

UNIPoint: Universally Approximating Point Processes Intensities

Alexander Soen, Alexander Mathews, Daniel Grixti-Cheng, Lexing Xie

The Australian National University

alexander.soen@anu.edu.au, alex.mathews@anu.edu.au, a500846@anu.edu.au, lexing.xie@anu.edu.au

Abstract

Point processes are a useful mathematical tool for describing events over time, and so there are many recent approaches for representing and learning them. One notable open question is how to precisely describe the flexibility of point process models and whether there exists a general model that can represent *all* point processes. Our work bridges this gap. Focusing on the widely used event intensity function representation of point processes, we provide a proof that a class of learnable functions can universally approximate any valid intensity function. The proof connects the well known Stone-Weierstrass Theorem for function approximation, the uniform density of non-negative continuous functions using a transfer functions, the formulation of the parameters of a piece-wise continuous functions as a dynamic system, and a recurrent neural network implementation for capturing the dynamics. Using these insights, we design and implement UNIPoint, a novel neural point process model, using recurrent neural networks to parameterise sums of basis function upon each event. Evaluations on synthetic and real world datasets show that this simpler representation performs better than Hawkes process variants and more complex neural network-based approaches. We expect this result will provide a practical basis for selecting and tuning models, as well as furthering theoretical work on representational complexity and learnability.

1 Introduction

Temporal point processes (Daley and Vere-Jones 2007) are a preferred tool for describing events happening in irregular intervals, such as, earthquake modelling (Ogata 1988), social media (Zhao et al. 2015), and finance (Embrechts, Liniger, and Lin 2011). One common variant is the self-exciting Hawkes process with parametric kernel (Laub, Taimre, and Pollett 2015), which describes prior events triggering future events. However, misspecification of the kernel will likely result in poor performance (Mishra, Rizoiu, and Xie 2016). One may ask what are the most flexible classes of point process intensity functions? How can they be implemented computationally? Does a flexible representation lead to good performance?

There is a body of literature surrounding these three questions. Multi-layer neural networks are well known for being flexible function approximators. They are able to ap-

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

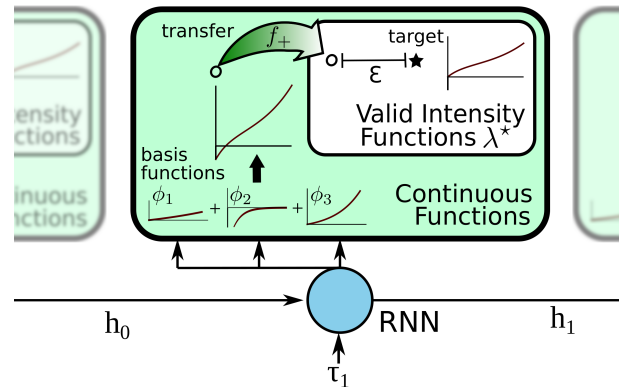


Figure 1: Overview of our method of universally approximating point processes. A RNN is used to parameterise a set of basis functions for each interarrival time τ_i . Then, the sum of basis functions is used to approximate a continuous function, which is composed with a transfer function f_+ to universally approximate all valid intensity functions.

proximate any Borel-measurable function on a compact domain (Cybenko 1989; Hornik, Stinchcombe, and White 1989). A number of neural architectures have been proposed for point processes. The Recurrent Marked Temporal Point Process model (RMTTPP) (Du et al. 2016) uses Recurrent Neural Networks (RNN) to encode event history, and defines the conditional intensity function by a parametric form. Common choices of such parametric forms include an exponential function (Du et al. 2016; Upadhyay, De, and Rodriguez 2018) or a constant function (Li et al. 2018; Huang, Wang, and Mak 2019). Variants of the RNN have been explored, including NeuralHawkes (Mei and Eisner 2017) that makes the RNN state a functions over time; as well as Transformer Hawkes (Zuo et al. 2020) and Self-attention Hawkes (Zhang et al. 2019) which uses attention mechanisms instead of recurrent units. However, a conceptual gap on the flexibility of the neural point process representation still remains. Piece-wise exponential functions (Du et al. 2016; Upadhyay, De, and Rodriguez 2018) only encode intensities that are monotonic between events. The functional RNN representation (Mei and Eisner 2017) is flexible but uses many more parameters. Transformers (Zuo et al. 2020;

Zhang et al. 2019) are generic sequence-to-sequence function approximators (Yun et al. 2020), but the functional form of the Transformer Hawkes point process intensity function is not an universal approximator. Furthermore, intensity functions are non-negative and discontinuous at event times which means neural network approximation results cannot be applied directly.

Recent results shed light on alternative point process representations. Omi, Ueda, and Aihara (2019) uses a positive weight monotone neural network to learn the compensator (the integral of the intensity function). Although it is a generic approximator for compensators, it might assign non-zero probability to invalid inter-arrival times as the compensator can be non-zero at time zero. (Shchur, Biloš, and Günnemann 2020) represents inter-arrival times using normalising flow and mixture models, which can universally approximate any density. However, by defining the point process with the event density, the model cannot account for event sequences which stop naturally (see Section 3). These approaches are promising alternatives but are not a full replacement for intensity functions, which are preferred since they are intuitive and can be superimposed.

In this work, we propose *a class of neural networks that can approximate any point process intensity function to arbitrary accuracy*, along with a proof showing the role of three key constituents: a set of uniformly dense basis functions, a positive transfer function, and an approximator for arbitrary dynamic systems. We implement this proposal using RNNs, the output of which is used to parameterise a set of basis functions upon arrival of each event, as shown in Figure 1. Named *UNIPoint*, the proposed model performs well across synthetic and real world datasets in comparison to the Hawkes process and other neural variants. This work provides a general yet parsimonious representation for temporal point processes, and so forms a solid basis for future development in point process representations that incorporate rich contextual information into event models.

Our primary contributions are:

- A novel architecture that can approximate any point process intensity function to arbitrary accuracy.
- A theoretical guarantee for the flexible point process representation that builds upon the theory of universally approximating continuous functions and dynamic systems.
- UNIPoint — the neural network implementation of the proposed architecture with strong empirical results on both synthetic and real world datasets. Reference code is available online¹.

Notation

$C(X, Y)$ denotes the class of continuous functions mapping from domain X to range Y . Denote \mathbf{R} as the set of real numbers, \mathbf{R}_+ as the non-negative reals and \mathbf{R}_{++} as the strictly positive reals. Define the composition of a function f and a class of functions \mathcal{F} as $f \circ \mathcal{F} = \{f \circ g : g \in \mathcal{F}\}$. The sigmoid function $[1 + \exp(-x)]^{-1}$ is denoted as $\sigma(x)$.

¹<https://github.com/alexandersoen/unipoint>

2 Preliminary: Temporal Point Processes

A temporal point process is an ordered set of event times $\{t_i\}_{i=0}^N$. We typically describe a point process by its conditional intensity function $\lambda(t | \mathcal{H}_{t-})$ which can be interpreted as the instantaneous probability of an event occurring at time t given event history \mathcal{H}_{t-} , consisting of the set of all events before time t . This can be written as (Daley and Vere-Jones 2007):

$$\lambda(t | \mathcal{H}_{t-}) \doteq \lim_{h \downarrow 0^+} \frac{\mathbf{P}(N[t, t+h] > 0 | \mathcal{H}_{t-})}{h}, \quad (1)$$

where $N[t_1, t_2]$ is the number of events occurring between two arbitrary times $t_1 < t_2$. Equation 1 restricts the conditional intensity function to non-negative functions. Given history \mathcal{H}_{t-} , the conditional intensity is a deterministic function of time t . Following standard convention, we refer to the conditional intensity function as simply the intensity function, abbreviating $\lambda(t | \mathcal{H}_{t-})$ to $\lambda^*(t)$.

Point processes can be specified by choosing a functional form for the intensity function. For example, the Hawkes process, one of the simplest interacting point process (Bacry, Mastromatteo, and Muzy 2015), can be defined as follows:

$$\lambda^*(t) = \mu + \sum_{t_i < t} \varphi(t - t_i), \quad (2)$$

where μ specifies the background intensity and $\varphi(t - t_i)$ is the triggering kernel which characterises the self-exciting effects of prior events t_i .

The likelihood of a point process is (Daley and Vere-Jones 2007)

$$L = \left[\prod_{i=1}^N \lambda^*(t_i) \right] \exp \left(- \int_0^T \lambda^*(s) ds \right), \quad (3)$$

where the negated term in the exponential is known as the compensator function $\Lambda^*(t) = \int_0^T \lambda^*(s) ds$.

3 Universal Approximation of Intensities

To represent the influence of past events on future events, point process intensity functions $\lambda^*(t)$ are often continuous between events $(t_{i-1}, t_i]$; with discontinuities only possible at events. For example, the intensity function of the Hawkes process has discontinuities at each event, Eq. (2). Intuitively, this piece-wise continuous characterisation of the intensity function encodes the belief that the process only significantly changes its behaviour when new information (an event) is observed. As such, there are two behaviours of a point process we need to approximate: (1) the continuous intensity function segment between consecutive events, given a fixed event history; and (2) the change in the point process intensity function when an event occurs, so that we can approximate the jump dynamics between events.

We consider an intensity function $\lambda^*(t)$ with fixed observation period $(0, T]$. The intensity function can be segmented by the event times of an event sequence $(t_0, t_1], (t_1, t_2], \dots, (t_{N-1}, t_N], (t_N, t_{N+1}]$, where $t_{N+1} = T$. Given a piece-wise continuous intensity function, the segmented intensity function is continuous: $u_i(\tau) = \lambda^*(t)$ for

$t \in (t_{i-1}, t_i]$, where $\tau = t - t_{i-1} \in (0, t_i - t_{i-1}]$. Thus to approximate the intensity function between consecutive events, we learn a function $\hat{u}(\tau; p_i)$, parameterised by p_i , to approximate *any* of the segmented intensity functions $u_i(\tau)$, where each segment only differs in parameterisation. Then to approximate the jump dynamics of the intensity function we utilise the RNN approximation of a dynamic system, which dictates how the parameters p_i change over time.

To quantify the quality of an approximation, we use the uniform metric between two functions $f, g : X \rightarrow \mathbf{R}$,

$$d(f, g) = \sup_{x \in X} |f(x) - g(x)|. \quad (4)$$

This metric is the maximum difference of the two functions over a shared (compact) domain X . The uniform metric has been used to prove universal approximation properties for neural networks (Hornik, Stinchcombe, and White 1989; Deba0 1993) and RNNs (Schäfer and Zimmermann 2007). Given classes of functions \mathcal{F} and \mathcal{G} , \mathcal{F} is a universal approximator of \mathcal{G} if for any $\varepsilon > 0$ and $g \in \mathcal{G}$, there exists an $f \in \mathcal{F}$ such that $d(f, g) < \varepsilon$. An equivalently expression is: \mathcal{F} is uniformly dense in \mathcal{G} .

Remark. Although we refer to Hawkes point processes as the primary example of a point process to approximate throughout the paper, following the example of (Mei and Eisner 2017), we note that as long as the point process has continuous intensity function between events, our approximation analysis will hold. Thus, in addition to Hawkes processes, the methods proposed in our work can approximate point processes including self correcting processes and non-homogenous Poisson processes with continuous densities.

Approximation Between Two Events

To approximate the time shifted non-negative functions $u_i(\tau)$, we first introduce transfer functions f_+ (Definition 1). We then prove that the class of composed function $f_+ \circ \mathcal{F}$ preserves uniform density (Theorem 1). Given this theorem, we provide a method for constructing uniformly dense classes with sums of basis functions $\Sigma(\phi)$ (Definition 2) which are in turn uniformly dense after composing with f_+ (Corollary 1). We further provide a set of suitable basis functions (Table 1).

Formally, we define the *M-transfer functions* which maps negative outputs of a function to positive values.

Definition 1. A function $f_+ : \mathbf{R} \rightarrow \mathbf{R}_+$ is a *M-transfer function* if it satisfies the following:

1. f_+ is *M-Lipschitz continuous*;
2. $\mathbf{R}_{++} \subseteq f_+[\mathbf{R}]$;
3. And f_+ is strictly increasing on $f_+^{-1}[\mathbf{R}_{++}]$.

Definition 1 provides a wide range of functions. In practice, it is convenient to use softplus function $f_{\text{SP}}(x) = \log(1 + \exp(x))$ which is a 1-transfer function — commonly used in other neural point processes (Mei and Eisner 2017; Omi, Ueda, and Aihara 2019; Zuo et al. 2020). Alternatively, $f_+(x) = \max(0, x)$ could be used; however, this is not differentiable at $x = 0$ which can cause issues in practice. Intuitively, M-transfer function are increasing functions which map to all positive values and have bounded steepness.

When a Gaussian process is used to define an inhomogeneous Poisson process, the link functions serve a similar role to ensure valid intensity functions (Lloyd et al. 2015). However, many of these link function violate the conditions of being a M-transfer function (Donner and Oppor 2018), i.e., the exponential link function $f_+(x) = \exp(x)$ and squared link function $f_+(x) = x^2$ are not M-Lipschitz continuous as they have unbounded derivatives; whereas the sigmoid link function $f_+(x) = \sigma(x)$ is a bounded function (violating condition 2).

Using *M-transfer functions*, we can show that a uniformly dense class of unbounded functions will be uniformly dense for strictly positive functions under composition. These functions are defined with domain $K \subset \mathbf{R}$, a compact subset, which can be set as $K = [0, T]$ for intensity functions.

Theorem 1. Given a class of functions \mathcal{F} which is uniformly dense in $C(K, \mathbf{R})$ and a *M-transfer function* f_+ , the composed class of functions $f_+ \circ \mathcal{F}$ is uniformly dense in $C(K, \mathbf{R}_{++})$ for any compact subset $K \subset \mathbf{R}$.

Proof. Let $f \in C(K, \mathbf{R}_{++})$ and $\varepsilon > 0$ be arbitrary. Since f_+ is strictly increasing and continuous on the preimage of \mathbf{R}_{++} then f_+^{-1} exists, is continuous, and restricted to subdomain \mathbf{R}_{++} . Thus, there exists some $g \in C(K, \mathbf{R})$ such that $f = f_+ \circ g$.

As \mathcal{F} is dense with respect to the uniform metric, for ε/M there exists some $h \in \mathcal{F}$ such that $d(h, g) < \varepsilon/M$. Thus for any $x \in K$,

$$\begin{aligned} |(f_+ \circ h)(x) - f(x)| &= |(f_+ \circ h)(x) - (f_+ \circ g)(x)| \\ &\leq M|h(x) - g(x)| < \varepsilon. \end{aligned}$$

We have $d(f_+ \circ h, f) < \varepsilon$. □

To approximate $u_i(\tau)$ using Theorem 1 we need a family of functions which are able to approximate functions in $C(K, \mathbf{R})$. We consider the family of functions consisting of the sum of basis functions $\phi(\cdot; p_j)$, where $p_j \in \mathcal{P}$ denotes the parameterisation of the basis function ϕ .

Definition 2. Denote $\Sigma(\phi)$ as the class of functions corresponding to the sum of basis functions $\phi : \mathbf{R} \times \mathcal{P} \rightarrow \mathbf{R}$, with parameter space \mathcal{P} , as follows:

$$\left\{ \hat{u} : \mathbf{R} \rightarrow \mathbf{R} \mid \hat{u}(x) = \sum_{j=1}^J \phi(x; p_j), p_j \in \mathcal{P}, J \in \mathbf{N} \right\}.$$

The parameter space \mathcal{P} of a basis function is determined by the parametric form of a chosen basis function $\phi(x; p_j)$. For example, the class composed of exponential basis functions could be defined with parameter space $\mathcal{P} = \mathbf{R}^2$ with functions $\{\phi : \mathbf{R} \rightarrow \mathbf{R} \mid \phi(x) = \alpha \exp(\beta x), \alpha, \beta \in \mathbf{R}\}$. Definition 2 encompasses a wide range of function classes, including neural networks with sigmoid (Cybenko 1989; Hornik, Stinchcombe, and White 1989; Deba0 1993) or rectified linear unit activations (Sonoda and Murata 2017).

The Stone-Weierstrass Theorem provides sufficient conditions for finding basis function for universal approximation.

Theorem 2 (Stone-Weierstrass Theorem (Rudin et al. 1964; Royden and Fitzpatrick 1988)). *Suppose a subalgebra \mathcal{A} of $C(K, \mathbf{R})$, where $K \subset \mathbf{R}$ is a compact subset, satisfies the following conditions:*

1. *For all $x, y \in K$, there exists some $f \in \mathcal{A}$ such that $f(x) \neq f(y)$;*
2. *For all $x_0 \in K$, there exists $f \in \mathcal{A}$ such that $f(x_0) \neq 0$.*

Then \mathcal{A} is uniformly dense in $C(K, \mathbf{R})$.

Thus, by using Theorem 1 and the Stone-Weierstrass theorem, Theorem 2, we arrive at Corollary 1, which gives sufficient conditions for basis functions ϕ to ensure that $f_+ \circ \Sigma(\phi)$ is a universal approximator for $C(K, \mathbf{R}_{++})$.

Corollary 1. *For any compact subset $K \subset \mathbf{R}$ and for any M -transfer function f_+ , if a basis function $\phi(\cdot; p)$ parametrised by $p \in \mathcal{P}$ satisfies the following conditions:*

1. $\Sigma(\phi)$ *is closed under product;*
2. *For any distinct points $x, y \in K$, there exists some $p \in \mathcal{P}$ such that $\phi(x; p) \neq \phi(y; p)$;*
3. *For all $x_0 \in K$, there exists some $p \in \mathcal{P}$ such that $\phi(x_0; p) \neq 0$.*

Then $f_+ \circ \Sigma(\phi)$ is uniformly dense in $C(K, \mathbf{R}_{++})$.

The first condition of Corollary 1 is given such that the set of basis functions $\Sigma(\phi)$ is a subalgebra of $C(X, \mathbf{R})$. The later two conditions are the required preconditions for the Stone-Weierstrass Theorem to hold.

Given the conditions of Corollary 1, some interesting choices for valid basis functions $\phi(x; p)$ are the exponential basis function $\phi_{\text{EXP}}(x) = \alpha \exp(\beta x)$ and the power law basis function $\phi_{\text{PL}}(x) = \alpha(1+x)^{-\beta}$. These basis functions are similar to the exponential and power law Hawkes triggering kernels, which have seen widespread use in many domains (Ogata 1988; Bacry, Mastromatteo, and Muzy 2015; Laub, Taimre, and Pollett 2015; Rizoio et al. 2017).

We note that the class of intensity functions in Theorem 1 and Corollary 1 are strictly positive continuous functions. However, these results generalise to non-negative continuous functions as our definition of intensity functions permits arbitrarily low intensity in $u_i(\tau)$ — where switching from arbitrarily low intensities to zero intensity results in arbitrarily low error with respect to the uniform metric on $(0, T]$.

In Table 1, we provide a selection of interesting basis functions to universally approximate $u_i(\tau) \in C(K, \mathbf{R}_{++})$. One should note that Corollary 1 only provides sufficient conditions, where some of the basis function in Table 1 do not satisfy the precondition. For example, the sigmoid basis function $\phi_{\text{SIG}}(x) = \alpha \sigma(\beta x + \delta)$, $(\alpha, \beta, \delta) \in \mathbf{R}^3$ does not allow $\Sigma(\phi_{\text{SIG}})$ to be closed under product and thus does not satisfy the conditions of Corollary 1. However, the sum of sigmoid basis functions is equivalent to the class of single hidden layer neural networks (Hornik, Stinchcombe, and White 1989; Deba0 1993). Thus, in addition to an appropriate transfer function it does have the universal approximation property for non-negative continuous functions through Theorem 1. Additionally, other basis functions used to define point process intensity functions can be used, such as

Basis Function	Functional Form ϕ	Parameter Space \mathcal{P}
$\phi_{\text{EXP}}^\dagger$	$\alpha \exp(\beta x)$	$(\alpha, \beta) \in \mathbf{R}^2$
ϕ_{PL}^\dagger	$\alpha(1+x)^{-\beta}$	$(\alpha, \beta) \in \mathbf{R} \times \mathbf{R}_+$
$\phi_{\text{COS}}^\dagger$	$\alpha \cos(\beta x + \delta)$	$(\alpha, \beta, \delta) \in \mathbf{R}^3$
$\phi_{\text{SIG}}^\ddagger$	$\alpha \sigma(\beta x + \delta)$	$(\alpha, \beta, \delta) \in \mathbf{R}^3$
ϕ_{ReLU}^*	$\max(0, \alpha x + \beta)$	$(\alpha, \beta) \in \mathbf{R}^2$

Table 1: Basis function universal approximators for intensity functions between two consecutive events. \dagger indicates functions that satisfy Corollary 1; \ddagger one proven in (Cybenko 1989); and $*$ one proven in (Sonoda and Murata 2017).

radial basis functions (Tabibian et al. 2017) that are not generally closed under product but have universal approximation properties (Park and Sandberg 1991).

Approximation for Event Sequences

The approximations to $u_i(\tau)$ use a set of parameters, e.g. (α, β, δ) in Table 1. We denote these parameters vectors as $p_i \in \mathcal{P}$, and the approximated function segment as $\hat{u}_i(\tau; p_i)$. Since each segment $\hat{u}_i(\tau; p_i)$ is uniquely determined by p_i , and the union of all segments approximates $\lambda^*(t)$, we would only need to capture the dynamics in p_i .

We express p_i as the output of a dynamic system.

$$\begin{aligned} s_{i+1} &= g(s_i, t_i) \\ p_i &= \nu(s_i), \end{aligned} \quad (5)$$

where s_{i+1} is the internal state of the dynamic system, g updates the internal state at each step, and ν maps from the internal state to the output.

Theorem 3 (RNN Universal Approximation (Schäfer and Zimmermann 2007)). *Let $g : \mathbf{R}^J \times \mathbf{R}^I \rightarrow \mathbf{R}^J$ be measurable and $\nu : \mathbf{R}^J \rightarrow \mathbf{R}^n$ be continuous, the external inputs $x_i \in \mathbf{R}^I$, the inner states $s_i \in \mathbf{R}^J$, and the outputs $p_i \in \mathbf{R}^n$ (for $i = 1, \dots, N$). Then, any open dynamic system of the form of Eq. (5) can be approximated by an RNN, with sigmoid activation function, to arbitrary accuracy.*

Given that RNNs approximate p_i , we use continuity condition on basis ϕ and in turn \hat{u} to show how to universally approximate an intensity function with an RNN.

Theorem 4. *Let $\{t_i\}_{i=0}^N$ be a sequence of events with $t_i \in [0, T]$ and $\lambda^*(t)$ be an intensity function. Given a parametric family of functions $\mathcal{F} = \{\hat{u}(\cdot; p) : p \in \mathcal{P}\}$ which is uniformly dense in $C([0, T], \mathbf{R}_{++})$ and $\hat{u}(x; p)$ continuous with respect to p for all $x \in [0, T]$. Then there exists a recurrent neural network*

$$\begin{aligned} h_i &= \sigma(Wh_{i-1} + vt_{i-1} + b) \\ \hat{p}_i &= Ah_i && \text{for } t \in (t_{i-1}, t_i] \\ \hat{\lambda}(t) &= \hat{u}(\tau; \hat{p}_i) && \text{and } \tau = t - t_{i-1}, \end{aligned} \quad (6)$$

where σ is a sigmoid activation function and $[W, v, b, A]$ are weights of appropriate shapes, such that $\hat{\lambda}(t)$ approximates $\lambda^*(t)$ with arbitrary precision for all $(0, T]$.

Proof. Let $\varepsilon > 0$ be arbitrary. For any interval $(t_{i-1}, t_i]$, we know from the uniform density of \mathcal{F} that there exists a p_i such that

$$\sup_{\tau \in [0, T]} |\hat{u}_i(\tau; p_i) - u_i(\tau)| \leq \frac{\varepsilon}{2}. \quad (7)$$

By the continuity conditions of \hat{u} , it follows that for each p_i and any $\tau \in [0, T]$ there exists δ_i such that

$$\|p_i - \hat{p}_i\| < \delta_i \implies |\hat{u}(\tau; p_i) - \hat{u}(\tau; \hat{p}_i)| < \frac{\varepsilon}{2} \quad (8)$$

by taking the minimum over δ_τ 's in the $(\varepsilon/2, \delta_\tau)$ -condition of continuity for all $\tau \in [0, T]$ (where the subscript emphasises the range of τ for fixed i).

The LHS of Eq. (8) is the precision needed in our RNN approximator for each interval $(t_{i-1}, t_i]$. We take the minimum approximation discrepancy over the sequence of \hat{p}_i 's, $\delta := \min_i \delta_i$ and use an RNN with precision δ to bound the approximation quality due to \hat{p}_i 's using Theorem 3,

$$\sup_{\tau \in [0, T]} |\hat{u}(\tau; p_i) - \hat{u}(\tau; \hat{p}_i)| < \frac{\varepsilon}{2}. \quad (9)$$

Using the triangle inequality of the uniform metric, we can combine and bound the discrepancies due to \hat{u} in Eq. (7) and those due to \hat{p}_i in Eq. (9),

$$\sup_{\tau \in [0, T]} |u_i(\tau) - \hat{u}(\tau; \hat{p}_i)| < \varepsilon. \quad (10)$$

Eq. (10) holds for all $i \in \{1, \dots, N\}$. Thus uniform density condition for $\lambda^*(t)$ also holds for the piece-wise approximator $\hat{\lambda}(t)$ given by Eq. (6) over the entire sequence. \square

From Theorem 4 and Corollary 1, universal approximation with respect to the uniform metric follows immediately when using basis functions which are continuous with respect to their parameter space, for example Table 1.

Extensions and discussions. While the original work on learning the compensator function (Omi, Ueda, and Aihara 2019) does not provide theoretical backings for its proposal, we note that Theorem 4, combined with universal approximation capabilities of monotone neural networks (Sill 1998), can be used to show that the class of monotonic (increasing) neural networks provide universal approximation for compensator functions. The guarantee described here does not explicitly account for additional dimensions or marks. To extend Theorem 4 in this manner, we consider replacing basis functions $\phi(x)$, which has domain \mathbf{R} , to basis functions with extended domain $\mathbf{R} \times K$ where K is a compact set. For example, K can be a set of discrete finite marks in the case of approximated marked temporal point processes. The universal approximation property would then generalise as long as $\sum(\phi)$ is dense in $C([0, T] \times K, \mathbf{R}_{++})$ and continuous in the parameter space of the basis functions. Likewise, if we want to approximate a spatial point process, we let $K = \mathbf{R}^2$ and find an appropriate set of basis functions with domain $\mathbf{R} \times \mathbf{R}^2$.

It is worth mentioning two distinctions from the intensity free approach (Shchur, Biloš, and Günnemann 2020). First, although density approximation allows for direct event

time sampling, the log-normal mixture representation assumes that an event will always occur on \mathbf{R}_+ — specifically, events cannot naturally stop. Instead, the intensity function representation allow for events to stop with probability $1 - \mathbf{P}(\tau < \infty) = \exp(-\Lambda^*(\infty))$. In other-words, $1 - \mathbf{P}(\tau < \infty)$ is the probability of events not occurring in finite time, which is non-zero when the intensity function decays and stays at zero. Furthermore, the intensity free approach proposed one functional form (log-normal mixture) for approximating densities, whereas we show that a variety of basis functions all fulfil the goal of universal approximation.

4 Implementation with Neural Networks

We propose *UNIPoint*, a neural network architecture implementing a fully flexible intensity function. Let $\{t_i\}_{i=0}^N$ be a sequence of events with corresponding interarrival times $\tau_i = t_i - t_{i-1}$. Let M be the size of the hidden state of the RNN, and $\phi(\cdot; \cdot)$ be the chosen basis function with parameter space \mathcal{P} . Let P denote the dimension of the parameter space. The approximation guarantees (given in Corollary 1) hold in the limit of an infinite number of basis functions, in practice the number of basis functions is a hyper-parameter, denoted as J . This network has four key components.

Recurrent Neural Network. We use a simple RNN cell (Elman 1990), though other popular variants would also work, e.g., LSTM, or GRU. The recurrent unit produces hidden state vector h_i from h_{i-1} the previous hidden state and τ_{i-1} the normalised interarrival time (divided by standard deviation):

$$h_i = f(W h_{i-1} + v \tau_{i-1} + b) \quad (11)$$

Here W, v, b , and h_0 are learnable parameters. f is any activation function compatible with RNN universal approximation, i.e., sigmoid σ (Schäfer and Zimmermann 2007).

Basis Function Parameters are generated using a linear transformation that maps the hidden state vector of the RNN $h_i \in \mathbf{R}^M$ to parameters $p_i = (p_{i1}, \dots, p_{iJ})$,

$$p_{ij} = A_j h_i + B_j, \quad t \in (t_{i-1}, t_i], \quad j \in \{1, \dots, J\}. \quad (12)$$

Here A_j and B_j are learnable parameters and $p_{ij} \in \mathcal{P}$.

Eq. (11) and Eq. (12) defines the RNN which approximates a point processes' underlying dynamic system. The error contribution of these two equations is upper bounded by the sum of their individual contributions (Schäfer and Zimmermann 2007, Theorem 2).

Intensity Function. Using parameters p_{i1}, \dots, p_{iJ} , the intensity function with respect to time since the last event $\tau = t - t_{i-1}$ is defined as:

$$\hat{\lambda}(\tau) = f_{\text{SP}} \left[\sum_{j=1}^J \phi(\tau; p_{ij}) \right], \quad \tau \in (0, t_i - t_{i-1}], \quad (13)$$

where $f_{\text{SP}}(x) = \log(1 + \exp(x))$ is the softplus function.

Loss Function. We use the point process negative log-likelihood, as per Eq. (3). In most cases the integral cannot be calculated analytically so instead we calculate it numerically using Monte-Carlo integration (Press et al. 2007), see

Training settings and the online appendix (Soen et al. 2020, Section F).

Our use of RNNs to encode event history is similar to other neural point process architectures. We note that (Du et al. 2016) only supports monotonic intensities. Our representation is more parsimonious than (Mei and Eisner 2017) since the hidden states need not be functions over time, yet the output can still universally approximate any intensity function. (Omi, Ueda, and Aihara 2019) produce monotonically increasing compensator functions but can have invalid inter-arrival times.

5 Evaluation

We compare the performance of UNIPoint models to various simple temporal point processes and neural network based models on three synthetic datasets and three real world datasets. For the simple temporal point processes we consider self-exciting intensity functions which are piece-wise monotonic (Self-Correcting process (Isham and Westcott 1979) and Exponential Hawkes process (Hawkes 1971)) and non-monotonic (Decaying Sine Hawkes process). The details of dataset preprocessing, model settings and parameter sizes can be found in the appendix (Soen et al. 2020, Section A and B).

Synthetic Datasets

We synthesise datasets from simple temporal point process models, generating 2,048 event sequences each containing 128 events. This results in roughly 262,000 events, which is of the same magnitude tested in (Omi, Ueda, and Aihara 2019). Self-correcting process and exponential Hawkes process datasets have previously been used in other neural point process studies (Du et al. 2016; Omi, Ueda, and Aihara 2019; Shchur, Biloš, and Günnemann 2020). We consider a decaying sine Hawkes process to test whether the models capture non-monotonic self-exciting intensity functions. The following synthetic datasets are used:

Self-Correcting Process. The intensity function is

$$\lambda^*(t) = \exp\left(\nu t - \sum_{t_i < t} \gamma\right),$$

where $\nu = 1$ and $\gamma = 1$.

Exponential Hawkes Process. The intensity function is a Hawkes process with exponential decaying triggering kernel, given by

$$\lambda^*(t) = \mu + \alpha\beta \sum_{t_i < t} \exp(-\beta(t - t_i)),$$

where $\mu = 0.5$, $\alpha = 0.8$, and $\beta = 1$.

Decaying Sine Hawkes Process. The intensity function is a Hawkes process with a sinusoidal triggering kernel product with an exponential decaying triggering kernel:

$$\lambda^*(t) = \mu + \gamma \sum_{t_i < t} (1 + \sin(\alpha(t - t_i))) \exp(-\beta(t - t_i)),$$

where $\mu = 0.5$, $\alpha = 5\pi$, $\beta = 2$, and $\gamma = 1$.

Real World Dataset

We further evaluate the performance of our model with three real world datasets. Although these dataset originally have marks/event types, we ignore such information to test UNIPoint. The real world datasets used are:

MOOC². A dataset of student interactions in online courses (Kumar, Zhang, and Leskovec 2019), previously used for evaluating neural point processes (Shchur, Biloš, and Günnemann 2020). Events correspond to different types of interaction, e.g., watching videos.

Reddit². A dataset of user posts on a social media platform (Kumar, Zhang, and Leskovec 2019), previously used for evaluating neural point processes (Shchur, Biloš, and Günnemann 2020). Each event sequence corresponds to a user’s post behaviour.

StackOverflow (Du et al. 2016). A dataset of events which consists of users gaining badges on a question-answer website. Only users with at least 40 badges between 01-01-2012 and 01-01-2014 are considered.

Baselines

The following traditional and neural network point process models are compared to our models. We implement all but the NeuralHawkes baseline. We also compare to TransformerHawkes (Zuo et al. 2020) but the results are sensitive to model settings, the observations from which are discussed in the appendix (Soen et al. 2020, Section C).

Exponential Hawkes Process The point process likelihood is optimised to determine parameter μ , α , and β in intensity function

$$\lambda^*(t) = \mu + \alpha\beta \sum_{t_i < t} \exp(-\beta(t - t_i)).$$

Power Law Hawkes Process. The point process likelihood is optimised to determine parameters μ , α , and β in intensity function

$$\lambda^*(t) = \mu + \alpha \sum_{t_i < t} (t - t_i + \delta)^{-(1+\beta)}.$$

The δ parameter is fixed at 0.5 to compensate for the difficulty of the power law intensity function being infinity when $t - t_i + \delta = 0$ (Bacry, Mastromatteo, and Muzy 2015).

RMTTP (Du et al. 2016). We implement the RMTTP neural network architecture as a baseline. The intensity function of RMTTP

$$\lambda^*(t) = \exp(v^T h_i + w(t - t_{i-1}) + b) \quad (14)$$

is defined with respect to the RNN hidden state h_i . We use a RNN size of 48 for testing.

FullyNeural (Omi, Ueda, and Aihara 2019). We also implement the fully neural network point process. The integral of the intensity function (compensator) is defined as a neural network with RNN hidden state and event time input. We use a RNN size of 48 and fully connected layer of size 48 to produce the compensator.

²<https://github.com/srijankr/jodie/>

Dataset		Synthetic			Real World		
Models		SelfCorrecting	ExpHawkes	DecayingSine	MOOC	Reddit	StackOverflow
Baseline	ExpHawkes	$-0.994 \pm .001$	$0.044 \pm .037$	$-0.838 \pm .019$	$3.578 \pm .060$	$-0.100 \pm .039$	$-1.031 \pm .025$
	PLHawkes	$-0.994 \pm .001$	$0.036 \pm .037$	$-0.845 \pm .019$	$0.532 \pm .070$	$-0.787 \pm .035$	$-0.918 \pm .024$
	RMTTP	$-0.776 \pm .003$	$0.054 \pm .038$	$-0.864 \pm .020$	$2.040 \pm .098$	$-0.336 \pm .031$	$-0.864 \pm .022$
	FullyNeural	$-0.789 \pm .003$	$0.059 \pm .037$	$-0.833 \pm .020$	$4.699 \pm .054^\dagger$	$0.206 \pm .046^\dagger$	$-0.810 \pm .022$
	NeuralHawkes	$-0.777 \pm .006^\dagger$	$0.066 \pm .037^\dagger$	$-0.821 \pm .021^\dagger$	$4.641 \pm .110$	$0.201 \pm .048$	$-0.801 \pm .023^\dagger$
UNIPoint	ExpSum	$-0.774 \pm .008^\ddagger$	$0.056 \pm .042$	$-0.828 \pm .020$	$3.114 \pm .125$	$0.151 \pm .045$	$-0.812 \pm .023$
	PLSum	$-0.779 \pm .006$	$0.064 \pm .038^\ddagger$	$-0.829 \pm .020$	$4.939 \pm .085^\ddagger$	$0.162 \pm .046$	$-0.814 \pm .023$
	ReLUsum	$-0.780 \pm .007$	$0.059 \pm .039$	$-0.828 \pm .021$	$4.676 \pm .075$	$0.221 \pm .046^\ddagger$	$-0.810 \pm .023$
	CosSum	$-0.777 \pm .008$	$0.062 \pm .039$	$-0.828 \pm .020$	$4.471 \pm .075$	$0.139 \pm .044$	$-0.814 \pm .023$
	SigSum	$-0.776 \pm .007$	$0.064 \pm .038$	$-0.827 \pm .020^\ddagger$	$4.346 \pm .076$	$0.170 \pm .045$	$-0.814 \pm .023$
MixedSum	$-0.779 \pm .007$	$0.062 \pm .038$	$-0.828 \pm .020$	$4.928 \pm .085$	$0.201 \pm .047$	$-0.804 \pm .023^\ddagger$	

Table 2: Averaged log-likelihood scores with corresponding 95% confidence intervals. A higher score is better; the best of the baselines are indicated by \dagger and the best of the UNIPoint models are indicated by \ddagger . Bold indicates results when the difference between \dagger and \ddagger are *significantly better* (t-test $p = 0.05$).

NeuralHawkes³ (Mei and Eisner 2017). We utilise the reference implementation for NeuralHawkes (Mei, Qin, and Eisner 2019), which provides a neural network architecture that encodes the decaying nature of Hawkes process exponential kernels in the LSTM of the model. We use a LSTM size of 48 and default parameters for other model settings.

Training Settings

We fit a UNIPoint model for each of the five basis function types described in Table 1 with softplus transfer functions and 64 basis functions with learnable parameters. The mixture of basis functions, MixedSum, is used, with 32 power law and 32 ReLU basis functions. We study effects of the number of basis functions in the appendix (Soen et al. 2020, Section E). We fit models for all synthetic and real world datasets, with a 60 : 20 : 20 train-validation-test split. Our models are implemented in PyTorch⁴.

During training, we use a single sample per event interval to calculate the loss function as we find using multiple samples does not improve performance, as shown in the appendix (Soen et al. 2020, Section F). All UNIPoint models tested employ an RNN with 48 hidden units, a batch size of 64, and are trained using Adam (Kingma and Ba 2014) with $L2$ weight decay set to 10^{-5} . The validation set is used for early stopping: training halts if the validation loss does not improve by more than 10^{-4} for 100 successive minibatches. The training for one of the real world datasets (e.g., StackOverflow) takes approximately 1 day.

We further test UNIPoint using LSTMs in Appendix G and an alternative transfer function in Appendix H.

Evaluation Metrics

Holdout Log-likelihood. We calculate the log-likelihood of event sequences using Eq. (3). We numerically calculate the integral term with Monte-Carlo integration (Press et al. 2007) if it cannot be calculated analytically.

³<https://github.com/hmeijatju/neural-hawkes-particle-smoothing>

⁴<https://pytorch.org> (Paszke et al. 2017)

Total Variation. We use total variation as it mimics the uniform metric as they both depend on the difference between the true and approximate intensity function. It is defined as $TV(f, g) = \int |f(s) - g(s)|^2 ds$. Total variation can only be used on synthetic datasets where the true intensity is known. To calculate it, we use Monte-Carlo integration (Press et al. 2007). We do not compute total variation for NeuralHawkes as the reference implementation does not allow the intensity function to be evaluated over fixed event histories.

6 Results

Table 2 reports log-likelihoods of all models across the three synthetic and three real world datasets. Figure 2 reports the total variations of intensity functions for the synthetic datasets and relative log-likelihood (calculated by subtracting the log-likelihood of UNIPoint ReLUsum) for the three real world datasets. The total variation scores are only available for synthetic datasets since calculating the total variation requires a ground truth intensity function.

Synthetic datasets. Contrasting the log-likelihood and total variation metrics reveal interesting insights about model performance. The SelfCorrecting dataset has a piece-wise monotonically increasing intensity function. Both metrics indicate that ExpHawkes, PLHawkes, and RMTTP underperform the other approaches by a large margin, since they are restricted to piece-wise monotone intensity functions. All UNIPoint variants perform well, achieving average likelihoods within 0.01 of each other. ExpSum is the best variant, possibly due to its exponential shape matching that of the ground-truth SelfCorrecting intensity function.

For the ExpHawkes dataset, the ExpHawkes baseline has the lowest total variation (close to zero, as expected) but not the best holdout log-likelihood. This indicates that models with good log-likelihood scores still have the potential to overfit given the wrong intensity function representation. Despite UNIPoint’s guarantees with infinite basis functions, ExpSum shows significantly better total variation scores than other UNIPoint models here — showing that, selection of basis functions for specific datasets is important.

For DecayingSine, the intensity between events are non-

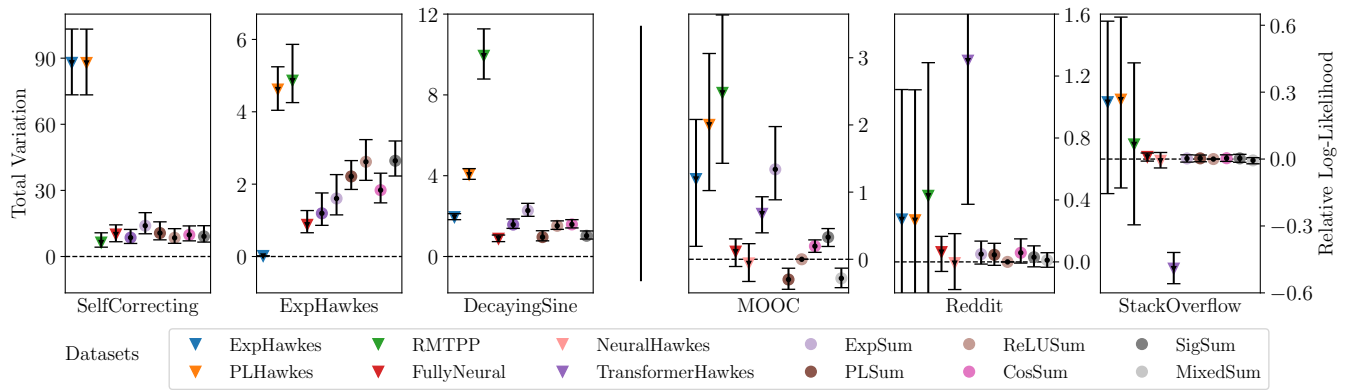


Figure 2: Total variation of intensity functions for synthetic datasets (left) and relative log-likelihood of event sequences for real world datasets standardised by subtracting the score of ReLUSum UNIPoint (right). *Lower score is better*. Markers correspond to the mean of the score and error bars to the interquartile range. A missing marker indicate a mean above the visible axis range.

monotonic. All UNIPoint variants perform comparably on both the log-likelihood and total variation metric. The FullyNeural approach performs comparably with the UNIPoint variants on total variation, but is inferior on log-likelihood. This is likely due to it assigning non-zero probabilities to negative event times. NeuralHawkes has the best log-likelihood for this dataset, but the difference with respect to SigSum is not significant.

In addition, we visualise intensity functions learnt by UNIPoint and other approaches, see the appendix (Soen et al. 2020, Section D). The neural baseline models learn similar intensity functions to UNIPoint in ExpHawkes. However in the case of the MOOC dataset, RMTTP learns an intensity function that is different to those learnt by the other neural point processes. Meanwhile, FullyNeural does not exhibit strong decaying components in the intensity function.

Real-world datasets. For all three real-world datasets, baselines ExpHawkes, PLHawkes, and RMTTP significantly under-perform in comparison to the rest of the approaches. This likely occurs due to their inability to support non-monotone intensity functions in inter-event intervals.

We observe that UNIPoint variants are significantly better than the baselines for MOOC and Reddit. UNIPoint is second best (to NeuralHawkes) on StackOverflow dataset, but the difference is not statistically significant. NeuralHawkes performs strongly on the StackOverflow dataset, potentially because it has the closest architecture to the UNIPoint ExpSum variant, while also being more complex. In particular NeuralHawkes has time decaying hidden states and LSTM recurrent units rather than the a perceptron recurrent unit and vector-formed hidden state of UNIPoint. The StackOverflow dataset has a longer average sequence length than MOOC and Reddit, which would advantage the LSTM recurrent units over the standard RNN — since the RNN is more likely to suffer from vanishing or exploding gradients than the LSTM which allows for long-term dependencies (Hochreiter and Schmidhuber 1997). Details on dataset characteristics can be found in the appendix (Soen et al. 2020, Section A). One peculiar result is the performance of

ExpSum in the MOOC dataset. The reason for the poor performance is that the exponential basis function is unstable with large interarrival times, which can cause numeric overflow or underflow. The performance of UNIPoint variants depend greatly on the particular basis function used for each dataset. We find that no single type of basis function ensures that a UNIPoint model performs best over all datasets. For example, in the MOOC dataset, ignoring ExpSum, the UNIPoint models have log-likelihood scores from 4.346 ± 0.076 to 4.939 ± 0.085 .

Using mixture of basis function, MixedSum provides good overall performance. Among the UNIPoint variants, it is either the best or a close second across all datasets, suggesting that using different types of basis functions improves model flexibility in practice even with a fixed parameter budget. We also observe an improvement in performance when more basis functions are used, see Appendix E.

Overall, our evaluations demonstrate the power of UNIPoint for modelling complex intensity function that are not piece-wise monotone. Results on real-world datasets show models with flexible intensity functions outperform Hawkes processes. Open questions remain on which neural architectures, among the ones with universal approximation power, strike the best balance of representational power, parsimony, and learnability.

7 Conclusion

We develop a new method for universally approximating the conditional intensity function of temporal point processes. This is achieved by breaking down the intensity function into piece-wise continuous functions and approximating each segment with a sum of basis functions, followed by a transfer function. We also propose UNIPoint, a neural implementation of the approximator. Evaluations on synthetic and real world benchmarks demonstrates that UNIPoint consistently outperform the less flexible alternatives. Future work include: investigating methods for selecting and tuning different basis functions and further theoretical work on representation complexity, expressiveness and learnability.

Acknowledgments

This research was supported in part by the Australian Research Council Project DP180101985 and AOARD Project 20IOA064. This research was supported by use of the Nectar Research Cloud, a collaborative Australian research platform supported by the National Collaborative Research Infrastructure Strategy (NCRIS).

References

- Bacry, E.; Mastromatteo, I.; and Muzy, J.-F. 2015. Hawkes processes in finance. *Market Microstructure and Liquidity* 1(01): 1550005.
- Cybenko, G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems* 2(4): 303–314.
- Daley, D. J.; and Vere-Jones, D. 2007. *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media.
- Debao, C. 1993. Degree of approximation by superpositions of a sigmoidal function. *Approximation Theory and its Applications* 9(3): 17–28.
- Donner, C.; and Opper, M. 2018. Efficient Bayesian inference of sigmoidal Gaussian Cox processes. *The Journal of Machine Learning Research* 19(1): 2710–2743.
- Du, N.; Dai, H.; Trivedi, R.; Upadhyay, U.; Gomez-Rodriguez, M.; and Song, L. 2016. 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*, 1555–1564. ACM.
- Elman, J. L. 1990. Finding structure in time. *Cognitive science* 14(2): 179–211.
- Embrechts, P.; Liniger, T.; and Lin, L. 2011. Multivariate Hawkes processes: an application to financial data. *Journal of Applied Probability* 48(A): 367–378.
- Hawkes, A. G. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58(1): 83–90.
- Hochreiter, S.; and Schmidhuber, J. 1997. LSTM can solve hard long time lag problems. In *Advances in neural information processing systems*, 473–479.
- Hornik, K.; Stinchcombe, M.; and White, H. 1989. Multilayer feedforward networks are universal approximators. *Neural networks* 2(5): 359–366.
- Huang, H.; Wang, H.; and Mak, B. 2019. Recurrent poisson process unit for speech recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6538–6545.
- Isham, V.; and Westcott, M. 1979. A self-correcting point process. *Stochastic processes and their applications* 8(3): 335–347.
- Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980.
- Kumar, S.; Zhang, X.; and Leskovec, J. 2019. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1269–1278.
- Laub, P. J.; Taimre, T.; and Pollett, P. K. 2015. Hawkes processes. *arXiv preprint arXiv:1507.02822*.
- Li, S.; Xiao, S.; Zhu, S.; Du, N.; Xie, Y.; and Song, L. 2018. Learning temporal point processes via reinforcement learning. In *Advances in neural information processing systems*, 10781–10791.
- Lloyd, C.; Gunter, T.; Osborne, M.; and Roberts, S. 2015. Variational inference for Gaussian process modulated Poisson processes. In *International Conference on Machine Learning*, 1814–1822.
- Mei, H.; and Eisner, J. M. 2017. The neural hawkes process: A neurally self-modulating multivariate point process. In *Advances in Neural Information Processing Systems*, 6754–6764.
- Mei, H.; Qin, G.; and Eisner, J. 2019. Imputing Missing Events in Continuous-Time Event Streams. In *Proceedings of the International Conference on Machine Learning*.
- Mishra, S.; Rizoïu, M.-A.; and Xie, L. 2016. Feature driven and point process approaches for popularity prediction. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*, 1069–1078.
- Ogata, Y. 1988. Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical association* 83(401): 9–27.
- Omi, T.; Ueda, N.; and Aihara, K. 2019. Fully Neural Network based Model for General Temporal Point Processes. In *Advances in Neural Information Processing Systems 32*, 2120–2129. Curran Associates, Inc.
- Park, J.; and Sandberg, I. W. 1991. Universal approximation using radial-basis-function networks. *Neural computation* 3(2): 246–257.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic Differentiation in PyTorch. In *NIPS Autodiff Workshop*.
- Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; and Flannery, B. P. 2007. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press.
- Rizoïu, M.-A.; Xie, L.; Sanner, S.; Cebrian, M.; Yu, H.; and Van Hentenryck, P. 2017. Expecting to be HIP: Hawkes intensity processes for social media popularity. In *Proceedings of the 26th International Conference on World Wide Web*, 735–744.
- Royden, H. L.; and Fitzpatrick, P. 1988. *Real analysis*, volume 32. Macmillan New York.
- Rudin, W.; et al. 1964. *Principles of mathematical analysis*, volume 3. McGraw-hill New York.

- Schäfer, A. M.; and Zimmermann, H.-G. 2007. Recurrent neural networks are universal approximators. *International journal of neural systems* 17(04): 253–263.
- Shchur, O.; Biloš, M.; and Günnemann, S. 2020. Intensity-Free Learning of Temporal Point Processes. In *International Conference on Learning Representations*.
- Sill, J. 1998. Monotonic networks. In *Advances in neural information processing systems*, 661–667.
- Soen, A.; Mathews, A.; Grixti-Cheng, D.; and Xie, L. 2020. Appendix - UNIPoint: Universally Approximating Point Processes Intensities. *CoRR* abs/2007.14082: 11–13. URL <https://arxiv.org/pdf/2007.14082.pdf#page=11>.
- Sonoda, S.; and Murata, N. 2017. Neural network with unbounded activation functions is universal approximator. *Applied and Computational Harmonic Analysis* 43(2): 233–268.
- Tabibian, B.; Valera, I.; Farajtabar, M.; Song, L.; Schölkopf, B.; and Gomez-Rodriguez, M. 2017. Distilling information reliability and source trustworthiness from digital traces. In *Proceedings of the 26th International Conference on World Wide Web*, 847–855.
- Upadhyay, U.; De, A.; and Rodriguez, M. G. 2018. Deep reinforcement learning of marked temporal point processes. In *Advances in Neural Information Processing Systems*, 3168–3178.
- Yun, C.; Bhojanapalli, S.; Rawat, A. S.; Reddi, S.; and Kumar, S. 2020. Are Transformers universal approximators of sequence-to-sequence functions? In *International Conference on Learning Representations*.
- Zhang, Q.; Lipani, A.; Kirnap, O.; and Yilmaz, E. 2019. Self-attentive Hawkes processes. *arXiv preprint arXiv:1907.07561*.
- Zhao, Q.; Erdogdu, M. A.; He, H. Y.; Rajaraman, A.; and Leskovec, J. 2015. Seismic: A self-exciting point process model for predicting tweet popularity. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 1513–1522.
- Zuo, S.; Jiang, H.; Li, Z.; Zhao, T.; and Zha, H. 2020. Transformer Hawkes Process. *arXiv preprint arXiv:2002.09291*.