

Modeling the Intensity Function of Point Process via Recurrent Neural Networks

Shuai Xiao,¹ Junchi Yan,^{2,3*} Xiaokang Yang,¹ Hongyuan Zha,⁴ Stephen M. Chu³
¹ Shanghai Jiao Tong University ² East China Normal University ³ IBM Research – China ⁴ Georgia Tech
 {benjaminforever,xkyang}@sjtu.edu.cn, jcyan@sei.ecnu.edu.cn, schu@us.ibm.com, zha@cc.gatech.edu

Abstract

Event sequence, asynchronously generated with random timestamp, is ubiquitous among applications. The precise and arbitrary timestamp can carry important clues about the underlying dynamics, and has lent the event data fundamentally different from the time-series whereby series is indexed with fixed and equal time interval. One expressive mathematical tool for modeling event is point process. The intensity functions of many point processes involve two components: the background and the effect by the history. Due to its inherent spontaneousness, the background can be treated as a time series while the other need to handle the history events. In this paper, we model the background by a Recurrent Neural Network (RNN) with its units aligned with time series indexes while the history effect is modeled by another RNN whose units are aligned with asynchronous events to capture the long-range dynamics. The whole model with event type and timestamp prediction output layers can be trained end-to-end. Our approach takes an RNN perspective to point process, and models its background and history effect. For utility, our method allows a black-box treatment for modeling the intensity which is often a pre-defined parametric form in point processes. Meanwhile end-to-end training opens the venue for reusing existing rich techniques in deep network for point process modeling. We apply our model to the predictive maintenance problem using a log dataset by more than 1000 ATMs from a global bank headquartered in North America.

Introduction

Event sequence is becoming increasingly available in a variety of applications such as e-commerce transactions, social network activities, conflicts, and equipment failures etc. The event data can carry rich information not only about the event attribute (e.g. type, participator) but also the timestamp $\{z_i, t_i\}_{i=1}^N$ indicating when the event occurs. Being treated as a random variable when the event is stochastically generated in an asynchronous manner, the timestamp makes the

event sequence fundamentally different from the time series (Montgomery, Jennings, and Kulahci 2015) with equal and fixed time interval, whereby the time point only serves as a role for index $\{y_t\}_{t=1}^T$. A major line of research (Aalen, Borgan, and Gjessing 2008) has been devoted to study event sequence, especially exploring the timestamp information to model the underlying dynamics of the system, whereby point process (Snyder and Miller 2012) has been a powerful and compact framework in this direction.

Recently there are many machine learning based models for scalable point process modeling. We attribute the progressions in this direction in part to the smart mathematical reformulations and optimization techniques e.g. (Lewis and Mohler 2011; Zhou, Zha, and Song 2013b; 2013a), as well as novel parametric forms for the conditional intensity function (Shen et al. 2014; Ertekin, Rudin, and McCormick 2015; Xu et al. 2016) as carefully designed by researchers' prior knowledge to capture the character of the dataset in study. However, one major limitation of the parametric forms of point process is due to their specialized and restricted expression capability for arbitrary distributed event data which trends to be oversimplified or even infeasible for capturing the problem complexity in real applications. Moreover, it runs at the risk of model underfitting due to the misjudgement on model choice. Recent works e.g. (Zhou, Zha, and Song 2013b) start to turn to non-parametric form to fit the structure of a point process, but their method is under the Hawkes process formulation, while this formulation runs at the risk of model mis-choice.

In this paper, we view the conditional intensity of a point process as a nonlinear mapping between the predicted transient occurrence intensity of events with different types, and the model input information of event participators, event profile and the system history. Such a nonlinear mapping is expected to be complex and flexible enough to model various characters of real event data for its application utility.

In fact, deep learning models, such as Convolutional Neural Networks (CNNs) (LeCun et al. 1998), Recurrent Neural Networks (RNNs) (Pascanu, Mikolov, and Bengio 2013) have attracted wide attention in recent vision, speech and language communities, and many of them has dominated the competing results on perceptual benchmark tasks e.g. (Rusakovsky et al. 2015). In particular, we turn to RNNs as a natural way to encode such nonlinear and dynamic mapping,

*Correspondence author is Junchi Yan. This research was partially supported by The National Key Research and Development Program of China (2016YFB1001003), NSFC (61602176, 61672231, 61527804, 61521062), STCSM (15JC1401700, 14XD1402100), China Postdoctoral Science Foundation Funded Project (2016M590337), the 111 Program (B07022) and NSF (IIS-1639792, DMS-1620345).
 Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

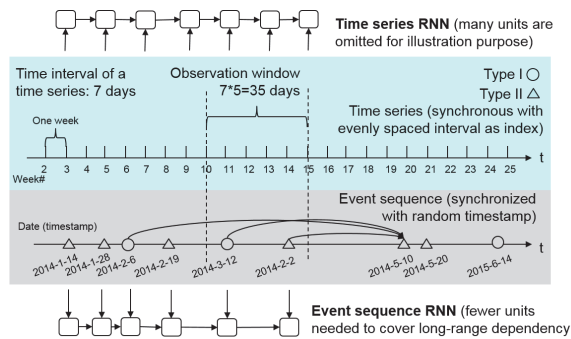


Figure 1: Time series and event sequence can be synergistically modeled. The former can be used to timely capture the recent window for the time-varying features, while the latter can capture the long-range dependency over time. Note the dependency in an event sequence can be easily captured by an event sequence LSTM with less than 5 steps, while it takes too much more steps if using a time series with a fixed time interval e.g. 7 days in the figure (Note many unit steps in the top time series are omitted in the figure for clarity).

in an effort for modeling an end-to-end nonlinear intensity mapping without any prior knowledge.

Key idea and highlights Our model interprets the conditional intensity function of a point process as a nonlinear mapping, which is synergistically established by a composite neural network with two RNNs as its building blocks. As illustrated in Fig.1, time series (top row) and event sequence (bottom row) are distinct to each other in that time series is more suitable to carry the synchronously (i.e. in a fixed pace) and regularly updated or constant profile features, while the event sequence can compactly catch event driven, more abrupt information, which can affect the condition intensity function over longer time period. More specifically, the highlights of this paper are:

1) We first make an observation that many conditional intensity functions can be viewed as an integration of two effects: i) spontaneous background component inherently affected by the internal (time-varying) attributes of the individual and the event type; ii) effects from history events. Meanwhile, most information in real world can also be covered by continuously updated features like age, temperature, and asynchronous event data such as clinical records, failures. This motivates us to devise a general approach.

2) Then we use an RNN whose units are aligned with the time points of a time series as its units, and an RNN whose units are aligned with events. The time series RNN can timely track the spontaneous background while the event sequence RNN is used to efficiently capture the long-range dependency over history with arbitrary time intervals. This allows to fit arbitrary dynamics of point process which otherwise will be very difficult or often impossible to be specified by a parameterized model under certain assumptions.

3) To our best knowledge, this is the first work to *fully* interpret and instantiate the conditional intensity function with fused time series and event sequence RNNs. This opens up the room for connecting the neural network techniques to

traditional point process that emphasizes more on specific model driven by domain knowledge. More importantly the introduction of a full RNN treatment lessen the efforts for the design of (semi-)parametric point process model and its complex learning algorithms which often call for special tricks that prohibiting the wide use for practitioners. In contrast, neural networks and specifically RNN, is becoming off-the-shelf tools and getting widely used recently.

4) Our model is simple, general and can be end-to-end trained. We target to a predictive maintenance problem. The data is from a global bank headquartered in North America, consisting decades of thousands of event logs for a large number of Automated Teller Machines (ATMs). The state-of-the-art performance on failure type and timestamp prediction corroborates its suitability to real-world applications.

Related Work and Motivation

We view the related concepts and work in this section, which is mainly focused on Recurrent Neural Networks (RNNs) and their applications in time series and sequences data, respectively. Then we give our point of view on existing point process methods and their connection to RNNs. All these observations indeed motivate the work of this paper.

Recurrent neural network The building blocks of our model are the Recurrent Neural Networks (RNNs) (Elman 1990; Pascanu, Mikolov, and Bengio 2013) and its modern variant Long Short-Term Memory (LSTM) units (Hochreiter and Schmidhuber 1997; Graves 2013). RNNs are dynamical systems whose next state and output depend on the present network state and input, which are more general models than the feed-forward networks. RNNs have long been explored in perceptual applications for many decades, however it can be very difficult to train RNNs to learn long-range dynamics perhaps in part due to the vanishing and exploding gradients problem. LSTMs provide a solution by incorporating memory units that allow the network to learn when to forget previous hidden states and when to update hidden states given new information. Recently, RNNs and LSTMs have been successfully applied in large-scale vision (Gregor et al. 2015), speech (Graves, rahman Mohamed, and Hinton 2014) and language (Sutskever, Vinyals, and Le. 2014) problems.

RNNs for series data From application perspective, we view RNNs works by two scenarios as particularly considered in this paper: i) RNNs for synchronized series with evenly spaced interval e.g. time series or indexed sequence with pure order information e.g. language; ii) asynchronous sequence with random timestamp e.g. event data.

i) *Synchronized series*: RNNs have been a long time a natural tool for standard time series modeling and prediction (Connor, Martin, and Atlas 1994; Han et al. 2004; Chandra and Zhang 2012; Chen, Chang, and Chang 2013), whereby the indexed series data point is fed as input to an (unfold) RNN. In a broader sense, video frames can also be treated as time series and RNN are widely used in recent visual analytics works (Jain et al. 2016; Tripathi et al. 2016) and so for speech (Graves, rahman Mohamed, and Hinton 2014). RNNs are also intensively adopted for sequence modeling tasks (Chung et al. 2014; Bengio et al. 2015) when only order information is considered.

ii) *Asynchronous event*: In contrast, event sequence with timestamp about their occurrence, which are asynchronously and randomly distributed over the continuous time space, is another typical input type for RNNs (Du et al. 2016; Choi et al. 2016; Esteban et al. 2016) and (Che et al. 2016) (despite its title for ‘time series’). One key differentiation against the first scenario is that the timestamp or time duration between events (together with other features) is taken as input to the RNNs. By doing so, (long-range) event dependency can be effectively encoded.

Point process Point process has been a principled framework for modeling event data (Aalen, Borgan, and Gjessing 2008). The dynamics of the point process can be well captured by its conditional intensity function whose definition is briefly reviewed here: for a short time window $[t, t + dt)$, $\lambda(t)$ represents the rate for the occurrence of a new event conditioned on the history $\mathcal{H}_t = \{z_i, t_i | t_i < t\}$:

$$\lambda(t) = \lim_{\Delta t \rightarrow 0} \frac{\mathbb{E}(N(t + \Delta t) - N(t) | \mathcal{H}_t)}{\Delta t} = \frac{\mathbb{E}(dN(t) | \mathcal{H}_t)}{dt}$$

where $\mathbb{E}(dN(t) | \mathcal{H}_t)$ is the expectation of the number of events happened in the interval $(t, t + dt]$ given the historical observations \mathcal{H}_t . The conditional intensity function has played a central role in point processes and many popular processes vary on how it is parameterized.

1) *Poisson process* (Kingman 1992): the homogeneous Poisson process has a very simple form for its intensity function: $\lambda(t) = \lambda_0$. Poisson process and its time-varying generalization are both assumed to be independent of the history.

2) *Reinforced poisson processes* (Pemantle 2007; Shen et al. 2014): the model captures the ‘rich-get-richer’ mechanism characterized by a compact intensity function, which is recently used for popularity prediction (Shen et al. 2014).

3) *Hawkes process* (Hawkes 1971): Hawkes process has received wide attention recently in social network analysis (Zhou, Zha, and Song 2013a), viral diffusion (Yang and Zha 2013) and criminology (Lewis et al. 2010) etc. It explicitly uses a triggering term to model the excitation effect from history events and is originally motivated to analyze the earthquake and its aftershocks (Ogata 1988).

4) *Reactive point process* (Ertekin, Rudin, and McCormick 2015): it can be regarded as a generalization for the Hawkes process by adding a self-inhibiting term to account for the inhibiting effects from history events.

5) *Self-correcting process* (Isham and Westcott 1979): its background part increases steadily, while it is decreased by a constant $e^{-\alpha} < 1$ every time a new event appears.

We reformulate these intensity functions in their general form in Table 1. It tries to separate the spontaneous background component and history event effect explicitly.

Predictive maintenance Predictive maintenance (Mobley 2002) is a sound testbed for our model which refers to a practice that involves equipment risk prediction to allow for proactive scheduling of corrective maintenance. Such an early identification of potential concerns helps deploy limited resources more cost effectively, reduce operations costs and maximize equipment uptime (Grall et al. 2002). Predictive maintenance is adopted in a wide variety of applications such as fire inspection (Madaio et al. 2016), data cen-

Table 1: Conditional intensity functions of point processes.

Model	Background	History event effect
Poisson process	$\mu(t)$	0
Reinforced poisson process	0	$\gamma(t) \sum_{t_i < t} \delta(t_i < t)$
Hawkes process	$\mu(t)$	$\sum_{t_i < t} \gamma(t, t_i)$
Reactive point process	$\mu(t)$	$\sum_{t_i < t} \gamma_1(t, t_i) - \sum_{t_i < t} \gamma_2(t, t_i)$
Self-correcting process	0	$\exp(\mu t - \sum_{t_i < t} \gamma(t, t_i))$

Note: $\delta(t)$ is Dirac function, $\gamma(t, t_i)$ is time-decaying kernel and $\mu(t)$ can be constant or time-varying function.

ter (Sirbu and Babaoglu 2015) and electrical grid (Ertekin, Rudin, and McCormick 2015) management. For its practical importance in different scenarios and relative rich event data for modeling, we target our model to a real-world dataset of more than 1,000 automated teller machines (ATMs) from a global bank headquartered in North America.

Network Structure and End-to-End Learning

Taking a sequence $\{\mathbf{x}\}_{t=1}^T$ as input, the RNN generates the hidden states $\{\mathbf{h}\}_{t=1}^T$ and outputs a sequence (Elman 1990; Pascanu, Mikolov, and Bengio 2013). Specifically, we implement our RNN with Long Short Term Memory (LSTM) (Hochreiter and Schmidhuber 1997; Graves 2013) for its popularity and well-known capability for efficient long-range dependency learning. In fact other RNN variant e.g. Gated Recurrent Units (GRU) (Chung et al. 2014) can also be alternative choice. We reiterate the formulation of LSTM:

$$\begin{aligned} \mathbf{i}_t &= \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{V}_i \mathbf{c}_{t-1} + \mathbf{b}_i), \\ \mathbf{f}_t &= \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{V}_f \mathbf{c}_{t-1} + \mathbf{b}_f), \\ \mathbf{c}_t &= \mathbf{f}_t \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_c \mathbf{x}_t + \mathbf{U}_c \mathbf{h}_{t-1} + \mathbf{b}_c), \\ \mathbf{o}_t &= \sigma(\mathbf{W}_o \mathbf{x}_t + \mathbf{U}_o \mathbf{h}_{t-1} + \mathbf{V}_o \mathbf{c}_t + \mathbf{b}_o), \\ \mathbf{h}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \end{aligned}$$

where \odot denotes element-wise multiplication and the recurrent activation σ is the Logistic Sigmoid function. The above system can be reduced into an LSTM equation:

$$(\mathbf{h}_t, \mathbf{c}_t) = \text{LSTM}(\mathbf{x}_t, \mathbf{h}_{t-1} + \mathbf{c}_{t-1})$$

We consider two types of input: i) continuously and evenly distributed time-series data e.g. temperature; ii) event data whose occurrence time interval is random. The network is comprised by two RNNs using evenly spaced time series $\{y_t\}_{t=1}^T$ to model the background intensity of events occurrence and event sequence $\{z_i, t_i\}_{i=1}^N$ to capture long-range event dependency. As a result, we have:

$$(\mathbf{h}_t^y, \mathbf{c}_t^y) = \text{LSTM}_y(\mathbf{y}_t, \mathbf{h}_{t-1}^y + \mathbf{c}_{t-1}^y), \quad (1)$$

$$(\mathbf{h}_t^z, \mathbf{c}_t^z) = \text{LSTM}_z(\mathbf{z}_t, \mathbf{h}_{t-1}^z + \mathbf{c}_{t-1}^z), \quad (2)$$

$$\mathbf{e}_t = \tanh(\mathbf{W}_f [\mathbf{h}_t^y, \mathbf{h}_t^z] + \mathbf{b}_f), \quad (3)$$

$$\mathbf{U}_t = \text{softMax}(\mathbf{W}_U \mathbf{e}_t + \mathbf{b}_U), \quad (4)$$

$$\mathbf{u}_t = \text{softMax}(\mathbf{W}_u [\mathbf{e}_t, \mathbf{U}_t] + \mathbf{b}_u) \quad (5)$$

$$s_t = \mathbf{W}_s \mathbf{e}_t + b_s, \quad (6)$$

where U and u denotes the main type and subtype of events respectively. s is the timestamp associated with each event.

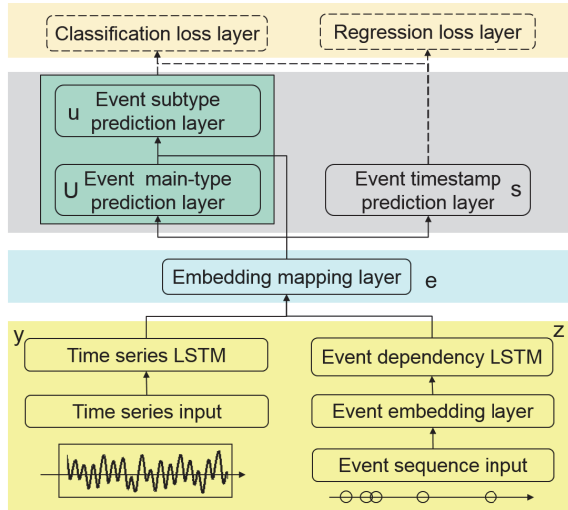


Figure 2: Our network can be trained end-to-end. Time series and event sequence are fed into two RNNs (LSTM) which are connected to an embedding mapping layer that fuses the information from two LSTMs. Then three prediction layers are used to output the predicted main type, subtype of events, and the associated timestamp. Cross-entropy with the time penalty loss by Eq.7 and square loss are respectively used for event type and timestamp prediction.

The total loss is the sum of the time prediction loss and the cross-entropy loss for event type:

$$\sum_{j=1}^N \left(-W_U^j \log(U_t^j) - w_u^j \log(u_t^j) - \log \left(f(s_t^j | h_{t-1}^j) \right) \right) \quad (7)$$

where N is the number of training samples indexed by j , and s_t^j is the timestamp for the coming event, while h_{t-1}^j is the history information. The underlying rationale for the third term is that we not only encourage correct classification of the coming event type, but also reinforce the corresponding timestamp of the event shall be close to the ground truth. We adopt a Gaussian penalty function with a fixed $\sigma^2 = 10$:

$$f(s_t^j | h_{t-1}^j) = \frac{1}{\sqrt{2\pi\sigma}} \exp \left(-\frac{(s_t^j - \tilde{s}_t^j)^2}{2\sigma^2} \right)$$

The output \tilde{s}_t^j from the timestamp prediction layer is fed to the classification loss layer to compute the above penalty given the actual timestamp s_t^j for sample i .

Following the importance weighting methodology for skewed data of model training (Rosenberg 2012), the weight parameters W, w for both main-type and subtype are set as the inverse of the sample number ratio in that type against the total size of samples, in order to weight more on those classes with fewer training samples. For the loss of independent main-type or subtype prediction as shown in Fig.3(b), we set the weight parameter w and W to zero respectively.

We adopt RMSprop gradients (Dauphin et al. 2015) which have been shown to work well on training deep networks to learn these parameters.

Table 2: Statistics of main/sub-type of event count per ATM, and timestamp interval in days for all ATMs (in brackets).

data	type	total	max	min	mean	std
Training set	Ticket	2226(-)	10(137.04)	0(1.21)	2.09(31.70)	1.85(25.14)
	Error	28434(-)	168(153.90)	0(0.10)	26.70(6.31)	18.38(9.74)
	PRT	9204(-)	88(210.13)	0(0.10)	8.64(12.12)	11.37(21.41)
	CNG	7767(-)	50(200.07)	0(0.10)	7.29(15.49)	6.59(23.87)
	IDC	4082(-)	116(206.61)	0(0.10)	3.83(23.85)	5.84(30.71)
	COMM	3371(-)	47(202.79)	0(0.10)	3.16(22.35)	3.90(29.36)
	LMTP	2525(-)	81(207.93)	0(0.10)	2.37(22.86)	4.41(34.56)
MISC	1485(-)	32(204.41)	0(0.10)	1.39(24.27)	2.54(34.38)	
Testing set	Ticket	1164(-)	15(148.00)	0(0.13)	2.52(26.30)	2.41(25.22)
	Error	11799(-)	104(193.75)	1(0.10)	25.59(6.47)	17.71(10.09)
	PRT	4089(-)	60(205.48)	0(0.10)	8.86(11.45)	10.77(20.44)
	CNG	3134(-)	66(196.93)	0(0.10)	6.79(15.90)	6.91(25.22)
	IDC	1645(-)	35(205.75)	0(0.10)	3.57(25.15)	4.29(31.62)
	COMM	1366(-)	53(205.17)	0(0.10)	2.96(22.88)	4.08(30.09)
	LMTP	939(-)	20(186.87)	0(0.10)	2.04(26.63)	2.76(36.96)
MISC	626(-)	21(190.75)	0(0.10)	1.36(25.14)	2.59(34.91)	

Experiments on Real-world Data

We use failure prediction for predictive ATMs maintenance as a typical example of event based point process modeling. We have no prior knowledge on the dynamics of the complex system and the task can involve arbitrarily working schedules and heterogeneous mix of conditions. It takes much cost or even impractical to devise specialized models.

Problem and real data description

In maintenance support services, when a device fails, the equipment owner raises a maintenance service ticket and technician will be assigned to repair the failure. In fact, the history log and relevant profile information about the equipment can be indicative signals for the coming failures.

The studied dataset is comprised of the event logs involving error reporting and failure tickets, which is originally collected from a large number of ATMs owned by an anonymous global bank headquartered in North America. The bank is also a customer of the technical support service department of a Fortune 500 IT company.

ATM models The training data consists of 1085 ATMs and testing data has 469 ATMs, in total 1557 Wincor ATMs that cover 5 ATM machine models: ProCash 2100 RL (980, 430), 1500 RL (19, 5), 2100 FL (53, 21), 1500 FL (26, 10), and 2250XE RL (7, 3). The numbers in the bracket indicate the number of machines for training and testing.

Event type There are two main types ‘ticket’ and ‘error’ from Sep. 2014 to Mar. 2015. Statistics is presented in Table 2. Moreover ‘error’ is divided into 6 subtypes regarding in which component the error occurs: 1) printer (PRT), 2) cash dispenser module (CNG), 3) internet data center (IDC), 4) communication part (COMM), 5) printer monitor (LMTP), 6) miscellaneous e.g. hip card module, usb (MISC).

Features The input features for the two RNNs are: 1) *Time series RNN*: For each sub-window of length 7 days, for the time series RNN, we extract features including: i) the inventory information: ATM models, age, location, etc; ii) Event statistics, including tickets events from maintenance records, and errors from system log. Their occurrence frequencies are used as features. The concatenation of the above two categories of features serves as the features for

each sub-window i.e. time series point. **2) Event sequence RNN:** event type and the time interval between two events.

Model setting We use a single layer LSTM of size 32 with Sigmoid gate activations and tanh activation for hidden representation. The embedding layer is fully connected and it uses tanh activation and outputs a 16 dimensional vector. One-hot or embedding can be used for event type representation. For a large number of types, embedding representation is compact and efficient. For time series RNN, we set the length of each sub-window (i.e. the evenly spaced time interval) to be 7 days and the number of sub-window to be 5. In this way, our observation length is 35 days for time series. For event-dependency, the length of event sequence can be arbitrarily long. Here we take it by 7.

We also test degraded versions of our model as follows:

1) Time series RNN: the input is event sequence (the right half in the yellow part of Fig.2) is removed. Note this design is in spirit similar to many LSTM models (Jain et al. 2016; Tripathi et al. 2016) used for video analytics, whereby the frame sequence can be treated as time series as the input to LSTMs. **2) Event (sequence) RNN:** the RNN whose input is time series (the left half in the yellow part of Fig.2) is removed; **3) Intensity RNN:** two RNN are fused as shown in Fig.2. For the above three methods, the output layer is directly the fine-grained subtype of events with no hierarchical structure as shown in the top left part of Fig.2) in dark green. We also term three ‘hierarchical’ versions whose two hierarchical prediction layers in Fig.2 are used: **4) Time series hRNN, 5) Event (sequence) hRNN, 6) Intensity hRNN.**

In addition, we compare three major peer methods. For Logistic model, the input are the concatenation of feature vectors for all active time series RNN sub-windows (set to 5 in this paper). For RMTTP and Hawkes process, we train the model on the event sequences with associated information. In fact, RMTTP will further process the event data into the similar input information to our event RNN.

1) Logistic model: we use Logistic regression for event timestamp prediction and use another independent Logistic classification model for event type prediction.

2) Recurrent Marked Temporal Point Processes (RMTTP): (Du et al. 2016) uses neural network to model the event dependency flexibly. The method can only sample transient time series features when an event happens and use partially parametric form for the base intensity.

3) Hawkes Process: To enable multi-type event prediction, we use a Multi-dimensional Hawkes process. Similar to (Zhou, Zha, and Song 2013a), we also add a sparsity regularization term on the mutual infection matrix but the low-rank assumption is removed as we only have 6 subtypes.

Evaluation metrics We use several popular prediction metrics for performance evaluation. For the coming event type prediction, we adopt *Precision*, *Recall*, *F1 Score* and *Confusion matrix* over 2 main types (‘error’, ‘ticket’) as well as *Confusion matrix* over 6 subtypes under ‘error’. Note all these metrics are computed for each type, and then are averaged over all types. For event time prediction, we use the *Mean Absolute Error (MAE)* which measures the absolute difference between the predicted time point and the actual one. These settings are similar to (Du et al. 2016).




Figure 3: Hierarchical layer and Flat independent layer.

Table 3: Ablation test of our method and peer methods i.e. multi-dimensional Hawkes process, recurrent Hawkes process and Logistic classification (for type) and regression (for event timestamp). Numbers are averaged over types.

	model	main-type	subtype	hierarchical output main-type	output subtype
precision	Time series RNN	0.673	0.554	0.582	0.590
	Event sequence RNN	0.671	0.570	0.623	0.614
	Intensity RNN	0.714	0.620	0.642	0.664
	Hawkes process	0.457	0.387	—	—
	Logistic prediction	0.883	0.385	—	—
	RMTTP	0.581	0.574	—	—
recall	Time series RNN	0.853	0.522	0.738	0.608
	Event sequence RNN	0.821	0.543	0.770	0.621
	Intensity RNN	0.905	0.614	0.805	0.661
	Hawkes process	0.493	0.394	—	—
	Logistic prediction	0.795	0.273	—	—
	RMTTP	0.691	0.583	—	—
F1 score	Time series RNN	0.707	0.533	0.571	0.605
	Event sequence RNN	0.703	0.555	0.651	0.610
	Intensity RNN	0.765	0.616	0.662	0.663
	Hawkes process	0.473	0.386	—	—
	Logistic prediction	0.832	0.269	—	—
	RMTTP	0.584	0.572	—	—
MAE (in days)	Time series RNN	4.37	4.48	4.26	4.41
	Event sequence RNN	4.24	4.42	4.21	4.37
	Intensity RNN	4.13	4.20	4.02	4.13
	Hawkes process	5.26	5.46	—	—
	Logistic prediction	4.52	4.61	—	—
	RMTTP	4.28	4.32	—	—
F1 score+	Time series RNN	0.768	0.547	0.572	0.603
	Event sequence RNN	0.705	0.597	0.639	0.646
	Intensity RNN	0.825	0.661	0.684	0.708
	Hawkes process	0.467	0.451	—	—
	Logistic prediction	0.846	0.286	—	—
	RMTTP	0.584	0.619	—	—
MAE+	Time series RNN	4.21	3.78	4.05	3.97
	Event sequence RNN	4.16	3.84	4.12	4.01
	Intensity RNN	4.12	3.57	4.21	4.11
	Hawkes process	5.42	3.93	—	—
	Logistic prediction	4.5	4.24	—	—
	RMTTP	4.26	3.99	—	—

To evaluate the type and timestamp prediction jointly, we devise two more strict metrics. For type prediction, we narrow down to test the samples whose timestamp prediction error $MAE < 3$ days and we compute the new $F1\ score+$. For timestamp, we recompute the new $MAE+$ only for the samples whose coming event is correctly predicted.

Platform The code is based on Theano running on a Linux server with 32G memory, 2 CPUs with 6 cores for each: Intel(R) Xeon(R) CPU E5-2603 v3 @ 1.60GHz. We also use 4 GPU: GeForce GTX TITAN X for acculturation.

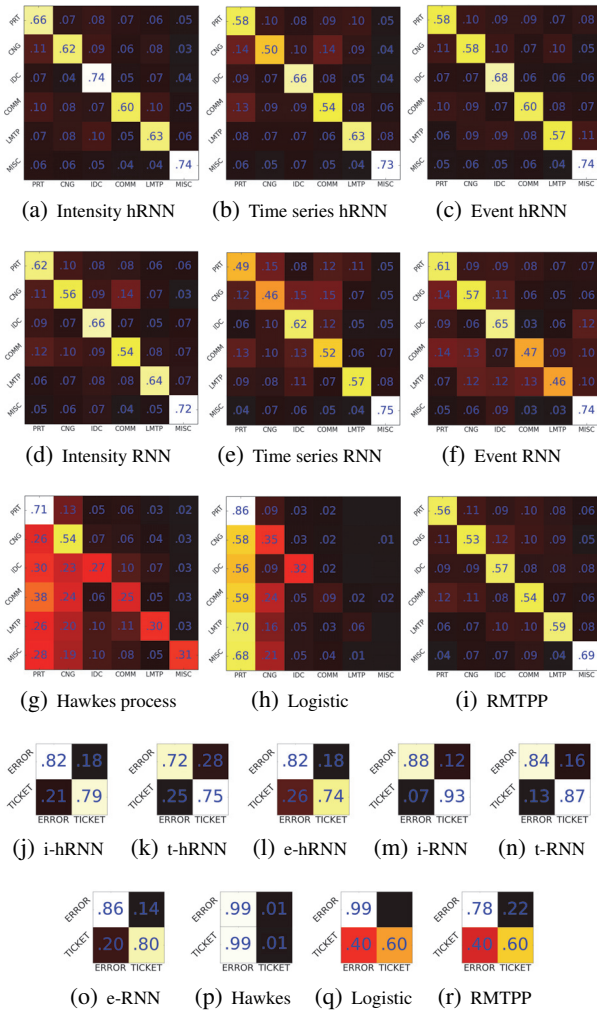


Figure 4: Confusion matrixes for sub/main type. Three methods in top and middle row use the hierarchical and flat structure as in Fig.3 respectively. Zoom in for better view.

Results and discussion

All evaluations are performed on the testing dataset distinctive to training set whose statistics are shown in Table 2.

Averaged performance Table 3 shows the averaged performance among various types of events. As shown in Fig.3, we test two architectures of the event type prediction layer, i.e. hierarchical predictor (Fig.3(a)) and flat independent predictors (Fig.3(b)). The main type includes ‘ticket’ and ‘error’ and the subtype include ‘ticket’ and the other six subtypes under ‘error’ as we describe earlier in the paper.

Confusion matrix The confusion matrix for the six subtypes under ‘error’ event, as well as for the two main types ‘ticket’ and ‘error’ are shown in Fig.4 by various methods.

We make observations and analysis based on the results:

1) As shown in 3, for main-type, the flat architecture that directly predicts the main types outperforms the hierarchical one in different settings of the input RNN as well as varying evaluation metrics. This can be explained that the loss func-




Figure 5: The evolving of point process modeling.

tion focuses on the main-type misclassification only. While for the subtype prediction, the hierarchical layer performs better since it fuses the output from the main-type prediction layer and the embedding layer as shown in Fig.3(a).

2) No surprisingly, for both event type and timestamp prediction, our main approach, i.e. intensity RNN that fuses two RNNs outperforms its counterparts time series RNN and event sequence RNN by a notable margin. While the event RNN also often performs better than the time series counterpart. This suggests at least in the studied dataset, history event effects are important for the future event occurrence.

3) Our main method intensity RNN is almost always superior against other methods except for the main-type prediction task, whereby the Logistic classification model performs better. However for more challenging tasks i.e. subtype prediction and event timestamp prediction, our method significantly outperforms especially for subtype prediction task. Interestingly, all point process based models obtain better results on this task which suggests the point process models are more promising compared with classical classification models. Indeed, our methodology provides an end-to-end learning mechanism without any pre-assumption in modeling point process. All these empirical results on real-world tasks suggest the efficacy of our approach.

Conclusion

We use Fig.5 to conclude and further position our model in the development of (implicit and explicit) modeling the intensity function of point process. In fact, Hawkes process uses a full explicit parametric model and RMTTP misses the dense time series features to model time-varying base intensity and assumes a partially parametric form for it. We make a further step by a full implicit mapping model. Our model (see Fig.2) is simple, general and can be learned end-to-end with standard backpropagation and opens up new possibilities for borrowing the advances in neural network learning to the area of point process modeling and applications. The representative study in this paper has clearly suggests its high potential to real-world problems, even we have no domain knowledge on the problem at hand. This is in contrast to existing point process models where an assumption about the dynamics is often need to be specified beforehand.

References

- Aalen, O.; Borgan, O.; and Gjessing, H. 2008. *Survival and event history analysis: a process point of view*. Springer Science & Business Media.
- Bengio, S.; Vinyals, O.; Jaitly, N.; and Shazeer, N. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*.

- Chandra, R., and Zhang, M. 2012. Cooperative coevolution of el-man recurrent neural networks for chaotic time series prediction. *Neurocomputing* 86:116–123.
- Che, Z.; Purushotham, S.; Cho, K.; Sontag, D.; and Liu, Y. 2016. Recurrent neural networks for multivariate time series with missing values. In *arXiv:1606.01865*.
- Chen, P. A.; Chang, L. C.; and Chang, F. J. 2013. Reinforced recurrent neural networks for multi-step-ahead flood forecasts. *Journal of Hydrology* 497:71–79.
- Choi, E.; Bahadori, M. T.; Schuetz, A.; Stewart, W. F.; and Sun, J. 2016. Doctor ai: Predicting clinical events via recurrent neural networks. In *arXiv:1511.05942v10*.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *arXiv:1412.3555*.
- Connor, J. T.; Martin, R. D.; and Atlas, L. E. 1994. Recurrent neural networks and robust time series prediction. *IEEE Transactions on Neural Networks* 5(2):240–254.
- Dauphin, Y. N.; de Vries, H.; Chung, J.; and Bengio, Y. 2015. Rmsprop and equilibrated adaptive learning rates for non-convex optimization. In *arXiv:1502.04390*.
- 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 *KDD*.
- Elman, J. L. 1990. Finding structure in time. *Cognitive Science* 14:179–211.
- Ertekin, S.; Rudin, C.; and McCormick, T. H. 2015. Reactive point processes: A new approach to predicting power failures in underground electrical systems. *The Annals of Applied Statistics* 9(1):122–144.
- Esteban, C.; Staeck, O.; Yang, Y.; and Tresp, V. 2016. Predicting clinical events by combining static and dynamic information using recurrent neural networks. In *arXiv:1602.02685*.
- Grall, A.; Dieulle, L.; Berenguer, C.; and Roussignol, M. 2002. Continuous-time predictive-maintenance scheduling for a deteriorating system. *IEEE transactions on Reliability* 51(2):141–150.
- Graves, A.; rahman Mohamed, A.; and Hinton, G. 2014. Towards end-to-end speech recognition with recurrent neural networks. In *ICML*.
- Graves, A. 2013. Generating sequences with recurrent neural networks. In *arXiv:1308.0850*.
- Gregor, K.; Danihelka, I.; Graves, A.; Rezende, D.; and Wierstra, D. 2015. Draw: A recurrent neural network for image generation. In *ICML*.
- Han, M.; Xi, J.; Xu, S.; and Yin, F.-L. 2004. Prediction of chaotic time series based on the recurrent predictor neural network. *IEEE Transactions on Signal Processing* 52(12):3409–3416.
- Hawkes, A. G. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Isham, V., and Westcott, M. 1979. A self-correcting pint process. *Advances in Applied Probability* 37:629–646.
- Jain, A.; Singh, A.; Koppula, H. S.; Soh, S.; and Saxena, A. 2016. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In *ICRA*.
- Kingman, J. F. C. 1992. *Poisson processes*. volume 3. Oxford university press.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324.
- Lewis, E., and Mohler, E. 2011. A nonparametric em algorithm for multiscale hawkes processes. *Journal of Nonparametric Statistics*.
- Lewis, E.; Mohler, G.; Brantingham, P. J.; and Bertozzi, A. 2010. Self-exciting point process models of insurgency in iraq. *UCLA CAM Reports* 10 38.
- Madaio, M.; Chen, S.; Haimson, O.; Zhang, W.; Cheng, X.; and Hinds-Aldrich, M. 2016. Firebird: Predicting fire risk and prioritizing fire inspections in atlanta. In *KDD*.
- Mobley, R. K. 2002. *An introduction to predictive maintenance*. Butterworth-Heinemann.
- Montgomery, D. C.; Jennings, C. L.; and Kulahci, M. 2015. *Introduction to time series analysis and forecasting*. John Wiley & Sons.
- Ogata, Y. 1988. Statistical models for earthquake occurrences and residual analysis for point processes. *J. Amer. Statist. Assoc.* 83(401):9–27.
- Pascanu, R.; Mikolov, T.; and Bengio, Y. 2013. On the difficulty of training recurrent neural networks. In *ICML*.
- Pemantle, R. 2007. A survey of random processes with reinforcement. *Probability Survey* 4(0):1–79.
- Rosenberg, A. 2012. Classifying skewed data: Importance weighting to optimize average recall. In *INTERSPEECH*, 2242–2245.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115(3):211–252.
- Shen, H.; Wang, D.; Song, C.; and Barabási, A. 2014. Modeling and predicting popularity dynamics via reinforced poisson processes. In *AAAI*.
- Sirbu, A., and Babaoglu, O. 2015. A holistic approach to log data analysis in high-performance computing systems: The case of ibm blue gene/q. *Euro-Par 2015: Parallel Processing Workshops* 631–643.
- Snyder, D. L., and Miller, M. I. 2012. *Random point processes in time and space*. Springer Science & Business Media.
- Sutskever, I.; Vinyals, O.; and Le., Q. V. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Tripathi, S.; Lipton, Z. C.; Belongie, S.; and Nguyen, T. 2016. Context matters: Refining object detection in video with recurrent neural networks. In *BMVC*.
- Xu, H.; Wu, W.; Nemati, S.; and Zha, H. 2016. Patient flow prediction via discriminative learning of mutually-correcting processes. *IEEE transactions on Knowledge and Data Engineering*.
- Yang, S.-h., and Zha, H. 2013. Mixture of mutually exciting processes for viral diffusion. In *ICML*.
- Zhou, K.; Zha, H.; and Song, L. 2013a. Learning social infectivity in sparse low-rank networks using multi-dimensional hawkes processe. In *AISTATS*.
- Zhou, K.; Zha, H.; and Song, L. 2013b. Learning triggering kernels for multi-dimensional hawkes processes. In *ICML*.