

# Hawkes Process Based on Controlled Differential Equations

Minju Jo\*, Seungji Kook\* and Noseong Park

Yonsei University, Seoul, South Korea

{alfsow112,202132139,noseong}@yonsei.ac.kr

## Abstract

Hawkes processes are a popular framework to model the occurrence of sequential events, i.e., occurrence dynamics, in several fields such as social diffusion. In real-world scenarios, the inter-arrival time among events is *irregular*. However, existing neural network-based Hawkes process models not only i) fail to capture such complicated irregular dynamics but also ii) resort to heuristics to calculate the log-likelihood of events since they are mostly based on neural networks designed for regular discrete inputs. To this end, we present the concept of Hawkes process based on controlled differential equations (HP-CDE), by adopting the neural controlled differential equation (neural CDE) technology which is an analogue to *continuous* RNNs. Since HP-CDE continuously reads data, i) irregular time-series datasets can be properly treated preserving their uneven temporal spaces, and ii) the log-likelihood can be exactly computed. Moreover, as both Hawkes processes and neural CDEs are first developed to model complicated human behavioral dynamics, neural CDE-based Hawkes processes are successful in modeling such occurrence dynamics. In our experiments with 4 real-world datasets, our method outperforms existing methods by non-trivial margins.

## 1 Introduction

Real-world phenomena typically correspond to the occurrence of sequential events with *irregular* time intervals and *numerous* event types, ranging from online social network activities to personalized healthcare and so on [Zhao *et al.*, 2015; Enguehard *et al.*, 2020; Stoyan and Penttinen, 2000; Mohler *et al.*, 2011; Ogata, 1999]. Hawkes processes and Poisson point process are typically used to model those sequential events [Hawkes, 1971; Miles, 1970; Streit, 2010]. However, their basic assumptions are too stringent to model such complicated dynamics, e.g., all past events should influence the occurrence of the current event. To this end, many advanced techniques have been proposed for the past several years, ranging from classical recurrent neural network (RNN) based models such as RMTTP [Du *et al.*, 2016] and

Model	Exact log-likelihood	How to model dynamics
NHP, SAHP, THP	X	Discrete
HP-CDE	O ( $\lambda^*$ is continuous.)	Continuous & robust to irregular dynamics

Table 1: Comparison of neural network-based Hawkes process models.  $\lambda^*$  denotes the conditional intensity function (cf. Eqs. (4), (6), and (7)).

NHP [Mei and Eisner, 2017] to recent transformer models like SAHP [Zhang *et al.*, 2020] and THP [Zuo *et al.*, 2020]. Even so, they still do not treat data in a fully continuous way but resort to heuristics, which is sub-optimal in processing irregular events [Chen *et al.*, 2018; Choi *et al.*, 2021; Yildiz *et al.*, 2019]. Likewise, their heuristic approaches to model the continuous time domain impede solving the multivariate integral of the log-likelihood calculation in Eq. (4), leading to approximation methods such as the Monte Carlo sampling (cf. Table 1). As a consequence, the strict constraint and/or the inexact calculation of the log-likelihood may induce inaccurate predictions.

In this work, therefore, we model the occurrence dynamics based on differential equations, not only directly handling the sequential events in a continuous time domain but also exactly solving the integral of the log-likelihood. One more inspiration of using differential equations is that they have shown several non-trivial successes in modeling human behavioral dynamics [Poli *et al.*, 2019; Rubanova *et al.*, 2019; Jeon *et al.*, 2021] — in particular, we are interested in controlled differential equations. To our knowledge, therefore, we first answer the question of whether occurrence dynamics can be modeled as controlled differential equations.

Controlled differential equations (CDEs [Lyons *et al.*, 2004]) are one of the most suitable ones for building human behavioral models. CDEs were first developed by a financial mathematician to model complicated dynamics in financial markets which is a typical application domain of Hawkes processes since financial transactions are temporal point processes. In particular, neural controlled differential equations (neural CDEs [Kidger *et al.*, 2020]), whose initial value problem (IVP) is written as below, are a set of techniques to learn

CDEs from data with neural networks:

$$\begin{aligned} \mathbf{h}(t_b) &= \mathbf{h}(t_a) + \int_{t_a}^{t_b} f(\mathbf{h}(t); \theta_f) dZ(t) \\ &= \mathbf{h}(t_a) + \int_{t_a}^{t_b} f(\mathbf{h}(t); \theta_f) \frac{dZ(t)}{dt} dt, \end{aligned} \quad (1)$$

where  $f$  is a CDE function, and  $\mathbf{h}(t)$  is a hidden vector at time  $t$ .  $Z(t)$  is a continuous path created from discrete sequential observations (or events)  $\{(\mathbf{z}_j, t_j)\}_{j=a}^b$  by an appropriate algorithm<sup>1</sup>, where in our case,  $\mathbf{z}_j$  is a vector containing the information of  $j$ -th occurrence, and  $t_j \in [t_a, t_b]$  contains the time-point of the occurrence, i.e.,  $t_j < t_{j+1}$ . Note that neural CDEs keep reading the time-derivative of  $Z(t)$  over time, denoted  $\dot{Z}(t) := \frac{dZ(t)}{dt}$ , and for this reason, neural CDEs are in general, considered as *continuous* RNNs. In addition, NCDEs are known to be superior in processing irregular time series [Lyons *et al.*, 2004].

Given the neural CDE framework, we propose **Hawkes Process** based on **C**ontrolled **D**ifferential **E**quations (HP-CDE). We let  $\mathbf{z}_j$  be the sum of the event embedding and the positional embedding and create a path  $Z(t)$  with the linear interpolation method which is a widely used interpolation algorithm for neural CDEs (cf. Figure 2). To get the exact log-likelihood, we use an ODE solver to calculate the non-event log-likelihood. Calculating the non-event log-likelihood involves the integral problem in Eq. (4), and our method can solve it exactly since conditional intensity function  $\lambda^*$ , which indicates an instantaneous probability of an event, is defined in a continuous manner over time by the neural CDE technology. In addition, we have three prediction layers to predict the event log-likelihood, the event type, and the event occurrence time (cf. Eqs. (8), (12), (13) and Figure 3).

We conduct event prediction experiments with 4 datasets and 4 baselines. Our method shows outstanding performance in all three aspects: i) event type prediction, ii) event time prediction, and iii) log-likelihood. Our contributions are as follows:

1. We model the *continuous* occurrence dynamics under the framework of neural CDE whose original theory was developed for describing *irregular non-linear* dynamics. Many real-world Hawkes process datasets have irregular inter-arrival times of events.
2. We then exactly solve the integral problem in Eq. (4) to calculate the non-event log-likelihood, which had been done typically through heuristic methods before our work.

## 2 Preliminaries

### 2.1 Multivariate Point Processes

Multivariate point processes are a generative model of an event sequence  $X = \{(k_j, t_j)\}_{j=1}^N$  and  $x_j = (k_j, t_j)$  indicates  $j$ -th event in the sequence. This event sequence is a subset of an event stream under a continuous time interval

<sup>1</sup>One can use interpolation algorithms or neural networks for creating  $Z(t)$  from  $\{(\mathbf{z}_j, t_j)\}_{j=a}^b$  [Kidger *et al.*, 2020].

$[t_1, t_N]$ , and an observation  $x_j$  at time  $t_j$  has an event type  $k_j \in \{1, \dots, K\}$ , where  $K$  is total number of event types. The arrival time of events is defined as  $t_1 < t_2 < \dots < t_N$ . The point process model learns a probability for every  $(k, t)$  pair, where  $k \in \{1, \dots, K\}$ ,  $t \in [t_1, t_N]$ .

The key feature of multivariate point processes is the intensity function  $\lambda_k(t)$ , i.e., the probability that a type- $k$  event occurs at the infinitesimal time interval  $[t, t + dt)$ . The Hawkes process, one popular point process model, assumes that the intensity  $\lambda_k(t)$  of type  $k$  can be calculated by past events before  $t$ , so-called history  $\mathcal{H}_t$ , and its form is as follows:

$$\lambda_k^*(t) := \lambda_k(t|\mathcal{H}_t) = \mu_k + \sum_{j:t_j < t} \psi_k(t - t_j), \quad (2)$$

where  $\lambda^*(t) = \sum_{k=1}^K \lambda_k^*(t)$ ,  $\mu_k$  is the base intensity, and  $\psi_k(\cdot)$  is a pre-determined decaying function for type  $k$ . We use  $*$  to represent conditioning on the history  $\mathcal{H}_t$ . According to the formula, all the past events affect the probability of new event occurrence with different influences. However, the intensity converges to the base intensity if the decaying function becomes close to zero.

Currently, a deep learning mechanism is applied to Hawkes processes by parameterizing the intensity function. For instance, RNNs are used in the neural Hawkes process (NHP) [Mei and Eisner, 2017], and its intensity function is defined as follows:

$$\lambda^*(t) = \sum_{k=1}^K \phi_k(\mathbf{w}_k^\top \mathbf{h}(t)), \quad t \in [t_1, t_N], \quad (3)$$

where  $\phi_k(\cdot)$  is the softplus function,  $\mathbf{h}(t)$  is a hidden state from RNNs, and  $\mathbf{w}_k$  is a weight for each event type. The softplus function keeps intensity values positive. However, one downside of NHP is that RNN-based models assume that events have regular intervals. Thus, one of the main issues in NHP is how to fit a model to a continuous irregular time domain.

### 2.2 Neural Network-based Hawkes Processes

Hawkes processes are a popular temporal predicting framework in various fields since it predicts both *when*, which *type* of events would happen with mathematical approaches. It is especially widely used in sociology fields to capture the diffusion of information [Hardiman *et al.*, 2013; Kobayashi and Lambiotte, 2016; Da Fonseca and Zaatour, 2014], seismology fields to model when earthquakes and aftershocks occur, medical fields to track the status of patients [Choi *et al.*, 2015; Garetto *et al.*, 2021], and so on.

For enhancing the performance of Hawkes processes, a lot of deep learning approaches have been applied. The two basic approaches are the recurrent marked temporal point process (RMTTP [Du *et al.*, 2016]) and the neural Hawkes process (NHP [Mei and Eisner, 2017]). RMTTP is the first model that combines RNNs into point processes, and NHP is a Hawkes process model with an RNN-parameterized intensity function. Based on NHP, the self-attentive Hawkes process (SAHP [Zhang *et al.*, 2020]) attaches self-attention modules to reflect the relationships between events. Additionally, the

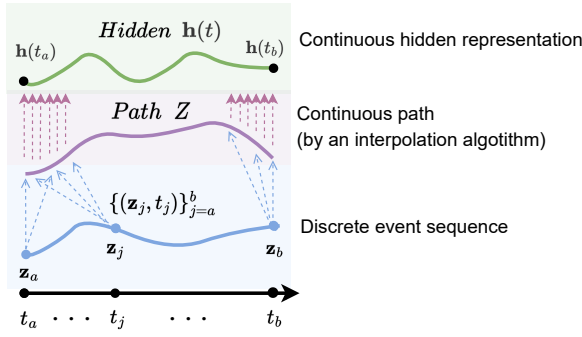


Figure 1: Visualization of the continuous hidden state of the neural CDE model

transformer Hawkes process (THP [Zuo *et al.*, 2020]) uses the transformer technology [Vaswani *et al.*, 2017], one of the most popular structures in natural language processing, to capture both short-term and long-term temporal dependencies of event sequences.

One important issue of neural network-based Hawkes process is how to handle irregular time-series datasets. To deal with this issue, NHP uses continuous-time LSTMs, whose memory cell exponentially decays. SAHP and THP both employ modified positional encoding schemes to represent irregular time intervals since the conventional encoding assumes regular spaces between events. However, all mentioned approaches still do not explicitly process irregular time-series. In contrast to them, our HP-CDE is robust to irregular time-series since the original motivation of neural CDEs is better processing irregular time-series by constructing continuous RNNs.

### 2.3 Neural Controlled Differential Equations as continuous RNNs

Neural controlled differential equations (neural CDEs) are normally regarded as a continuous analogue to RNNs since they process the time-derivative of the continuous path  $Z(t)$ . Especially, neural CDEs retain their continuous properties by using the interpolated path  $Z$  made of discrete data  $\{(z_j, t_j)\}_{j=a}^b$  and solving the Riemann-Stieltjes integral to get  $\mathbf{h}(t_b)$  from  $\mathbf{h}(t_a)$  as shown in Eq. (1) — in particular, this problem to derive  $\mathbf{h}(t_b)$  from the initial condition  $\mathbf{h}(t_a)$  is known as initial value problem (IVP) (cf. Figure 1). At first, to make the interpolated continuous path  $Z$ , linear interpolation or natural cubic spline interpolation is generally used among several interpolation methods. Then, we use existing ODE solvers to solve the Riemann-Stieltjes integral problem with  $\dot{\mathbf{h}}(t) := \frac{d\mathbf{h}(t)}{dt} = f(\mathbf{h}(t); \theta_f) \frac{dZ(t)}{dt}$ .

### 2.4 Maximum Likelihood Estimation in Temporal Point Process

Most of the neural temporal point process frameworks choose the maximum likelihood estimation (MLE) [Aitchison and Silvey, 1958] as one of the main training objectives. In order to enable the MLE training, getting the log-probability of every sequence  $X$  is required, which consists of formulas using

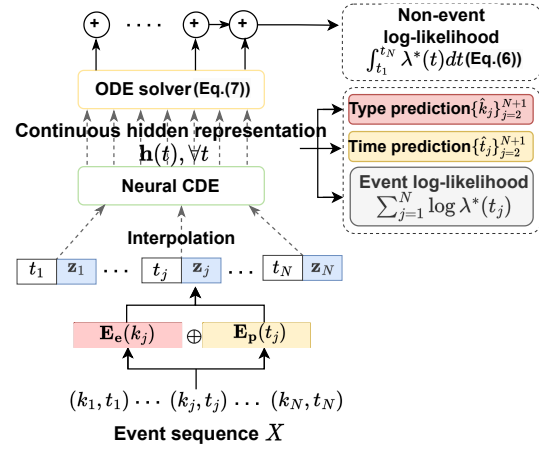


Figure 2: Our proposed HP-CDE architecture

intensity functions conditioned on the history  $\mathcal{H}_t = \{(k_j, t_j) : t_j < t\}$ . Thus, log-probability for any event sequence  $X$  whose events are observed in an interval  $[t_1, t_N]$  is as follows:

$$\log p(X) = \sum_{j=1}^N \log \lambda^*(t_j) - \int_{t_1}^{t_N} \lambda^*(t) dt, \quad (4)$$

where  $\sum_{j=1}^N \log \lambda^*(t_j)$  denotes the event log-likelihood and  $\int_{t_1}^{t_N} \lambda^*(t) dt$  means the non-event log-likelihood. Non-event log-likelihood represents sum of the infinite number of non-events' log-probabilities in  $[t_1, t_N]$ , except the infinitesimal times when the event occurs. In the case of the event log-likelihood, it is comparably easy to compute as the formula is simply a sum of the intensity functions. However, it is challenging to compute the non-event log-likelihood, due to its integral computation. Due to the difficulty, NHP, SAHP, THP and many other models use approximation methods, such as Monte Carlo integration [Robert and Casella, 2005] and numerical integration methods [Stoer and Bulirsch, 2013], to get the value. However, since those methods do not exactly solve the integral problem, numerical errors are inevitable.

## 3 Proposed Method

In this section, we describe our *explicitly continuous* Hawkes process model, called HP-CDE, based on the neural CDE framework which is considered as continuous RNNs. Owing to the continuous property of the proposed model, the exact log-likelihood, especially for the non-event log-likelihood part with its challenging integral calculation, can also be computed through ODE solvers. That is, our proposed model reads event sequences with irregular inter-arrival times in a continuous manner, and exactly computes the log-likelihood.

### 3.1 Overall Workflow

Figure 2 shows comprehensive designs of our proposed model, HP-CDE. The overall workflow is as follows:

1. Given the event sequence  $X = \{(k_j, t_j)\}_{j=1}^N$ , i.e., event type  $k_j$  at time  $t_j$ , the embeddings  $\{\mathbf{E}_e(k_j), \mathbf{E}_p(t_j)\}_{j=1}^N$

are made through the encoding processes, where  $\mathbf{E}_e(k_j)$  is an embedding of  $k_j$  and  $\mathbf{E}_p(t_j)$  is a positional embedding of  $t_j$ .

2. Then we use  $\{\mathbf{E}_e(k_j) \oplus \mathbf{E}_p(t_j)\}_{j=1}^N$  as the discrete hidden representations  $\{\mathbf{z}_j\}_{j=1}^N$ . In other words,  $\mathbf{z}_j = \mathbf{E}_e(k_j) \oplus \mathbf{E}_p(t_j)$ , i.e., the element-wise summation of the two embeddings.
3. An interpolation algorithm is used to create the continuous path  $Z(t)$  from  $\{(\mathbf{z}_j, t_j)\}_{j=1}^N$  — we augment the time information  $t_j$  to each  $\mathbf{z}_j$ .
4. Using the continuous path  $Z(t)$ , a neural CDE layer calculates the final continuous hidden representation  $\mathbf{h}(t)$  for all  $t$ . At the same time, an ODE solver integrates the continuous intensity function  $\lambda^*(t)$  which is calculated from  $\mathbf{h}(t)$  (cf. Eq. (7)) to calculate the non-event log-likelihood. In addition, there are three prediction layers to predict the event type, time, and log-likelihood (cf. Figure 3).

We provide more detailed descriptions for each step in the following subsections with the well-posedness of our model.

### 3.2 Embedding

We embed both the type and time of each event into separate vectors and then add them. To be more specific, we map each event type to an embedding vector  $\mathbf{E}_e(k)$ , which is trainable. With trigonometric functions, we embed the time information to a vector  $\mathbf{E}_p(t)$ , which is called positional encoding in transformer language models (cf. Appendix A). We use the sum of the two embeddings,  $\{\mathbf{E}_e(k_j) \oplus \mathbf{E}_p(t_j)\}_{j=1}^N$  as the discrete hidden representations  $\{\mathbf{z}_j\}_{j=1}^N$ , i.e.,  $\mathbf{z}_j = \mathbf{E}_e(k_j) \oplus \mathbf{E}_p(t_j)$ .

### 3.3 Occurrence Dynamics and Continuous Intensity Function

With  $\{\mathbf{z}_j\}_{j=1}^N$ , we calculate the *continuous* hidden representation  $\mathbf{h}(t_j)$  for any arbitrary  $j$ , where  $t_1 \leq t_j$ , based on the neural CDE framework as follows:

$$\mathbf{h}(t_j) = \mathbf{h}(t_1) + \int_{t_1}^{t_j} f(\mathbf{h}(t); \theta_f) \frac{dZ(t)}{dt} dt, \quad (5)$$

where  $Z(t)$  is a continuous path created by an interpolation algorithm from  $\{(\mathbf{z}_j, t_j)\}_{j=1}^N$ . The well-posedness<sup>2</sup> of neural CDEs is proved in [Lyons *et al.*, 2004, Theorem 1.3] under the Lipschitz continuity requirement (cf. Appendix B). Neural CDE layer is able to generate the continuous hidden representation  $\mathbf{h}(t_j)$ , where  $t_1 \leq t_j$ , even when the sequence  $\{(\mathbf{z}_j, t_j)\}_{j=1}^N$  is an irregular time-series, i.e., the inter-arrival time varies from one case to another.

This continuous property enables our model to exactly solve the integral problem of the non-event log-likelihood.

<sup>2</sup>The well-posedness of an initial value problem means that i) its unique solution, given an initial value, exists, and ii) its solutions continuously change as initial values change.

That is, the non-event log-likelihood can be re-written as the following ODE form:

$$\mathbf{a}(t_N) = \int_{t_1}^{t_N} \lambda^*(t) dt, \quad (6)$$

where the conditional intensity function of Eqs. (2) and (3) is, in our case, the sum of the conditional intensity functions of all event types as follows:

$$\lambda^*(t) = \sum_{k=1}^K \lambda_k^*(t), \quad \lambda_k^*(t) = \phi_k(\mathbf{W}_k^{\text{intst}\top} \mathbf{h}(t_j)), \quad (7)$$

where  $\mathbf{W}_k^{\text{intst}}$  is a weight matrix of intensity about type  $k$ , and therefore,  $\mathbf{W}_k^{\text{intst}\top} \mathbf{h}(t_j)$  is a linear projected representation which has the history of events before time  $t_j$ .  $\phi_k(x) := \beta_k \log(1 + \exp(x/\beta_k))$  is the softplus function with a parameter  $\beta_k$  to be learned. The softplus function is used to restrict the intensity function to have only positive values. Therefore, the log-probability of HP-CDE for any event sequence  $X$  is redefined from Eq. (4) as:

$$\log p(X) = \sum_{j=1}^N \log \lambda^*(t_j) - \mathbf{a}(t_N). \quad (8)$$

As a result, we can naturally define the following augmented ODE, where  $\mathbf{h}(t)$  and  $\mathbf{a}(t)$  are combined:

$$\frac{d}{dt} \begin{bmatrix} \mathbf{h}(t) \\ \mathbf{a}(t) \end{bmatrix} = \begin{bmatrix} f(\mathbf{h}(t); \theta_f) \frac{dZ(t)}{dt} \\ \lambda^*(t) \end{bmatrix} \quad (9)$$

and

$$\begin{bmatrix} \mathbf{h}(t_1) \\ \mathbf{a}(t_1) \end{bmatrix} = \begin{bmatrix} \pi(\mathbf{z}(t_1); \theta_\pi) \\ 0 \end{bmatrix}, \quad (10)$$

where  $\pi$  is a fully connected layer. The neural network  $f$  is defined as follows:

$$f(\mathbf{h}(t)) = \text{Tanh}(\pi_M(\text{ELU}(\cdots(\text{ELU}(\pi_1(\mathbf{h}(t))))))), \quad (11)$$

which consists of fully connected layers with the ELU or the hyperbolic tangent activation. The number of layers  $M$  is a hyperparameter.

In Zuo *et al.* [Zuo *et al.*, 2020], the generated hidden representations from the self-attention module of their transformer have discrete time stamps, and therefore, its associated intensity function definition is inevitably discrete. For that reason, they rely on a heuristic method, e.g., Monte Carlo method, to calculate the non-event log-likelihood. In our case, however, the physical time is modeled in a continuous manner and therefore, the exact non-event log-likelihood can be calculated as in Eq. (6).

### 3.4 Prediction Layer

Our model has three prediction layers as in other Hawkes process models: i) next event type, ii) next event time, and iii) the event log-likelihood (cf. Figure 3). We use Eq. (7) to calculate the event log-likelihood.




Figure 3: Prediction layer of HP-CDE

For the event type and time predictions, we predict  $\{\hat{t}_j\}_{j=2}^{N+1}$  and  $\{\hat{k}_j\}_{j=2}^{N+1}$  after reading  $X = \{(k_j, t_j)\}_{j=1}^N$ . For the event type prediction layer, we use the following method:

$$\begin{aligned} \hat{\mathbf{p}}_{j+1} &= \text{Softmax}(\mathbf{W}^{\text{type}} \mathbf{h}(t_j)), \\ \hat{k}_{j+1} &= \arg \max_k \hat{\mathbf{p}}_{j+1}(k), \end{aligned} \quad (12)$$

where  $\mathbf{W}^{\text{type}}$  is a trainable parameter and  $\hat{\mathbf{p}}_{j+1}(k)$  is the probability of type  $k$  at time  $t_{j+1}$ . For the event time prediction layer, we use the following definition:

$$\hat{t}_{j+1} = \mathbf{W}^{\text{time}} \mathbf{h}(t_j), \quad (13)$$

where  $\mathbf{W}^{\text{time}}$  is a trainable parameter.

### 3.5 Training Algorithm

Our loss definition consists of three parts. The first part is the following MLE loss, i.e. maximizing the log-likelihood (cf. Eq. (8)):

$$\max \sum_{i=1}^S \log p(X_i), \quad (14)$$

where  $S$  is the number of training samples. While training, the log-intensity of each observed event increases and the non-event log-likelihood decreases in the whole interval  $[t_1, t_N]$ .

The second loss is the event type loss function which is basically a cross-entropy term as follows:

$$\mathcal{L}_{\text{type}}(X) = \sum_{j=2}^{N+1} -\mathbf{k}_j^\top \log(\hat{\mathbf{p}}_j), \quad (15)$$

where  $\mathbf{k}_j$  is a one-hot vector for the event type  $k_j$ . In the case of the event time loss, we use the inter-arrival time  $\tau_i = t_i - t_{i-1}$  to compute the loss as follows:

$$\mathcal{L}_{\text{time}}(X) = \sum_{j=2}^{N+1} (\tau_j - \hat{\tau}_j)^2. \quad (16)$$

Therefore, the overall objective function of HP-CDE can be written as follows:

$$\min \sum_{i=1}^S -\alpha_1 \log p(X_i) + \mathcal{L}_{\text{type}}(X_i) + \alpha_2 \mathcal{L}_{\text{time}}(X_i), \quad (17)$$

where  $\alpha_1$  and  $\alpha_2$  are hyperparameters.

---

### Algorithm 1 How to train HP-CDE

---

**Input:** Training data  $\mathcal{D}_{\text{train}}$ , Iteration numbers  $\text{max\_iter}$

- 1: Initialize all the parameters of the embedding and the neural CDE layer
  - 2:  $\text{iter} \leftarrow 0$
  - 3: **while**  $\text{iter} < \text{max\_iter}$  **do**
  - 4:   Sample a mini-batch  $\{X_i\}_{i=1}^S \in \mathcal{D}_{\text{train}}$
  - 5:   Calculate the embedding vectors, i.e,  $\mathbf{E}_e(k_j)$ , and  $\mathbf{E}_p(t_j)$
  - 6:   Calculate the discrete hidden representation  $\mathbf{z}_j, \forall j$
  - 7:   Calculate the continuous hidden representation  $\mathbf{h}(t)$  using neural CDE and compute the non-event log-likelihood using ODE solver with Eq. (6) over time
  - 8:   Update the parameters with Eq. (17)
  - 9:   **if** the loss does not decrease for  $\delta$  iterations **then**
  - 10:     **exit**
  - 11:   **end if**
  - 12: **end while**
  - 13: **return** the trained parameters
- 

In Alg. (1), we show the training algorithm. We first initialize all the parameters. From our training data, we randomly build a mini-batch  $\{X_i\}_{i=1}^S$  in Line 4 — the optimal mini-batch size varies from one dataset to another. After feeding the constructed mini-batch into our model, we calculate the discrete and continuous hidden representations in Lines 6 and 7. With the loss in Eq. (17), we train our model. We repeat the steps  $\text{max\_iter}$  times.

## 4 Experiments

### 4.1 Experimental Environments

#### Experimental Settings

In this section, we compare the model performance of HP-CDE with 4 state-of-the-art baselines on 4 datasets. Each dataset is split into the training set and the testing set. The training set is used to tune the hyperparameters and the testing set is used to measure the model performance. We evaluate the models with three metrics: i) log-likelihood (LL) of  $X = \{(k_j, t_j)\}_{j=1}^N$ , ii) accuracy (ACC) on the event type prediction, and iii) root mean square error (RMSE) on the event time prediction. We train each model 100 epochs and report the mean and standard deviation of the evaluation metrics of five trials with different random seeds. We compare our model with various baselines (cf. Section 2.2): Recurrent Marked Temporal Point Process (RMTTP)<sup>3</sup>, Neural Hawkes Process (NHP)<sup>4</sup>, Self-Attentive Hawkes Process (SAHP)<sup>5</sup>, and Transformer Hawkes Process (THP)<sup>6</sup>. More details including hyperparameter configurations are in Appendix C.

<sup>3</sup><https://github.com/dunan/NeuralPointProcess>

<sup>4</sup><https://github.com/hongyuanmei/neural-hawkes-particle-smoothing>

<sup>5</sup>[https://github.com/QiangAIRresearcher/sahp\\_repo](https://github.com/QiangAIRresearcher/sahp_repo)

<sup>6</sup><https://github.com/SimiaoZuo/Transformer-Hawkes-Process>

Dataset	Model	LL $\uparrow$	ACC $\uparrow$	RMSE $\downarrow$	Memory usage(MB)	Training time(m)
MIMIC	RMTTP	-1.222 $\pm$ 0.080	0.823 $\pm$ 0.014	1.035 $\pm$ 0.023	3	0.004
	NHP	-0.647 $\pm$ 0.051	0.534 $\pm$ 0.015	0.976 $\pm$ 0.020	13	0.045
	SAHP	-0.859 $\pm$ 0.328	0.555 $\pm$ 0.171	1.138 $\pm$ 0.059	34	0.037
	THP	-0.233 $\pm$ 0.012	0.741 $\pm$ 0.021	0.856 $\pm$ 0.040	9	0.012
	HP-CDE	<b>2.573<math>\pm</math>0.201</b>	<b>0.847<math>\pm</math>0.007</b>	<b>0.726<math>\pm</math>0.042</b>	58	0.058
MemeTracker	RMTTP	NaN	0.006 $\pm$ 0.000	NaN	1,708	0.425
	NHP	-9.395 $\pm$ 2.814	0.044 $\pm$ 0.003	441.293 $\pm$ 0.233	5,096	12.263
	SAHP	2.160 $\pm$ 0.324	0.009 $\pm$ 0.000	521.672 $\pm$ 4.071	32,894	6.642
	THP	-5.717 $\pm$ 0.649	0.015 $\pm$ 0.000	446.477 $\pm$ 2.665	891	2.610
	HP-CDE	<b>3.846<math>\pm</math>0.626</b>	<b>0.151<math>\pm</math>0.005</b>	<b>441.223<math>\pm</math>3.480</b>	3,669	3.817
Retweet	RMTTP	NaN	0.490 $\pm$ 0.000	NaN	210	0.044
	NHP	-9.082 $\pm$ 0.125	0.547 $\pm$ 0.010	16,630.956 $\pm$ 0.217	750	17.820
	SAHP	1.904 $\pm$ 0.566	0.505 $\pm$ 0.067	16,648.339 $\pm$ 1.436	13,276	0.197
	THP	-7.347 $\pm$ 0.268	0.499 $\pm$ 0.013	<b>15,050.470<math>\pm</math>26.712</b>	1,582	0.142
	HP-CDE	<b>6.844<math>\pm</math>0.539</b>	<b>0.552<math>\pm</math>0.009</b>	15,849.218 $\pm$ 269.068	197	6.236
StackOverFlow	RMTTP	-1.894 $\pm$ 0.002	0.429 $\pm$ 0.000	1.321 $\pm$ 0.002	27	0.040
	NHP	-7.726 $\pm$ 0.581	0.434 $\pm$ 0.015	1.027 $\pm$ 0.027	449	3.556
	SAHP	-0.431 $\pm$ 0.225	0.244 $\pm$ 0.002	4.525 $\pm$ 1.098	11,080	0.147
	THP	-0.554 $\pm$ 0.001	0.449 $\pm$ 0.001	<b>0.973<math>\pm</math>0.001</b>	4,585	0.169
	HP-CDE	<b>7.348<math>\pm</math>0.466</b>	<b>0.452<math>\pm</math>0.001</b>	0.996 $\pm$ 0.017	44	6.878

Table 2: Experimental results.  $\uparrow$  (resp.  $\downarrow$ ) denotes that the higher (resp. lower) the better, and we use boldface to denote the best score.

Dataset	$K$	Sequence length			# Events
		Min	Average	Max	
MIMIC	75	2	4	26	1,930
MemeTracker	5000	1	3	31	123,639
Retweet	3	50	109	264	2,173,533
StackOverFlow	22	41	72	720	345,116

Table 3: Characteristics of datasets used in experiments

Model	Dataset	
	MIMIC	MemeTracker
RMTTP	0.385 $\pm$ 0.037	0.000 $\pm$ 0.000
NHP	0.126 $\pm$ 0.018	0.011 $\pm$ 0.002
SAHP	0.108 $\pm$ 0.112	0.000 $\pm$ 0.000
THP	0.162 $\pm$ 0.016	0.000 $\pm$ 0.000
HP-CDE	<b>0.452<math>\pm</math>0.035</b>	<b>0.069<math>\pm</math>0.004</b>

Table 4: F1 score ( $\uparrow$ ) for imbalanced datasets

### Datasets

To show the efficacy and applicability of our model, we evaluate using various real-world data. MemeTracker [Leskovec and Krevl, June 2014], Retweet [Zhao *et al.*, 2015], and StackOverFlow [Leskovec and Krevl, June 2014], are collected from Stackoverflow, web articles, and Twitter, respectively. We also use a medical dataset, called MIMIC [Johnson *et al.*, 2016]. We deliberately choose the datasets with various average sequence lengths and event type numbers  $K$  to show the general efficacy of our model. The average sequence length ranges from 3 to 109, and the number of event types  $K$  ranges from 3 to 5000 (cf. Table 3). That is, we cover not only from simple to complicated ones, but also from short-term to long-term sequences. Details of datasets are in Appendix C.3.

### 4.2 Experimental Results

We show the experimental results of each model on MIMIC, MemeTracker, Retweet, and StackOverFlow in Table 2. We analyze the results in three aspects: i) the event prediction, ii) the log-likelihood, and iii) the model complexity. Ablation and sensitivity analyses are in Appendix D and E.

#### Event Prediction

HP-CDE outperforms other baselines with regards to both the event type and the event time prediction in most cases as re-

ported in Table 2. To be specific, in terms of accuracy, HP-CDE shows the best performance in every dataset. These results imply that processing data in a continuous manner is important when it is in a continuous time domain. Even though HP-CDE only shows the lowest RMSE on datasets with short sequence length, MIMIC and MemeTracker, we provide the solution to lower RMSE of HP-CDE when using datasets with long sequence length in Section 4.3.

For the imbalanced datasets of MIMIC and MemeTracker, where only 20% of types occupy 90% and 70% of events each, we do the following additional analyses. Notably, HP-CDE attains an accuracy of 0.151 in MemeTracker, which is up to 243% higher than those of baselines, and an RMSE of 0.726 in MIMIC, about 15% lower. Furthermore, we use the macro F1 score to measure the quality of type predictions. As shown in Table 4, our model shows the best F1 score in both of the imbalanced datasets. Especially for MemeTracker, models with attention modules have relatively low F1 scores, indicating that when there exist too many classes and if they are imbalanced, attentions are overfitted to several frequently occurring classes. This phenomenon is also observed in Figure 4. In Figure 4, HP-CDE shows the most diverse predictions in terms of the number of predicted classes.




Figure 4: The number of classes in test data vs. the number of classes in correct event type predictions, i.e., hits. HP-CDE provides not only accurate but also diverse predictions.

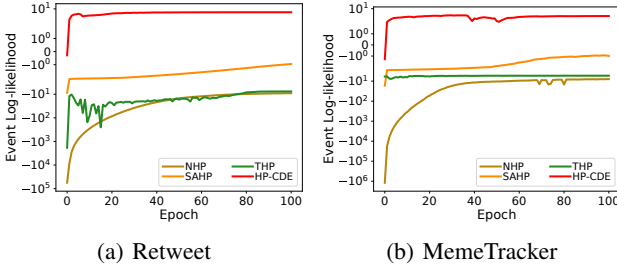


Figure 5: Training curves on Retweet and MemeTracker. HP-CDE shows the highest log-likelihood with the fastest convergence speed.

Particularly, in Figure 4 (b), HP-CDE successfully predicts for 1,164 classes among 2,604 classes, which is almost 50% of the classes in test data, whereas NHP, SAHP, and THP predict only for 217, 4, and 7 classes, respectively.

Regardless of the characteristics of datasets, e.g., the number of types, the degree of imbalance, and so on, our model shows outstanding prediction results, which prove the importance of continuous processing and computing the exact log-likelihood leading to more accurate learning of dynamics.

### Log-likelihood Calculation

As shown in Table 2, our models always show the best log-likelihood, outperforming others by large margins, on every dataset. One remarkable point is that our log-likelihood is always positive, while baselines show negative values in many cases. That is, in HP-CDE, the event log-likelihood exceeds the non-event log-likelihood at all times.

Figure 5 shows the training curves of models fitted on Retweet and MemeTracker in a log-scale. First of all, HP-CDE show the best log-likelihood at every training epoch. Overall, except THP, the log-likelihood of MemeTracker tends to be more unstable than that of Retweet, since MemeTracker has about 1,700 times more event types than Retweet.

### Memory Usage

Table 2 also recaps the model complexity. Exactly calculating the non-event log-likelihood using ODE solvers incurs additional memory usage, so that the model uses bigger memory than those of other sampling methods such as Monte Carlo sampling. Especially when the number of event types  $K$  is large, i.e., MIMIC and MemeTracker, the complexity of HP-CDEs increases as we exactly compute the non-event log-




Figure 6: Additional study on long-sequence datasets, comparing accuracy and RMSE of HP-CDE-AT to HP-CDE and THP.

likelihood for every event type. However, when  $K$  is relatively small, owing to the adjoint sensitivity method [Chen *et al.*, 2018; Kidger *et al.*, 2020], HP-CDE’s memory footprint notably decreases. For example, when using Retweet with  $K = 3$ , the space complexity of HP-CDE is almost 1% of that of THP.

### 4.3 Additional Study on the Long Sequence Length

While HP-CDE shows a good performance on the datasets with relatively short sequence lengths, i.e., MIMIC and MemeTracker, its RMSE results on others with longer sequence lengths, i.e., Retweet and StackOverflow, are slightly larger than those of THP’s. Therefore, to effectively deal with long sequence datasets, we put the self-attention part of transformer [Vaswani *et al.*, 2017] right before the neural CDE layer and name the model HP-CDE-AT. Experimental results of HP-CDE-AT in comparison with HP-CDE and THP, which shows the highest score among baselines, are summarized in Figure 6. According to Figure 6 (a), HP-CDE-AT achieves the smallest RMSE, improving the performance of the original HP-CDE model. Remarkably, in Figure 6 (b), HP-CDE-AT even shows the best performance on StackOverflow in both metrics, accuracy and RMSE. In conclusion, since HP-CDE-AT attains overall best results on longer datasets, HP-CDE-AT is one good option for long sequence datasets (cf. Appendix F).

## 5 Conclusions

Temporal point processes are frequently used in real-world applications to model occurrence dynamics in various fields. In particular, deep learning-based Hawkes process models have been extensively studied. However, we identified the two possible enhancements from the literature and presented HP-CDE to overcome the limitations. First, we use neural CDEs to model occurrence dynamics since one of their main application areas is to model uncertainties in human behaviors. Second, we exactly calculate the non-event log-likelihood which is one important part of the training objective. Existing work uses heuristic methods for it, which makes the training process unstable sometimes. In our experiments, consequently, our presented method significantly outperforms them and shows the most diverse predictions, i.e., the least overfitting.



## Acknowledgements

Noseong Park is the corresponding author. This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No. 2020-0-01361, Artificial Intelligence Graduate School Program at Yonsei University, 10%), and (2022-0-01032, Development of Collective Collaboration Intelligence Framework for Internet of Autonomous Things, 45%) and (No.2022-0-00113, Developing a Sustainable Collaborative Multi-modal Lifelong Learning Framework, 45%).

## Ethical Statement

MIMIC contains much personal health information. However, it was released after removing observations, such as diagnostic reports and physician notes, using a rigorously evaluated deidentification system to protect the privacy of the patients who have contributed their information. Therefore, our work does not have any related ethical concerns.

## Contribution Statement

Minju Jo and Seungji Kook contributed equally.

## References

- [Aitchison and Silvey, 1958] J. Aitchison and S. D. Silvey. Maximum-Likelihood Estimation of Parameters Subject to Restraints. *The Annals of Mathematical Statistics*, 29(3):813 – 828, 1958.
- [Chen *et al.*, 2018] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *NeurIPS*, 2018.
- [Choi *et al.*, 2015] Edward Choi, Nan Du, Robert Chen, Le Song, and Jimeng Sun. Constructing disease network and temporal progression model via context-sensitive Hawkes process. In *2015 IEEE International Conference on Data Mining*, pages 721–726. IEEE, 2015.
- [Choi *et al.*, 2021] Jeongwhan Choi, Jinsung Jeon, and Noseong Park. Lt-ocf: Learnable-time ode-based collaborative filtering. In *CIKM*, 2021.
- [Da Fonseca and Zaatour, 2014] José Da Fonseca and Riadh Zaatour. Hawkes process: Fast calibration, application to trade clustering, and diffusive limit. *Journal of Futures Markets*, 34(6):548–579, 2014.
- [Du *et al.*, 2016] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. Recurrent marked temporal point processes: Embedding event history to vector. *KDD '16*, page 1555–1564, New York, NY, USA, 2016. Association for Computing Machinery.
- [Enguehard *et al.*, 2020] Joseph Enguehard, Dan Busbridge, Adam Bozson, Claire Woodcock, and Nils Y. Hammerla. Neural temporal point processes for modelling electronic health records, 2020.
- [Garetto *et al.*, 2021] Michele Garetto, Emilio Leonardi, and Giovanni Luca Torrisi. A time-modulated Hawkes process to model the spread of covid-19 and the impact of counter-measures. *Annual reviews in control*, 51:551–563, 2021.
- [Hardiman *et al.*, 2013] Stephen J Hardiman, Nicolas Bercot, and Jean-Philippe Bouchaud. Critical reflexivity in financial markets: a Hawkes process analysis. *The European Physical Journal B*, 86(10):1–9, 2013.
- [Hawkes, 1971] Alan G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.
- [Jeon *et al.*, 2021] Jinsung Jeon, Soyoung Kang, Minju Jo, Seunghyeon Cho, Noseong Park, Seonghoon Kim, and Chiyoung Song. Lightmove: A lightweight next-poi recommendation for taxicab rooftop advertising. In *CIKM*, 2021.
- [Johnson *et al.*, 2016] Alistair Johnson, Tom Pollard, Lu Shen, Li-wei Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Celi, and Roger Mark. MIMIC-III, a freely accessible critical care database, 2016.
- [Kidger *et al.*, 2020] Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *Advances in Neural Information Processing Systems*, 33:6696–6707, 2020.
- [Kobayashi and Lambiotte, 2016] Ryota Kobayashi and Renaud Lambiotte. Tideh: Time-dependent Hawkes process for predicting retweet dynamics. In *Tenth International AAAI Conference on Web and Social Media*, 2016.
- [Leskovec and Krevl, June 2014] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- [Lyons *et al.*, 2004] Terry Lyons, M. Caruana, and T. Lévy. *Differential Equations Driven by Rough Paths*. Springer, 2004. École D’Eté de Probabilités de Saint-Flour XXXIV - 2004.
- [Mei and Eisner, 2017] Hongyuan Mei and Jason Eisner. The neural Hawkes process: A neurally self-modulating multivariate point process. In *Advances in Neural Information Processing Systems*, Long Beach, December 2017.
- [Miles, 1970] Roger E Miles. On the homogeneous planar poisson point process. *Mathematical Biosciences*, 6:85–127, 1970.
- [Mohler *et al.*, 2011] George O Mohler, Martin B Short, P Jeffrey Brantingham, Frederic Paik Schoenberg, and George E Tita. Self-exciting point process modeling of crime. *Journal of the American Statistical Association*, 106(493):100–108, 2011.
- [Ogata, 1999] Yosihiko Ogata. Seismicity analysis through point-process modeling: A review. *Seismicity patterns, their statistical significance and physical meaning*, pages 471–507, 1999.
- [Poli *et al.*, 2019] Michael Poli, Stefano Massaroli, Junyoung Park, Atsushi Yamashita, Hajime Asama, and



- Jinkyoo Park. Graph neural ordinary differential equations. *arXiv preprint arXiv:1911.07532*, 2019.
- [Robert and Casella, 2005] Christian P. Robert and George Casella. *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2005.
- [Rubanova *et al.*, 2019] Yulia Rubanova, Ricky T. Q. Chen, and David K Duvenaud. Latent ordinary differential equations for irregularly-sampled time series. In *NeurIPS*. 2019.
- [Stoer and Bulirsch, 2013] Josef Stoer and Roland Bulirsch. *Introduction to numerical analysis*, volume 12. Springer Science & Business Media, 2013.
- [Stoyan and Penttinen, 2000] Dietrich Stoyan and Antti Penttinen. Recent applications of point process methods in forestry statistics. *Statistical science*, pages 61–78, 2000.
- [Streit, 2010] Roy L Streit. The poisson point process. In *Poisson Point Processes*, pages 11–55. Springer, 2010.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, volume 30. Curran Associates, Inc., 2017.
- [Yildiz *et al.*, 2019] Cagatay Yildiz, Markus Heinonen, and Harri Lahdesmaki. Ode2vae: Deep generative second order odes with bayesian neural networks. In *NeurIPS*. 2019.
- [Zhang *et al.*, 2020] Qiang Zhang, Aldo Lipani, Omer Kirnap, and Emine Yilmaz. Self-attentive Hawkes process. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11183–11193. PMLR, 13–18 Jul 2020.
- [Zhao *et al.*, 2015] Qingyuan Zhao, Murat A. Erdogdu, Hera Y. He, Anand Rajaraman, and Jure Leskovec. 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*, page 1513–1522, 2015.
- [Zuo *et al.*, 2020] Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. Transformer hawkes process. ICML’20. JMLR.org, 2020.