a mixture of heterogeneous Gaussian diffusion kernel. Note that each Gaussian kernel is specified by neural networks, which summarizes the latent embedding information from data.

Proposition 1 (Integral of conditional intensity function). *Given ordered event times $0 = t_0 < t_1 < \dots < t_n < t_{n+1} = T$, for $i = 0, \dots, n$,*

$$\int_{t_i}^{t_{i+1}} \int_{\mathcal{S}} \lambda_{\theta}^*(\tau, r) dr d\tau = \lambda_0(t_{i+1} - t_i) |\mathcal{S}| + (1 - \epsilon) \frac{C}{\beta} \sum_{j: t_j < t_i} C_j \left(e^{-\beta(t_i - t_j)} - e^{-\beta(t_{i+1} - t_j)} \right),$$

where

$$C_j = \sum_{k=1}^K \phi_{s_j}^{(k)} \frac{\sigma_x^{(k)}(s_j) \sigma_y^{(k)}(s_j)}{|\Sigma_{s_j}^{(k)}|^{1/2}},$$

and the constant

$$\epsilon = \max_{j: t_j < t_{i+1}} \frac{\int_{t_i}^{t_{i+1}} \int_{\mathcal{S}} g(\tau, t_j, r, s_j) dr d\tau}{\int_{t_i}^{t_{i+1}} \int_{\mathbb{R}^2} g(\tau, t_j, r, s_j) dr d\tau}.$$

Since spatially, the kernel g is a Gaussian concentrated around s , and most events s_i locates in the relatively interior of \mathcal{S} when \mathcal{S} is chosen sufficiently large, we can ignore the marginal effect, and ϵ can become a number much smaller than 1. Due to the decreased activity in the region’s edges, the boundary effect is usually negligible [3]. Define $t_0 = 0$ and $t_{n+1} = T$. Since

$$\int_0^T \int_{\mathcal{S}} \lambda_{\theta}^*(\tau, r) dr d\tau = \sum_{i=0}^{n+1} \int_{t_i}^{t_{i+1}} \int_{\mathcal{S}} \lambda_{\theta}^*(\tau, r) dr d\tau,$$

using Proposition 1, we can write down the integral in the log-likelihood function in a closed-form expression.

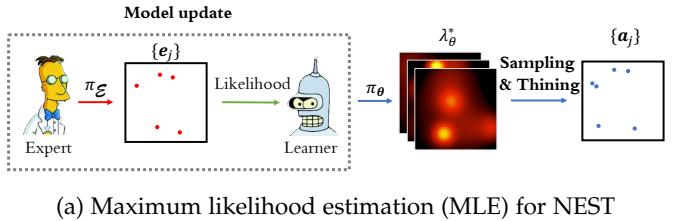
Finally, the optimal parameters can be thus obtained by $\hat{\theta} = \operatorname{argmax}_{\theta} \log \ell(\theta)$. Due to the non-convex nature of this problem, we solve the problem by stochastic gradient descent.

4.2 Imitation learning approach

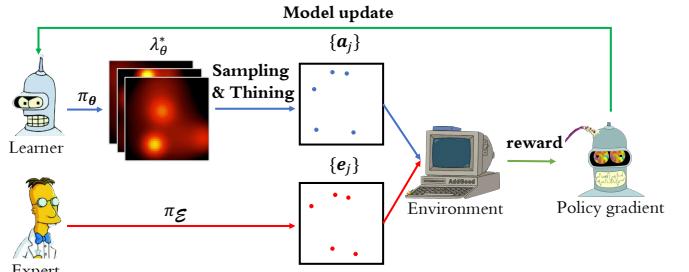
We now present a more flexible, imitation learning framework for model fitting. This approach’s primary benefit is that it does not rely on the pre-defined likelihood function; the reward function can be learned in a data-driven manner to optimally “imitate” the empirical distribution of the training data [28]. Hence, it is more robust to model mismatch. Moreover, we want to emphasize that, in our setting, the

learned reward function can be represented in a closed-form expression and conveniently estimated using samples. Thus, we can avoid the expensive inverse reinforce learning part for general imitation learning problems.

The imitation learning framework is described as follows. Assume a learner takes actions $a := (t, s) \in [0, T] \times \mathcal{S}$ sequentially in an environment according to a specific *policy*, and the environment gives feedbacks (using a reward function) to the learner via observing the discrepancy between the learner actions and the demonstrations (training data) provided by an expert. In our setting, both the learner’s actions and the demonstrations are over continuous-time and continuous space, which is a distinct feature of our problem.



(a) Maximum likelihood estimation (MLE) for NEST



(b) Imitation learning (IL) for NEST

Fig. 5: Comparison between the maximum likelihood and the imitation learning approaches. The main difference is that the MLE measure the “likelihood” of the data under a model. In contrast, our IL approach measures the actual divergence between the training data and the sequence generated from the model using MMD statistic without relying on model assumptions.

4.2.1 Policy parameterization

Our desired learner policy is a probability density function of possible actions given the history \mathcal{H}_t . We define such function as $\pi(t, s) : [0, T] \times \mathcal{S} \rightarrow [0, 1]$, which assigns a probability to the next event at any possible location and time. Let the last event time before T be t_n , and thus the next possible event is denoted as (t_{n+1}, s_{n+1}) . The definition of the policy is

$$\pi(t, s) = \mathbb{P}((t_{n+1}, s_{n+1}) \in [t, t + dt] \times B(s, \Delta s) | \mathcal{H}_t).$$

We will show that the policy can be explicitly related to the the conditional intensity function of the point process.

Lemma 1. *A spatial temporal point process which generate new samples according to $\pi(t, s)$ has the corresponding conditional intensity function*

$$\lambda_\theta^*(t, s) = \frac{\pi(t, s)}{1 - \int_0^t \int_{\mathcal{S}} \pi(\tau, r) d\tau dr}. \quad (4)$$

From Lemma 1, we can obtain the learner policy as the following proposition:

Proposition 2 (Learner policy related to conditional intensity). *Given the conditional intensity function $\lambda_\theta^*(t, s)$, the learner policy of a STPP on $[0, T] \times \mathcal{S}$ is given by*

$$\pi_\theta(t, s) = \lambda_\theta^*(t, s) \cdot \exp \left\{ - \int_{t_n}^t \int_{\mathcal{S}} \lambda_\theta^*(\tau, r) d\tau dr \right\}. \quad (5)$$

Thus, this naturally gives us a policy parameterization in a principled fashion: the learner policy π_θ is parameterized by θ based on the proposed heterogeneous Gaussian mixture diffusion kernel in Section 3. Note that using Proposition 1, which gives an explicit formula for the integral required in the exponent, the policy can be precisely evaluated even a deep neural network is included in the model.

4.2.2 Imitation learning objective

Now assume the training data is generated by an expert policy π_E , where the subscript E denotes “expert”. Given a reward function $r(\cdot)$, the goal is to find an optimal policy that maximizes the expected reward

$$\max_{\theta} J(\theta) := \mathbb{E}_{\mathbf{a} \sim \pi_\theta} \left[\sum_{i=1}^{n_a} r(a_i) \right],$$

where $\mathbf{a} = \{a_1, \dots, a_{n_a}\}$ is one sampled roll-out from policy π_θ . Note that n_a can be different for different roll-out samples.

4.2.3 Reward function

Consider the minimax formulation of imitation learning, which chooses the worst-case reward function that will give the maximum divergence between the rewards earned by the expert policy and the best learner policy:

$$\max_{r \in \mathcal{F}} \left(\mathbb{E}_{\epsilon \sim \pi_E} \left[\sum_{i=1}^{n_e} r(e_i) \right] - \max_{\pi_\theta \in \mathcal{G}} \mathbb{E}_{\mathbf{a} \sim \pi_\theta} \left[\sum_{i=1}^{n_a} r(a_i) \right] \right),$$

where \mathcal{G} is the family of all candidate policies π_θ and \mathcal{F} is the family class for reward function r in reproducing kernel Hilbert space (RKHS).

We adopt a data-driven approach to solve the optimization problem and find the worst-case reward. This

Algorithm 1: Efficient thinning algorithm for STPP

```

input  $\theta, \lambda_0, \beta, T, \mathcal{S}$ ;
output A set of events  $\alpha$  ordered by time;
Initialize  $\alpha = \emptyset, t = 0, s \sim \text{uniform}(\mathcal{S}), s_n = 0$ ;
while  $t < T$  do
     $u, D \sim \text{uniform}(0, 1); s \sim \text{uniform}(\mathcal{S})$ ;
     $\bar{\lambda} \leftarrow \lambda_0 + \sum_{(\tau, r) \in \alpha} \nu(t, \tau, s_n, r)$ ;
     $t \leftarrow t - \ln u / \bar{\lambda}$ ;
    Compute  $\lambda_\theta^*(t, s)$  from (4);
    if  $D \bar{\lambda} > \lambda_\theta^*(t, s)$  then
         $\alpha \leftarrow \alpha \cup \{(t, s)\}; s_n \leftarrow s$ ;
    end
end

```

is related to inverse reinforcement learning; we borrow the idea of MMD reward in [29, 13, 30, 7] and generalize it from a simple one-dimensional temporal point process to the more complex spatio-temporal setting. Suppose we are given training samples $\{e_j\}, j = 1, 2, \dots, M_E$, which are the demonstrations provided by the expert π_E , where each $e_j = \{e_0^{(j)}, e_1^{(j)}, \dots, e_{n_j}^{(j)}\}$ denotes a single demonstration. Also, we are given samples generated by the learner: let the trajectories generated by the learner π_θ denoted by $\{a_i\}, i = 1, 2, \dots, M_L$, where each trajectory $a_i = \{a_0^{(i)}, a_1^{(i)}, \dots, a_{n_i}^{(i)}\}$ denotes a single action trajectory. We will discuss how to generate samples in the following sub-section.

Using a similar argument as proving Theorem 1 in [7] and based on kernel embedding, we can obtain analytical expression for the worst-case reward function based on samples, which is given by

$$\hat{r}(a) \propto \frac{1}{M_E} \sum_{j=1}^{M_E} \sum_{u=1}^{n_j} k(e_u^{(j)}, a) - \frac{1}{M_L} \sum_{i=1}^{M_L} \sum_{v=1}^{n_i} k(a_v^{(i)}, a). \quad (6)$$

where $k(\cdot, \cdot)$ is a RKHS kernel. Here we use Gaussian kernel function, which achieves good experimental results on both synthetic data and real data.

Using samples generated from learner policy, the gradient of $J(\theta)$ with respect to θ can be computed by using policy gradient with variance reduction [7],

$$\nabla_\theta J(\theta) \approx \frac{1}{M_E} \sum_{j=1}^{M_E} \left[\sum_{i=1}^{n_j} (\nabla_\theta \log \pi_\theta(a_i) \cdot \hat{r}(a_i)) \right].$$

The gradient of policy $\nabla_\theta \log \pi_\theta(a_i)$ can be computed analytically in closed-form since it is specified by the conditional intensity in Proposition 2, and the architecture of the neural network fully specifies the conditional intensity as we discussed in Section 3.

4.2.4 Sampling from STPP using thinning algorithm

An important step in the imitation learning approach is to generate samples from our proposed model, i.e., $a \sim \pi_\theta$, given the history \mathcal{H}_t . Here, we develop an efficient sampling strategy to achieve good computational efficiency. We need to sample a point tuple $a = (t, s)$ according to the conditional intensity defined by (1). A default way to simulate point processes is to use the thinning algorithm [31, 32]. However,

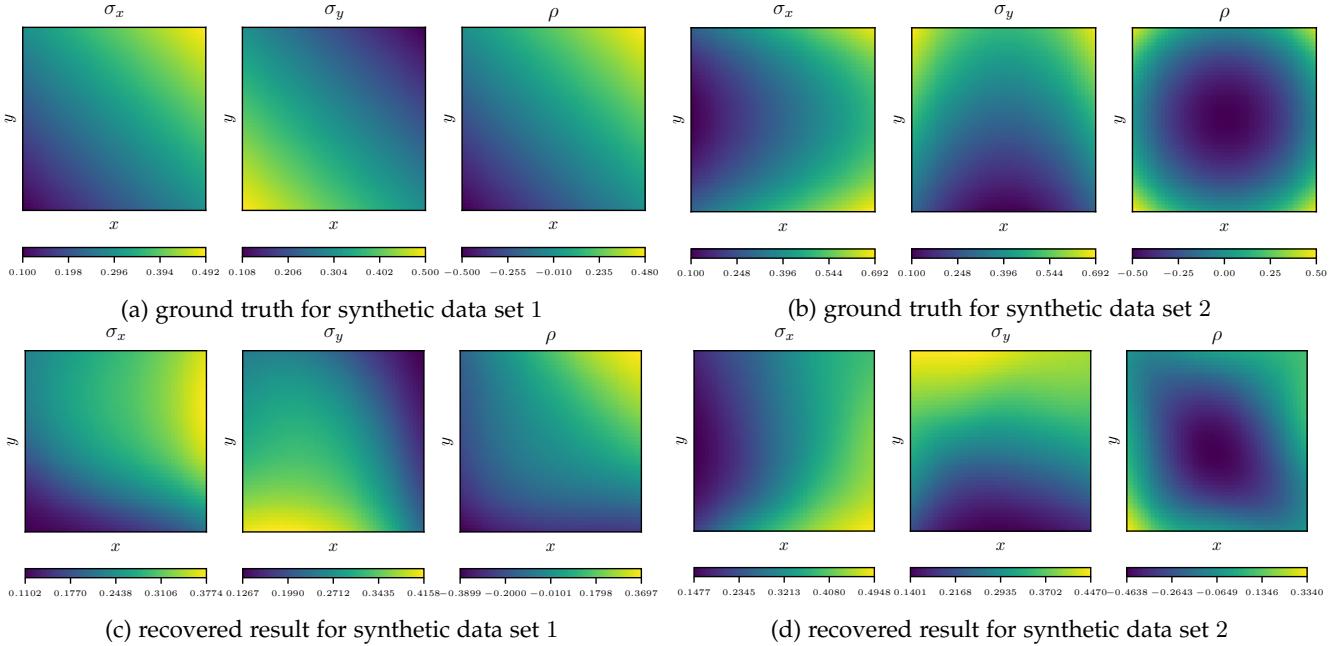


Fig. 6: Simulation results on two sets of synthetic data. (6a): The ground truth of the Gaussian parameter σ_x, σ_y, ρ in the synthetic data set 1; (6b): The ground truth of the Gaussian parameter σ_x, σ_y, ρ in the synthetic data set 2; (6c): The recovered Gaussian parameters using the synthetic data set 1; (6d): The recovered Gaussian parameters using the synthetic data set 2.

TABLE 1: Average MSE results on two synthetic data sets.

DATA SET	RANDOM	ETAS	NEST+IL	NEST+MLE	RLPP
SYNTHETIC 1 (SPACE-TIME)	.2781 (.2324, .3153)	.0433 (.0402, .0489)	.0075 (.0043, .0098)	.0134 (.0083, .0180)	N/A
SYNTHETIC 2 (SPACE-TIME)	.3327 (.2831, .3793)	.0512 (.0490, .0543)	.0124 (.0089, .0184)	.0321 (.0303, .0352)	N/A
SYNTHETIC 1 (TIME-ONLY)	.1734 (.1495, .1936)	.0135 (.0113, .0149)	.0021 (.0012, .0023)	.0048 (.0024, .0059)	.0146 (.0103, .0220)
SYNTHETIC 2 (TIME-ONLY)	.2147 (.1812, .2455)	.0323 (.0301, .0357)	.0036 (.0032, .0041)	.0055 (.0049, .0058)	.0341 (.0310, .0397)

TABLE 2: Log-likelihood on two real data sets.

DATA SET	ETAS	NEST $K = 1$	NEST $K = 5$	NEST $K = 10$	NEST $K = 20$
ROBBERY	22.34 (17.86, 25.12)	28.43 (26.13, 31.56)	34.27 (31.54, 36.53)	35.10 (32.48, 38.10)	35.32 (31.90, 39.34)
SEISMIC	138.1 (129.3, 145.3)	197.2 (190.3, 206.4)	224.1 (219.1, 229.1)	227.5 (223.3, 230.5)	226.9 (222.2, 231.4)

the vanilla thinning algorithm suffers from low sampling efficiency as it needs to sample in the space $|\mathcal{S}| \times [0, T]$ uniformly with the upper limit of the conditional intensity $\bar{\lambda}$ and only very few of candidate points will be retained in the end. In particular, given the parameter θ , the procedure's computing complexity increases exponentially with the size of the sampling space. To improve sampling efficiency, we propose an efficient thinning algorithm summarized in Algorithm 1. The "proposal" density is a non-homogeneous STPP, whose intensity function is defined from the previous iterations. This analogous to the idea of rejection sampling [12].

5 EXPERIMENTS

In this section, we evaluate our approaches by conducting experiments on both synthetic and real data sets.

5.1 Baselines and evaluation metrics

This section compares our proposed neural embedding spatio-temporal (NEST) with a benchmark and several state-of-the-art methods in the field. These include (1) Random uniform that randomly makes actions in the action space; (2) epidemic type aftershock-sequences (ETAS) with standard diffusion kernel, which is currently the most widely used approach in spatio-temporal event data modeling. For ETAS, the parameters are estimated by maximum likelihood estimate (MLE); (3) reinforcement learning point processes model (RLPP) [7] is for modeling temporal point process only, which cannot be easily generalized to spatio-temporal models. (4) NEST+IL is our approach using imitation learning; (5) NEST+MLE is our approach where the parameters are estimated by MLE. We also investigate our method's performance with a different number of Gaussian components (K) under the same experimental settings for the real data.

To evaluate the performance of algorithms (i.e., various generative models), we adopt two performance metrics: (1) the average mean square error (MSE) of the *one-step-ahead prediction*. The *one-step-ahead prediction* for the next event (t_{n+1}, s_{n+1}) is carried out by computing the expectation of the conditional probability function (or policy) defined in (5) given the past observation $\mathcal{H}_{t_{n+1}}$, i.e.,

$$\left[\hat{t}_{n+1} \right] = \left[\int_{t_n}^T \tau \int_{\mathcal{S}} \pi_{\theta}(\tau, \omega) d\omega d\tau \right].$$

Due to the close-form expression of π_{θ} in (5), the integration above can also be obtained analytically. (2) the maximum mean discrepancy (MMD) metric between the real observed sequences and the generated sequences from the models, as specified in (6). To obtain and compare the confidence interval of the performance metrics, we repeat each experiment 20 times. For synthetic data, we also directly compare

TABLE 3: Average absolute MMD per sequence on two real data sets.

DATA SET	RANDOM	ETAS	NEST+IL $K = 5$	NEST+MLE $K = 1$	NEST+MLE $K = 5$	NEST+MLE $K = 10$	NEST+MLE $K = 20$	RLPP
ROBBERY (SPACE-TIME)	108.0 (92.2, 125.4)	72.9 (66.1, 79.4)	68.4 21.1	71.6 (64.2, 79.5)	69.4 (64.7, 73.5)	69.1 (63.5, 74.1)	68.9 (61.1, 74.6)	N/A
SEISMIC (SPACE-TIME)	53.1 (40.0, 67.4)	32.2 (29.7, 36.4)	21.1 (19.5, 24.0)	30.9 (27.7, 32.8)	28.6 (24.3, 33.1)	27.9 (22.0, 34.3)	28.3 (23.8, 35.7)	N/A
ROBBERY (TIME ONLY)	126.8 (109.3, 150.1)	90.6 (85.2, 97.3)	82.2 21.2	88.1 (80.4, 95.1)	83.8 (77.3, 87.9)	83.1 (76.0, 86.4)	83.0 (76.2, 84.5)	83.1 (75.1, 89.2)
SEISMIC (TIME ONLY)	60.7 (55.1, 66.9)	51.0 (48.5, 53.1)	21.2 (18.8, 24.1)	35.3 (29.7, 39.2)	27.3 (24.6, 30.4)	25.1 (22.0, 28.9)	24.6 (21.6, 27.5)	23.3 (20.3, 28.2)

TABLE 4: Average MSE results on two real data sets.

DATA SET	RANDOM	ETAS	NEST+IL $K = 5$	NEST+MLE $K = 1$	NEST+MLE $K = 5$	NEST+MLE $K = 10$	NEST+MLE $K = 20$	RLPP
ROBBERY (SPACE-TIME)	.6323 (.6143, .6583)	.1425 (.1368, .1497)	.0503 .0221	.1144 (.0445, .0549)	.0649 (.1078, .1321)	.0610 (.0574, .0693)	.0601 (.0544, .0679)	.0601 (.0531, .0672)
SEISMIC (SPACE-TIME)	.2645 (.2457, .2833)	.0221 (.0180, .0277)	.0119 .0134	.0203 (.0173, .0252)	.0153 (.0133, .0181)	.0142 (.0132, .0174)	.0143 (.0128, .0182)	N/A
ROBBERY (TIME ONLY)	.4783 (.4553, .4940)	.0857 (.0788, .0914)	.0104 (.0082, .0141)	.0583 (.0492, .0651)	.0094 (.0064, .0134)	.0089 (.0052, .0123)	.0091 (.0056, .0143)	.0183 (.0126, .0242)
SEISMIC (TIME ONLY)	.1266 (.1001, .1562)	.0173 (.0134, .0210)	.0045 .0036	.0175 (.0142, .0211)	.0150 (.0131, .0198)	.0143 (.0124, .0177)	.0144 (.0122, .0183)	.0122 (.0112, .0154)

the recovered parameters against the true parameters used to construct the generator and generate the corresponding synthetic data.

5.2 Synthetic data

To validate the capability of the NEST in capturing the spatial information from discrete event data, we first evaluate our method on two synthetic data sets. As an ablation study, these two data sets are generated by two NEST with a single Gaussian component ($K = 1$), where their artificial parameters (ground truth) have been shown in Fig. 6a and 6b, respectively. The parameters in synthetic data set 1 (Fig. 6a) are linearly related to their spatial locations and parameters in synthetic data set 2 (Fig. 6b) are non-linearly related to their spatial locations. In both data sets, there are 5,000 sequences with an average length of 191, and 80%, 20% sequences are used for training and testing. The time and location of events have been normalized to $T = 10$ and $\mathcal{S} = [-1, +1] \times [-1, +1]$. We specify three hidden layers with 64 nodes per layer for the neural network defined by θ_h . The optimizer randomly takes 40 sequences as a batch during the training. Both models are trained by both MLE and IL approaches.

We fit our NEST models using these two synthetic data sets separately and obtain the corresponding parameters. Fig. 6c and Fig. 6d show the recovered parameters learned from two synthetic data sets. It shows that the spatial distribution of the recovered parameters resembles the true parameters. It also confirms that our model can capture the underlying spatial information from both linear and non-linear parameter space. In Table 1, we report the MSE of the one-step-ahead prediction and the 25% and 75% quantile of the performance metric repeated over 20 experiments (in brackets). Clearly, our models (NEST+ML and NEST+IL) consistently outperform other methods on two synthetic data sets.

5.3 Real data

Now we test our approaches on two real-world data sets: Atlanta 911 calls-for-services data (provide by the Atlanta

Police Department to us under a data agreement) and Northern California seismic data [33]. For the ease of comparison, we normalize the space region of both data sets to the same space $T \times \mathcal{S}$ where $T = (0, 10]$ and $\mathcal{S} = [-1, 1] \times [-1, 1]$. The detailed description of the two data sets is as follows.

Atlanta 911 calls-for-service data. The Atlanta Police Department provides the 911 calls-for-service data in Atlanta from the end of 2015 to 2017 (this data set has been previously used to validate the crime linkage detection algorithm [34, 35, 24]). We extract 7,831 reported robbery from the data set since robbers usually follow particular *modus operandi* (M.O.), where criminal spots and times tend to have a causal relationship with each other. Each robbery report is associated with the time (accurate to the second) and the geolocation (latitude and longitude) indicating when and where the robbery occurred. We consider each series of robbery as a sequence.

Northern California seismic data. The Northern California Earthquake Data Center (NCEDC) provides public time series data [33] that comes from broadband, short period, strong motion seismic sensors, GPS, and other geophysical sensors. We extract 16,401 seismic records with a magnitude larger than 3.0 from 1978 to 2018 in Northern California and partition the data into multiple sequences every quarter. To test our model, we only utilize time and geolocation in the record.

We first quantitatively compare our NEST+MLE and ETAS by evaluating the obtained log-likelihood per sequence for both real-world data sets. As shown in Table 2, our model attains a larger log-likelihood value on both data sets comparing to the state-of-the-art (ETAS), which confirms that our model can obtain a better fit on these two real data sets. We also compare NEST+MLE and NEST+IL with other baseline methods by measuring their average absolute MMD values and average MSE of the one-step-ahead prediction. The average absolute MMD value can be measured by computing the pairwise distances between the observed training data and the generated sequences according to (6). The absolute MMD value indicates the level of sim-

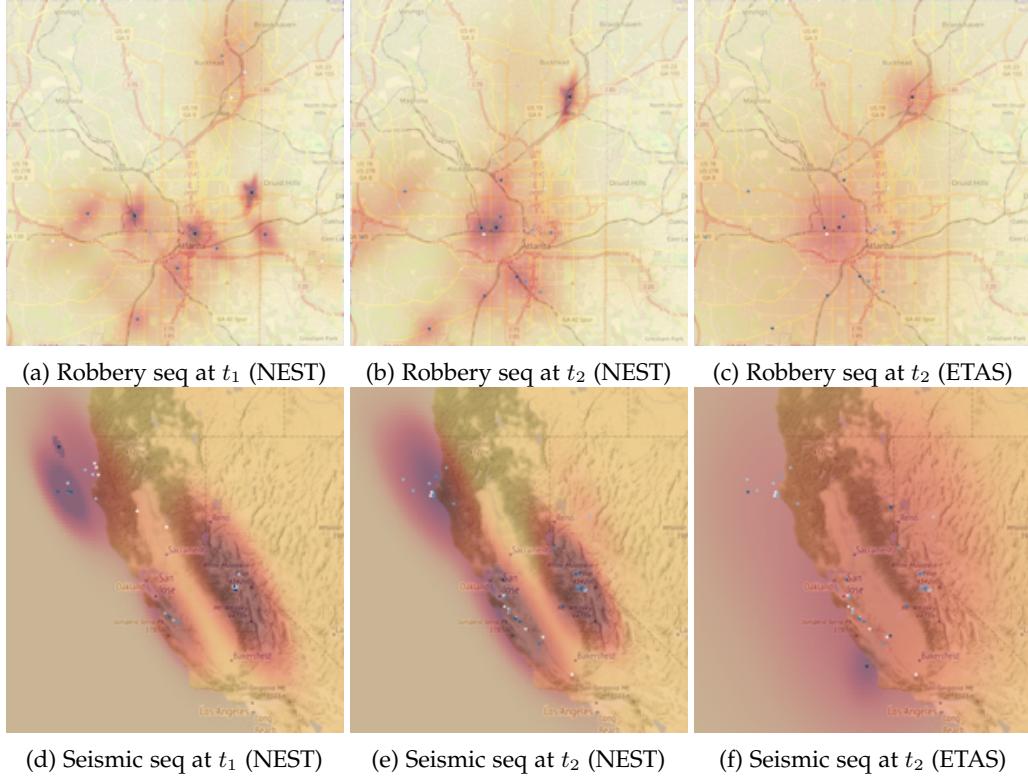


Fig. 7: Snapshots of the conditional intensity for two real data sequences (crime events in Atlanta and seismic events): (7a, 7b, 7c): Snapshots of the conditional intensity for a series of robberies in Atlanta. (7d, 7e, 7f): Snapshots of the conditional intensity for a series of earthquakes in Northern California. (7b, 7a, 7e, 7d) are generated by NEST+IL ($K = 5$) and (7f, 7c) are generated by ETAS, respectively. The color depth indicates the value of intensity. The region in darker red has a higher risk to have the next event happened again. The dots on the maps represent the occurred events in which dark blue represents the newly happened events, and white represents events that happened initially.

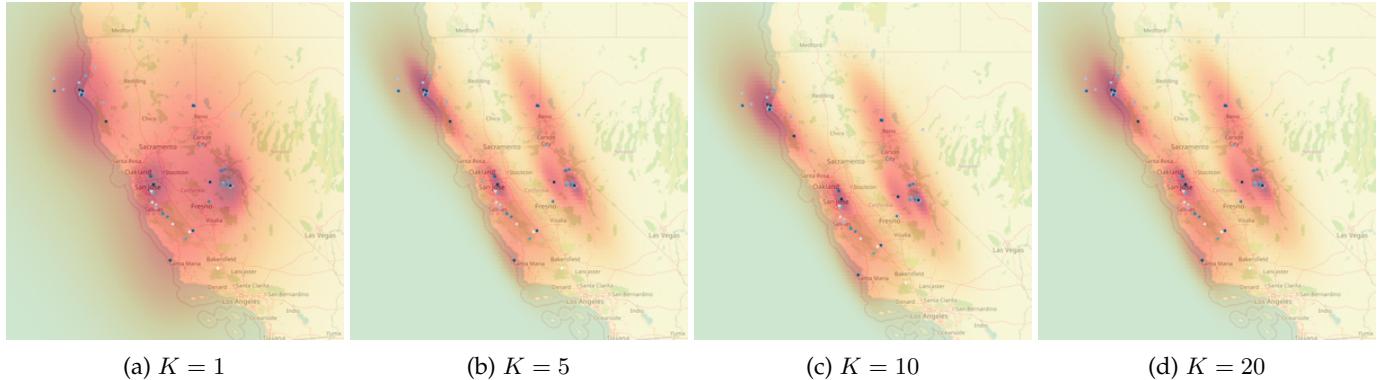


Fig. 8: Illustration of the conditional intensity function for seismic events in Northern California over different the number of components of the Gaussian mixture K , while other experimental settings are the same. As K increases from 1 to 5, the model can capture more details in conditional spatial intensities; also, when $K \geq 5$, there seems to be little difference.

ilarity of two arbitrary distributions. If two distributions are the same, then their absolute MMD value is zero. In the experiment, we randomly pick 100 pairs of generated sequences against observed sequences from the real data sets and compute their average absolute MMD value. We summarized the results of the average absolute MDD and average MSE and the 25% and 75% quantile of performance metrics repeated over 20 experiments (in brackets) in Table 3 and Table 4, respectively. The results show that both our

methods NEST+IL and NEST+MLE outperform the state-of-the-art (ETAS) regarding both metrics. We also show that our method has a competitive performance even without considering spatial information comparing to RLPP. From Fig. 8 and Table 2, 3, 4, we find that increasing the number of Gaussian component can improve the model performance. However, when $K \geq 5$, there is little difference (this may vary case by case).

To interpret the spatial dependence learned using our

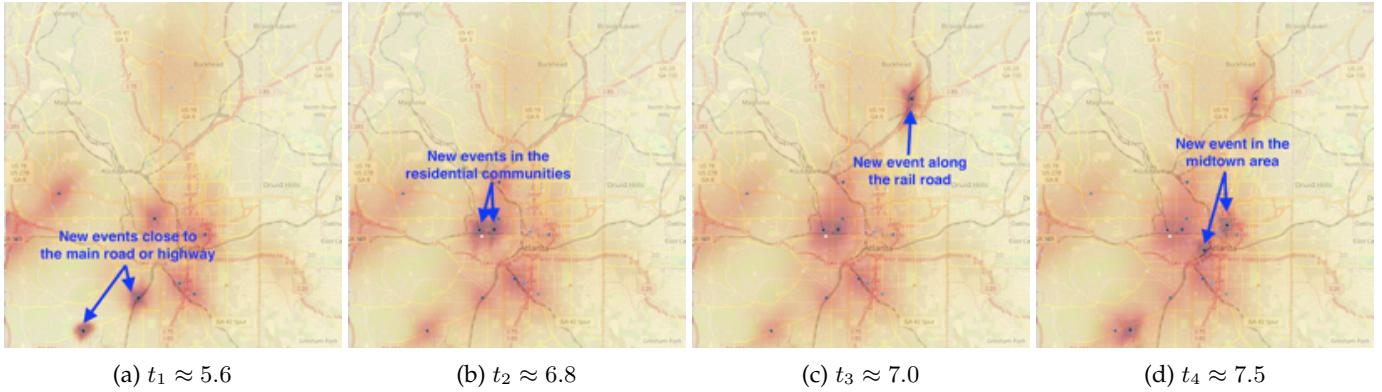


Fig. 9: Snapshots of the conditional intensity of a robbery sequence at four time points. The color depth indicates the value of intensity. The region in darker red has a higher risk to have robbery event happened again. It also shows that robbery events in different regions have very distinctive diffusion patterns, which may be related to the city's geographic and demographic features.

TABLE 5: Training time per batch on two real data sets.

DATA SET	ETAS	NEST+IL $K = 5$	NEST+MLE $K = 1$	NEST+MLE $K = 5$	NEST+MLE $K = 10$	NEST+MLE $K = 20$	RLPP
ROBBERY (SPACE-TIME)	12s	312s	34s	43s	60s	94s	N / A
SEISMIC (SPACE-TIME)	22s	457s	42s	69s	91s	110s	N / A
ROBBERY (TIME ONLY)	3s	142s	9s	16s	25s	34s	160s
SEISMIC (TIME ONLY)	3s	194s	12s	21s	34s	47s	216s

model, we visualize the spatial conditional intensity over space at a specific time frame. For the state-of-the-art, as shown in Fig. 7c, 7f, we can see that ETAS captured the general pattern of the conditional intensity over space, where regions with more events tend to have higher intensity. Comparing to the results shown in Fig. 7a, 7b, 7d, 7e, our NEST is able to capture more intricate spatial pattern at different locations and the shape of the captured diffusion also varies depending locations. The results clearly demonstrates that the Pacific-North America plate motion is along two parallel fault systems – the San Andreas fault, and Eastern California shear zone and Walker Lane.

For 911 calls-for-service data, shown in Fig. 7a, 7b, the spatial influence of some robbery events diffuse to the surrounding streets and the community blocks unevenly. For seismic data, shown in Fig. 7e, 7d, the spatial influence of some events majorly diffuses along the earthquake fault lines. We also visualize the fitted parameters of each Gaussian component of the NEST models on both data sets, as shown in Fig. ?? and Fig. ?? in the appendix.

Besides, Fig. 9 shows a series of snapshots of the conditional intensity for a sequence of robberies at different time points. Each subfigure shows that events in different urban regions have various shapes of spatial influence on their neighborhoods. For example, events near the highway or railroad tend to spread their influence along the road. In contrast, events around the Midtown area tend to spread their influence more evenly to the neighboring communities.

5.4 Learning efficiency

To compare the efficiencies between different learning strategies, we perform the experiments on the same data sets

given a single laptop's computational power with quad-core processors which speed up to 4.7 GHz, and record the training time per batch accordingly. In Table 5, we show that the training process of NEST+MLE is faster in general than NEST+IL; in the imitation learning framework, events generation is time-consuming (even using our proposed efficient thinning algorithm) compared to the learning process. Also, considering more Gaussian components in the NEST model can also increase the training time significantly. Comparing to other baseline methods [10, 7], including RLPP, which are usually over-parameterized on simple problems, our NEST can achieve better results with less training time using a fewer number of parameters since the deep neural network in NEST is only used to capture the non-linear spatial pattern in the embedding space.

6 CONCLUSION AND DISCUSSIONS

We have presented a generative model for spatio-temporal event data by designing heterogeneous and flexible spatio-temporal point process models whose parameters are parameterized by neural networks, representing complex-shaped spatial dependence. Unlike a full neural network approach, our parameterization based on Gaussian diffusion kernels can still enjoy interpretability: the shape of the spatial-temporal influence kernels, which are useful for deriving knowledge for domain experts such as seismic data analysis. Also, we present a flexible model-fitting approach based on imitation learning, a form of reinforcement learning, which does not rely on exact model specification since it is based on directly measuring the divergence between the empirical distributions of the training data and the data generated from data (in contrast to MLE, which measures the "likelihood" of data under a pre-specified model structure). The robustness of the imitation learning-based training approach for the NEST model is demonstrated using both synthetic and real data sets, in terms of the predictive capability (measured by MSE of the conditional intensity function, which measure the probability for an event to occur at a location in the future given the past), and the maximum-mean-divergence (measures how realistic the data generated from the model matches to the real training data). We have observed from

synthetic data and real-data that the imitation learning approach will achieve a significant performance gain than maximum likelihood and more robust for training NEST model. Our NEST can have a wide range of applications: it is a probabilistic generative model and can be used to evaluate the likelihood of a sequence and perform anomaly detection spatio-temporal event data, as well as making predictions about future incidents given the history information.

Future work may include extending to marked spatio-temporal processes, for instance, to consider the magnitude of the earthquake in the models in the current set-up. Moreover, it will be interesting to consider how to incorporate *known* spatial structural information in the spatio-temporal model. For instance, in the most recent earthquake model in California, geophysicists have already considered both the ETAS and the fault model/distribution [36].

ACKNOWLEDGEMENT

The work is partially funded by an NSF CAREER Award CCF-1650913, NSF CMMI-2015787, DMS-1938106, DMS-1830210.

REFERENCES

- [1] A. Reinhart, "A review of self-exciting spatio-temporal point processes and their applications," *Statist. Sci.*, vol. 33, no. 3, pp. 299–318, 08 2018.
- [2] Y. Ogata, "Statistical models for earthquake occurrences and residual analysis for point processes," *Journal of the American Statistical Association*, vol. 83, no. 401, pp. 9–27, 1988.
- [3] ——, "Space-time point-process models for earthquake occurrences," *Annals of the Institute of Statistical Mathematics*, vol. 50, no. 2, pp. 379–402, 1998.
- [4] S. Stein and M. Liu, "Long aftershock sequences within continents and implications for earthquake hazard assessment," *Nature*, vol. 462, no. 7269, pp. 87–89, 2009.
- [5] N. Du, H. Dai, R. Trivedi, U. Upadhyay, M. Gomez-Rodriguez, and L. 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*, ser. KDD '16. New York, NY, USA: ACM, 2016, pp. 1555–1564.
- [6] H. Mei and J. Eisner, "The neural hawkes process: A neurally self-modulating multivariate point process," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. USA: Curran Associates Inc., 2017, pp. 6757–6767.
- [7] S. Li, S. Xiao, S. Zhu, N. Du, Y. Xie, and L. Song, "Learning temporal point processes via reinforcement learning," in *Advances in Neural Information Processing Systems 31*. Curran Associates, Inc., 2018, pp. 10781–10791.
- [8] U. Upadhyay, A. De, and M. Gomez-Rodriguez, "Deep reinforcement learning of marked temporal point processes," in *Advances in Neural Information Processing Systems 31*, 2018.
- [9] S. Xiao, J. Yan, X. Yang, H. Zha, and S. M. Chu, "Modeling the intensity function of point process via recurrent neural networks," in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, ser. AAAI17. AAAI Press, 2017, p. 15971603.
- [10] S. Xiao, M. Farajtabar, X. Ye, J. Yan, L. Song, and H. Zha, "Wasserstein learning of deep generative point process models," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Curran Associates Inc., 2017, pp. 3250–3259.
- [11] S. Zhu, H. S. Yuchi, and Y. Xie, "Adversarial anomaly detection for marked spatio-temporal streaming data," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8921–8925.
- [12] Y. Ogata, "On lewis' simulation method for point processes," *IEEE Transactions on Information Theory*, vol. 27, no. 1, pp. 23–31, January 1981.
- [13] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. J. Smola, "A kernel method for the two-sample-problem," in *Advances in Neural Information Processing Systems 19*. MIT Press, 2007, pp. 513–520.
- [14] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. Cambridge: MIT press, 1998.
- [15] A. Y. Ng and S. Russell, "Algorithms for inverse reinforcement learning," in *in Proc. 17th International Conf. on Machine Learning*. Morgan Kaufmann, 2000, pp. 663–670.
- [16] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Advances in Neural Information Processing Systems 29*. D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Curran Associates, Inc., 2016, pp. 4565–4573.
- [17] Y. Duan, X. Chen, R. Houthooft, J. Schulman, and P. Abbeel, "Benchmarking deep reinforcement learning for continuous control," in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ser. ICML'16. JMLR.org, 2016, pp. 1329–1338.
- [18] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *Proceedings of The 33rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. F. Balcan and K. Q. Weinberger, Eds., vol. 48. New York, New York, USA: PMLR, 20–22 Jun 2016, pp. 1928–1937.
- [19] M. Short, A. Bertozzi, and P. Brantingham, "Nonlinear patterns in urban crime: Hotspots, bifurcations, and suppression," *SIAM Journal on Applied Dynamical Systems*, vol. 9, no. 2, pp. 462–483, 2010.
- [20] M. B. Short, P. J. Brantingham, A. L. Bertozzi, and G. E. Tita, "Dissipation and displacement of hotspots in reaction-diffusion models of crime," *Proceedings of the National Academy of Sciences*, vol. 107, no. 9, pp. 3961–3965, 2010.
- [21] E. W. Fox, M. B. Short, F. P. Schoenberg, K. D. Coronges, and A. L. Bertozzi, "Modeling e-mail networks and inferring leadership using self-exciting point processes," *Journal of the American Statistical Association*, vol. 111, no. 514, pp. 564–584, 2016.
- [22] J. R. Zipkin, F. P. Schoenberg, K. Coronges, and A. L. Bertozzi, "Point-process models of social network interactions: Parameter estimation and missing data recovery," *European Journal of Applied Mathematics*, vol. 27, no. 3, p. 502529, 2016.
- [23] E. Lewis, G. Mohler, P. J. Brantingham, and A. L. Bertozzi, "Self-exciting point process models of civilian deaths in iraq," *Security Journal*, vol. 25, no. 3, pp. 244–264, Jul 2012.
- [24] S. Zhu and Y. Xie, "Spatial-temporal-textual point processes with applications in crime linkage detection," 2019.
- [25] G. O. Mohler, M. B. Short, P. J. Brantingham, F. P. Schoenberg, and G. E. Tita, "Self-exciting point process modeling of crime," *Journal of the American Statistical Association*, vol. 106, no. 493, pp. 100–108, 2011.
- [26] A. G. HAWKES, "Spectra of some self-exciting and mutually exciting point processes," *Biometrika*, vol. 58, no. 1, pp. 83–90, 04 1971. [Online]. Available: <https://doi.org/10.1093/biomet/58.1.83>
- [27] F. Musmeci and D. Vere-Jones, "A space-time clustering model for historical earthquakes," *Annals of the Institute of Statistical Mathematics*, vol. 44, no. 1, pp. 1–11, Mar 1992.
- [28] J. Langford and B. Zadrozny, "Reducing t-step reinforcement learning to classification," 2003.
- [29] B. Kim and J. Pineau, "Maximum mean discrepancy imitation learning," in *Robotics: Science and Systems*, 2013.
- [30] G. K. Dziugaite, D. M. Roy, and Z. Ghahramani, "Training generative neural networks via maximum mean discrepancy optimization," in *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, ser. UAI'15. Arlington, Virginia, United States: AUAI Press, 2015, pp. 258–267.
- [31] E. Gabriel, B. Rowlingson, and P. Diggle, "stpp: An r package for plotting, simulating and analyzing spatio-temporal point patterns," *Journal of Statistical Software*, vol. 53, pp. 1–29, 04 2013.
- [32] D. J. Daley and D. Vere-Jones, *An introduction to the theory of point processes. Vol. II*, 2nd ed., ser. Probability and its Applications (New York). New York: Springer, 2008, general theory and structure.
- [33] N. C. E. D. C. U. B. S. Laboratory, "NCEDC," 2014.
- [34] S. Zhu and Y. Xie, "Crime event embedding with unsupervised feature selection," 2018.
- [35] ——, "Crime incidents embedding using restricted boltzmann machines," *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2376–2380, 2018.

[36] E. H. Field, T. H. Jordan, M. T. Page, K. R. Milner, B. E. Shaw, T. E. Dawson, G. P. Biasi, T. Parsons, J. L. Hardebeck, A. J. Michael, I. Weldon, Ray J., P. M. Powers, K. M. Johnson, Y. Zeng, K. R. Felzer, N. v. d. Elst, C. Madden, R. Arrowsmith, M. J. Werner, and W. R. Thatcher, "A Synoptic View of the Third Uniform California Earthquake Rupture Forecast (UCERF3)," *Seismological Research Letters*, vol. 88, no. 5, pp. 1259–1267, 07 2017.



Shixiang Zhu is currently pursuing the Ph.D. degree in Machine Learning in H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology. He received the B.S. and M.S. degree in computer science from Beijing University of Posts and Telecommunications in 2014 and 2017, respectively, and interned at Argonne National Laboratory in 2020. He is a recipient of the Doing Good with Good OR Student Paper Competition Award. His research interests include Machine Learning in spatio-temporal analysis, optimization in urban planning, and related applications.



Yao Xie is an associate professor and Harold R. and Mary Anne Nash Early Career Professor at Georgia Institute of Technology in the H. Milton Stewart School of Industrial and Systems Engineering, and an Associate Director of the Machine Learning Center. She received her Ph.D. in Electrical Engineering (minor in Mathematics) from Stanford University, M.Sc. in Electrical and Computer Engineering from the University of Florida, and B.Sc. in Electrical Engineering and Computer Science from University of Science and Technology of China (USTC). She was a Research Scientist at Duke University. Her research areas are statistics (in particular sequential analysis and sequential change-point detection), machine learning, and signal processing, in providing the theoretical foundation as well as developing computationally efficient and statistically powerful algorithms. She has worked on such problems in sensor networks, social networks, power systems, crime data analysis, and wireless communications. She received the National Science Foundation (NSF) CAREER Award in 2017. She is currently an Associate Editor for IEEE Transactions on Signal Processing, and Sequential Analysis: Design Methods and Applications.



Shuang Li is currently a postdoctoral researcher at Harvard University. She received her Ph.D. in Industrial Engineering from Georgia Institute of Technology in 2019, M.Sc. in Electrical Engineering from University of Massachusetts, Amherst, and B.Sc. in Automation from University of Science and Technology of China. She is the finalist for the 2018 INFORMS QSR best student paper competition, the finalist for the 2018 INFORMS Social Media Analytics Best Student Paper Competition, and the winner of Alice and John Jarvis Ph.D. Student Research Award, Honorable Mention, in 2016.



Zhigang Peng is a professor of geophysics in the School of Earth and Atmospheric Sciences at Georgia Institute of Technology. His research mainly focuses on active fault zone structures and earthquake physics. Over the past 13 years, he has received a total of more than \$2 million research funding from National Science Foundation (NSF), United States Geologic Survey (USGS), National Aeronautics and Space Administration (NASA), and Southern California Earthquake Center (SCEC). Dr. Peng has published more than 100 papers in peer-reviewed journals such as Nature Geoscience, Science, Journal of Geophysical Research (JGR), Geophysical Research Letters, etc. His H-index is 36 (from Google Scholar), and his research results have been widely reported by major news and scientific media around the world. He was an associate editor for both JGR and Bulletin of Seismological Society of America (BSSA) from 2011 to 2013, and the editor-in-chief for Seismological Research Letters from 2013 to 2019. He is currently serving as the Board of Directors for the Incorporated Research Institution for Seismology (IRIS), and the president of the Eastern Section of the Seismological Society of America. Dr. Peng is a recipient of the NSF CAREER award in 2010, the Charles Richter Award from the Seismological Society of America (SSA), and the Young Faculty award from the School of Earth and Atmospheric Sciences.