




# Identification and prediction of attacks to industrial control systems using temporal point processes

Giancarlo Fortino<sup>1</sup> · Claudia Greco<sup>1</sup> · Antonella Guzzo<sup>1</sup> · Michele Ianni<sup>1</sup> 

Received: 17 February 2022 / Accepted: 10 September 2022 / Published online: 23 September 2022  
© The Author(s) 2022

## Abstract

The task of identifying malicious activities in logs and predicting threats is crucial nowadays in industrial sector. In this paper, we focus on the identification of past malicious activities and in the prediction of future threats by proposing a novel technique based on the combination of Marked Temporal Point Processes (*MTTP*) and Neural Networks. Differently from the traditional formulation of Temporal Point Processes, our method does not make any prior assumptions on the functional form of the conditional intensity function and on the distribution of the events. Our approach is based the adoption of Neural Networks with the goal of improving the capabilities of learning arbitrary and unknown event distributions by taking advantage of the Deep Learning theory. We conduct a series of experiments using industrial data coming from gas pipelines, showing that our framework is able to represent in a convenient way the information gathered from the logs and predict future menaces in an unsupervised way, as well as classifying the past ones. The results of the experimental evaluation, showing outstanding values for precision and recall, confirm the effectiveness of our approach.

**Keywords** Marked temporal point processes · Security of industrial control systems · Anomaly detection · Threat prediction

## 1 Introduction

In recent years we have witnessed a rising paradigm shift in the Industrial Control Systems (*ICS*) landscape, since an ever-increasing number of technologies and devices are migrating from a traditional mechanical or electrotechnical often closed system, landing to integrated modern control systems. The new generation of ICS is based on a tight connection among components, made possible by the spread of both commercial off-the-shelf (*COTS*) products and open protocols, as well as by the price reduction of the hardware components needed for the development of these

architectures. However, if on the one hand, the interconnection and the remote availability of ICS led to a reduction of the costs for industrial operators and to an increased efficiency, on the other hand, the growing complexity is coupled with an increasing attack surface (Manadhata and Wing 2010) and on a change in the profile of the attacker, who is no more necessarily an insider, but can be an external actor exploiting the remote functionalities of the system. The rising damage due to cyberattacks hit 6 trillion dollars in 2021 (Morgan 2020), doubling the amount recorded in 2016. The need for efficient solutions that can help to protect sensitive information and devices is more topical than ever and led to concerns in the security community due to the potential impact on safety if the attacks are conducted against critical infrastructures (Johnson 2012; Kornecki and Zalewski 2010; Aven 2007).

The high number of software components generate an endless stream of data of a different nature, ranging from network captures to application logs, from hardware monitoring logs to user interaction, and so on. The analysis of the generated data represents a task of primary importance because it often allows to thwart attacks even before they happen (Babbin 2006). However, for this to occur, we cannot

---

✉ Michele Ianni  
michele.ianni@unical.it

Giancarlo Fortino  
giancarlo.fortino@unical.it

Claudia Greco  
claudia.greco@dimes.unical.it

Antonella Guzzo  
antonella.guzzo@unical.it

<sup>1</sup> DIMES, University of Calabria, Rende, Italy

rely only on the ability of system administrators and domain experts of noticing relevant actions in huge software logs, but we need automated techniques able to analyze in a fast and reliable way all the information available (Jang-Jaccard and Nepal 2014; Ianni and Masciari 2022; Zikopoulos et al. 2011).

Critical infrastructures represent nowadays an extremely attractive target for criminals, as demonstrated by recent attacks (e.g. Stuxnet and Flame) and talks or academic papers (Apa and Penagos 2013; Meixell and Forner 2013). The attacks target various kinds of ICS or SCADA (Supervisory Control And Data Acquisition) systems, including industrial sectors such as automotive (Koscher et al. 2010; Checkoway et al. 2011), aerospace (Bieber et al. 2012; Cockram and Lautieri 2007), electricity (Lee and Brewer 2009; Gumaï et al. 2020), oil and gas (Grøtan et al. 2007; Johnsen 2012), to cite a few. Due to the peculiar nature of the systems, it is straightforward to note that traditional approaches based on blacklisting (or whitelisting) activities are likely to fail, because, even if we are able to enumerate all malicious activities, we may incur in attacks that are the combination of several different actions that, when analyzed individually, are not considered menaces at all. There is a lot of information that can be inferred by the correct analysis of a log, and this knowledge can help to provide a roadmap to the origin of a specific threat, to identify the agents involved and even to predict future unauthorized behaviors, both from inside attackers and external ones (Schultz 2002; Kent and Soupaya 2006).

In this paper, we focus on the analysis of logs related to industrial systems from a security viewpoint. Our goal is to provide a novel approach to label malicious activity in logs and even predict potential future attacks. Our approach is based on *Marked Temporal Point Processes* (MTTPs), a probabilistic stochastic framework which showed its effectiveness in different domains (Yan et al. 2019), such as earthquake prediction, aftershocks, healthcare, financial trends, activity daily living prediction (Fortino et al. 2020, 2021) and so on, but that has found less attention than it deserves in security oriented scenarios. We highlight that the technique proposed in this paper can be applied to a wider scientific area, for both the prediction and the classification of what kind of event will take place at what time in the future. Our approach has the following contributions:

- A rigorous formalization of MTTPs to classify and predict malicious actions, by interpreting the actions described by different kinds of logs as discrete stochastic processes of asynchronous events.
- The dynamics of the processes do not follow a predetermined process model, but it is learned by the use of *Neural Networks* (NN), thus providing an unsupervised framework that is able to learn from logs.

- An experimental activity, carried out on a dataset related to gas pipelines shows the effectiveness of the approach.

The rest of the paper is organized as follows: Sect. 2 discusses the relevant literature in the field of malicious activity detection and prediction, Sect. 3 is an introduction to the theoretical foundations of MTTP and to two different techniques based on NNs. In Sect. 5 we describe the proposed approach and evaluate it. Lastly, in Sect. 6 we draw our conclusions and shed light on the perspectives of our future work.

## 2 Related work

The task of identifying relevant information which does not adhere to an expected behavior is referred to as anomaly (or outlier) detection and it has been studied, in statistics communities, since the 19th century (Edgeworth 1887). The problem gained attention with the passing of time and it has been tackled with different techniques and applied to a plethora of different domains, ranging from novelty detection (Markou and Singh 2003a, b) to noise removal (Chen et al. 1990). One of the domains where anomaly detection has been more useful is, without any doubt, cybersecurity, with application to intrusion detection (Kumar 2005; Guzzo et al. 2020), fraud identification (Aleskerov et al. 1997), assisted surveillance (even military surveillance), vulnerability discovery, exploitation of software flaws, and so on. The literature about anomaly detection is wide since it includes a great number of techniques and application domains. Historically, among these techniques, clustering (Jain and Dubes 1988; Ianni et al. 2020) is of particular importance. At first sight, clustering and anomaly detection seem to be poles apart, but looking more in detail, they are tightly connected, since most of the clustering-based anomaly detection algorithms are based on the assumption that malicious activities represent a small amount of data different by “normal” activities, thus not belonging to any cluster. Examples of algorithms based on this assumption are DBSCAN (Ester et al. 1996), ROCK (Guha et al. 2000), and SNN clustering (Ertoz et al. 2002). Another interesting assumption is that while “normal” data are close to the cluster centroids, anomalies are distant from them. The distance itself is used as an anomaly score (Smith et al. 2002; Kohonen 1990; Ramadas et al. 2003; Emamian et al. 2000; Brockett et al. 1998; Barbará et al. 2003).

Clustering-based approaches, however, are likely to fail when the anomalies in the data form clusters by themselves. To overcome this issue, many algorithms based on the density of the clusters have been proposed (Pires and Santos-Pereira 2005; Otey et al. 2003; Mahoney and Chan 2003; Jiang et al. 2001).

Apart from clustering, many other approaches, based on artificial intelligence, have been extensively adopted to identify malicious patterns in data logs and predict future menaces. Among them, in the early stages of research, fuzzy logic and genetic algorithms played an important role. In Dilek et al. (2015) the authors investigate many different approaches used for detecting attacks. The use of artificial immune systems (AIS) in the protection of IoT environments has been explored in Greensmith (2015). With the spread of Machine Learning and Deep Learning based techniques we have witnessed a change of course in the landscape of the technical paradigms proposed for attack identification and prediction. An extensive study of many Deep Learning based approaches (LSTM, RNN, CNN) for anomaly detection has been proposed in Ahmad et al. (2021). With specific reference to the IoT scenario in Tahsien et al. (2020) the authors list a series of Machine Learning based approaches to attack identification, whereas in Alsoufi et al. (2021) the attention is focused on Deep Learning based techniques. Interesting analyses on both security issues and protocols devoted to vulnerabilities mitigation can be found in Mahbub (2020), Frustaci et al. (2018) and Djenna et al. (2021).

In this work, we advocate the use of the mathematical framework of Temporal Point Processes (TPP) in order to identify and predict malicious activities in Industrial Control Systems. TPPs proved their effectiveness in modeling complex behaviours related to asynchronous events, even when the inter-event time does not follow a predetermined model. TPPs have already been successfully applied in many different application scenarios (Yan et al. 2019), with particular reference to the task of predicting activities, ranging from health-care analysis to activity daily living, from finance to earthquakes and aftershocks modeling, to cite a few. Traditional TPP are based on the definition of a conditional intensity function that can be used to describe the behaviour of the sequences under consideration. This is usually done by applying well known processes, such as Poisson (Kingman 1993), Hawkes (Hawkes 1971), Self-correcting (Isham and Westcott 1979) and Self-exciting (Hawkes and Oakes 1974). A recent shift of perspective (Yan 2019; Yan et al. 2019) suggests to avoid the choice of an a-priori intensity function, but rather learning it from the analysis of the event data related to past executions. Our work is based on this recent perspective, since we are proposing the use of two different Neural Network based variations of the traditional TPPs in order to dynamically learn the dynamics behind arbitrary and unknown event distributions. The first work proposing the use of Recurrent Neural Networks to jointly model the event information and the occurring times in order to learn a representation of the nonlinear dependencies is RMTTP (Du et al. 2016). The predictive capabilities of the RMTTP framework have been shown to outperform traditional parametric techniques in different scenarios, such

as trajectory prediction, financial analysis and evolution of diseases. However, it is worth noticing that the RMTTP framework is generic and can be adopted with a very wide spectrum of application domains. Another milestone in the application of Neural Networks to TPP has been reached with the proposal of ERPP (Xiao et al. 2017b), a new model where two different LSTMs are used to independently model event and time sequences and an embedding mapping layer is used to merge the information coming from both the RNNs. The reason behind the choice of managing sequences independently can be found on the assumption that, usually, time series are likely to be represented by regular patterns, whereas the event sequences are often characterized by abrupt occurrences, thus affecting, especially over long period of times, the shape of the conditional intensity function. Xiao et al. (2017a) presented an extension of Xiao et al. (2017b), adopting a multi-dimensional point process model where for each type of event a different intensity function is generated. The authors also added an attention mechanism to the network architecture, making the prediction model and the relational mining among the event dimensions easier to read. To our best knowledge, there are no other works in literature proposing TPP based techniques for the identification and prediction of attacks to Industrial Control Systems. In this work we present a learning framework whose components employ architectures proposed by Du et al. (2016) and Xiao et al. (2017b). Our thinking is that the convenient representation of event streams, the knowledge inferred by attack information and the resulting assessment of the efficacy can be appealing to further research.

### 3 Marked temporal point process prediction

Marked Temporal Point Process (shortly MTPP) are generative models for sequences of events with variable length in continuous time. These models allow to make predictions, find anomalies and learn useful representations of the data, based on the history of past events. In this section, we provide a formal background: the theoretical basis upon which we can use MTPP for the identification and prediction of Attacks to Industrial Control Systems.

#### 3.1 Theoretical foundations

A Marked Temporal Point Process is a stochastic process modeling the occurrence of a sequence of events, each one belonging to a specific category and taking place in a given time. MTPP is an evolution of the traditional TPP, in the sense that MTPP extends the notion of event point, by adding a mark  $a_j$  which can embed domain related additional information. For example, in the earthquake prediction

scenario, we can use marks to hold some additional information with the goal of improving the accuracy of the prediction model when forecasting the position and the magnitude of a seismic event. In fact, the mark can embed not only the prediction label, but any other information of interest associated to the event. In our case we make use of marks in Temporal Point Processes in order to build a more realistic and accurate model by taking into account the available information about the attacks targeting an ICS.

More formally, a MTPP is a generative model for a sequence of events  $e_j = (t_j, a_j)$ , with  $t_j \in \mathcal{R}^+$ ,  $j \in \mathcal{Z}^+$  and  $a_j \in \mathcal{A}$  where the domain of  $\mathcal{A}$  is application dependent. We denote as history  $H_t$ , the list of events, up to time  $t$ . Our conditional density function is defined as the probability that the next event will happen in the infinitesimal interval  $[t, t + dt]$  conditioned by the past history  $H_t$ . Formally, we define  $f^*(t, a) = f(t, a|H_t)$  with mark  $a$ , where  $*$  from (Daley and Vere-Jones 2003) means that the density is conditional on the past (right up to but not including the present) rather than writing explicitly that the function depends on the history.

To define the full distribution of MTPP, we specify the density functions  $f((t_1, a_1), \dots, (t_j, a_j))$  one by one, starting in the past, and we can derive the distribution of all events times and maker pairs by the joint density:

$$f(\{(t_j, a_j)\}_{j=1}^n) = \prod_j f(t_j, a_j|H_{t_{j-1}}) = \prod_j f^*(t_j, a_j) \tag{1}$$

However, in place of using different conditional distributions, a different way of characterizing a marked temporal point process is based on the use of the conditional intensity function that could be specified by the conditional density  $f(t, a|H_{t_n})$  and its corresponding cumulative distribution function  $F(t, a|H_{t_n})$  for any  $t > t_n$ .

Formally, the conditional intensity function is defined by:

$$\lambda^*(t, a) = \frac{f(t, a|H_{t_n})}{1 - F(t, a|H_{t_n})} \tag{2}$$

The conditional intensity function can be interpreted heuristically by considering an infinitesimal interval around  $t$ , i.e.  $[t, t + dt]$  and the number of points falling in it,  $N$ :

$$\lambda^*(t, a)dt da = E[N(dt \times da)|H_t] \tag{3}$$

that is, the mean number of points in a small time interval  $dt$  with the mark in a small interval  $dt$ .

To suitably model the event times of a temporal point process, different formulations of conditional intensity function have been proposed, aiming at capturing the phenomena of interest that better characterize the application scenarios in the given domain (Yan et al. 2019)

Traditionally, the most common formulation, the *homogeneous Poisson process*, rely on the assumption that  $\lambda$  is

fixed over time and it is independent from the history of the events  $H_t$ . Many real life scenarios are well described by the homogeneous Poisson process, e.g. customers arriving at a store, arrival of phone calls in a call center, and so on.

Differently, in *non homogeneous Poisson processes* formulation,  $\lambda(t)$  is a function of time and does not depend on the history of past events. Examples of real life processes modeled by a non homogeneous Poisson process are the daily measurements of ozone gas or of the exposure to high levels of noise.

Another formulation adopts *Hawkes processes* (Hawkes 1971), where the conditional intensity function is exponential, thus simulating that the occurrence of an event point increases the probability for it to be followed, immediately after, by other points. This is the case of earthquakes, where a seismic event is likely to be followed, after a short period of time, by other similar events. Another example is the modeling of preferential attachment when analyzing interactions in social networks.

In *Self-correcting processes* (Isham and Westcott 1979) the probability of arrival of new points decreases after an event point appears, as in the case of terrorist attacks.

Over the past decade many different works adopted parametric forms of the intensity functions, with the purpose of embedding the historical sequence of events. However, several limitations emerged from these approaches. It is straightforward to note that a crucial role is played by the approximation of the conditional intensity function (*CIF*): in order to achieve good performances in the accuracy of the modeled process the choice of the CIF family is fundamental.

It is important to underline that the aforementioned conditional intensity function families require strong assumptions on the functional forms of the generative processes, and usually these assumptions may be initially unknown or, in most of the cases, may not correspond to real world scenarios. Based on these considerations, recent studies exploited the state of the art deep learning algorithms in order to model, in a flexible and efficient way, complex behaviors. Recurrent Neural Networks have been successfully applied to TPPs, and it has been proved to be more effective in defining a conditional intensity function, w.r.t. predetermined parametric forms.

## 4 Problem formulation and modeling approach

In this section we provide a modeling for the identification and the prediction of attacks to Industrial Control Systems using the MTPP framework. Differently from other application domains, for example earthquake forecasting, in which it is possible to make a priori assumptions on the distribution

of the events, in our case each a priori assumption could introduce a dangerous bias in the predictive model. We then advocate the idea of using Neural Networks in order to approximate the stochastic generation process. Furthermore, differently from the traditional unmarked approach, MTPPs introduce a further element of complexity (and accuracy) in the specification of the model. To identify the appropriate architecture for the application context, it was essential to focus on the following research questions:

- RQ1: which type of neural network is more suitable as prediction model?;
- RQ2: is it worthwhile to jointly learn mark and time? or does considering distributions separately lead to better results?

**Problem definition:** *Given a stream  $S$  of temporal events,  $S = \{e_1 = (t_1, a_1), \dots, e_m = (t_m, a_m)\}$ , each one represented by a pair of an attack label its corresponding timestamp and such that if two events  $e_i$  and  $e_j$  with  $i > j$ , we have that  $t_i \geq t_j$ . Considering the stream up to the  $j$ -th observation, our problem is to predict the next observation ( $j+1$ ), i.e. both the type of the attack and the timestamp in which the attack will take place.*

### 4.1 Choosing a model to support multivariate sequence prediction

In the MTPP formulation, the most common way to learn a general representation that approaches the unknown dependence structure on the events of the sign and time throughout history relies on the use of RNNs. Indeed, finite-dimensional recurrent neural networks with sigmoidal activation units can simulate an universal Turing machine, which is able to perform an extremely rich family of computations. To answer RQ1, in our model formulation, we use, in place of a standard RNN, a Long Short Term Memory (LSTM), since it is known to be able to learn complex nonlinear relationships and, in addition, to be able to better discover long range temporal relationships thanks to the use of memory cells. LSTM are therefore the most successful type of Recurrent Neural Network capable of directly supporting multivariate sequence prediction problems.

### 4.2 Recurrent marked TPP to model mark of an event and its occurring time

To answer RQ2, we chose to implement and compare two different approaches: RMTTP (Du et al. 2016) and ERPP (Xiao et al. 2017b). Both of them make use of RNNs to model and automatically learn the conditional intensity function  $\lambda$ , without any particular prior assumptions. RMTTP (Du et al. 2016) is the seminal work demonstrating

that TPP based-RNN approach achieve better prediction performances w.r.t. classical TPP models (e.g. Point, Hawkes or Self-correcting processes).

In the RMTTP formulation,  $\lambda^*(t) = \exp(w^t(t - t_j) + v^T h_j + b^t)$ , with  $v^T$  a column vector and  $w^t, b^t$  scalars learned during the training of the network and  $h_j$  a vector embedding the history of the events. The exponential function is used as a non-linear transformation and guarantees that the intensity is positive.

Going beyond the analytic expression, the core idea in such a formulation is that  $\lambda$  summarizes three contributions:

1.  $(w^t(t - t_j))$  emphasizes the influence of the current event  $j$ ;
2.  $v^T h_j$  represents the accumulative influence of the past events, i.e. the history;
3.  $b^t$  gives a base intensity level for the occurrence of the next event.

The general formulation leaves room for some implementation choices that mostly depend on the application domain. For example, the embedding of the history can be done considering the timestamp or inter-event time  $\tau$  or its logarithm  $\log \tau$ . Mark embedding can be different, categorical marks are usually encoded with an embedding layer, while real marks are directly fed into the NN.

Another milestone in the use of RNNs in temporal point process modeling is ERPP (Xiao et al. 2017b). In ERPP the emphasis is placed on the distinction between time series and the sequence of events. Here both sequences are not jointly modeled, as in the case of RMTTP, but are handled separately. More in detail, the modeling is based on the assumption that time series are typically more suitable in capturing sequences with a synchronous and regularly updated or constant profile. Conversely, the sequence of events can compactly capture more abrupt event-specific information that can affect the conditional intensity function for a longer period of time. Even if the proposed approach is general, the architectural choice is based on the use of two distinct RNNs, one that models the time series and the other the events. The RNN time series can timely update the intensity function while the RNN event sequence is used to efficiently capture long-range dependence on history. They can interact with each other through a non-linear synergistic mapping.

### 4.3 Identification and prediction of attacks to ICS with RMTTP and ERPP

In order to model, using Marked Temporal Point Processes, the events related to ICSs, malicious activities are defined as marks, therefore we use an embedding layer for malicious activities and related attributes. At this point the two



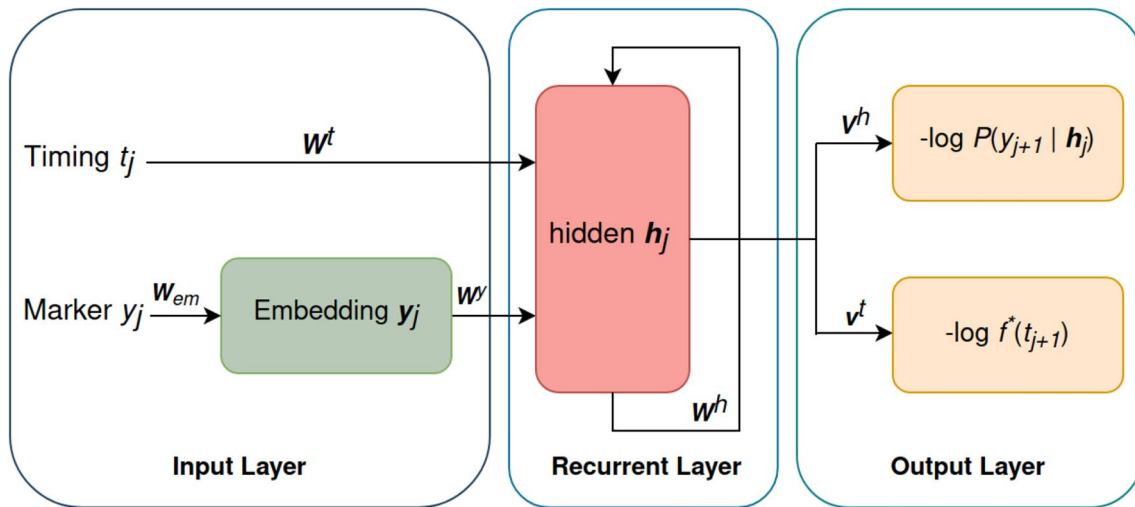


Fig. 1 RMTTP architecture

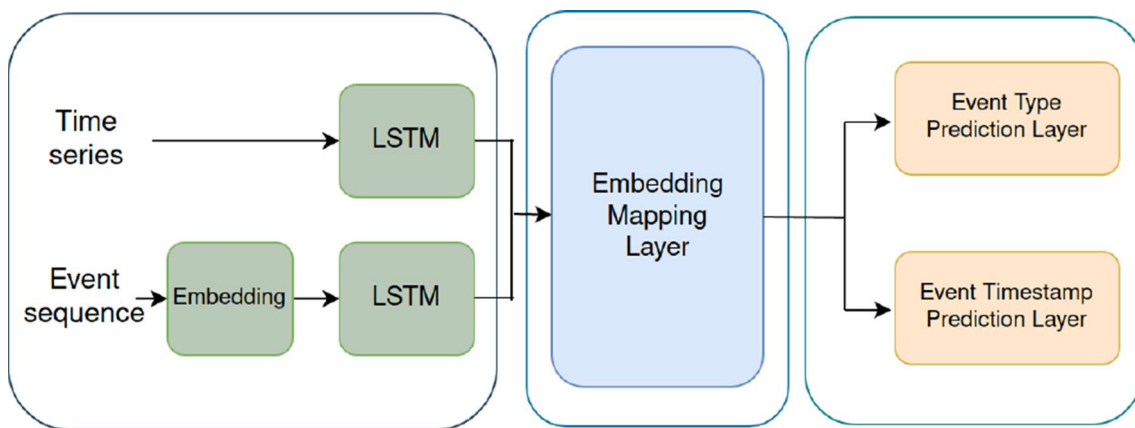


Fig. 2 ERPP architecture

approaches adopt different learning strategy: in RMTTP the feature vector of marks is merged with the temporal features and the resulting vector is inserted in the recurring level of an RNN. In RMTTP, model prediction is based on a unified representation of history dependence and a direct formulation 2 of the conditional intensity function  $\lambda^*(t_j + 1)$  which captures both time information from past events and event indicators. Since the prediction of marks also depends, in a non-linear way, on past temporal information, a unified approach can improve the performance of the model prediction when both times and events are correlated.

In ERPP, instead, timing vector and features marks are kept separate and fed into two different RNNs, meaning that we want to predict time and malicious activity separately. For both approaches the recurring level learns a  $h_j$  representation that incorporates and summarizes the nonlinear dependency of the event history. Then, on the basis of the

learned representation  $h_j$ , the learned model outputs the forecast for the next marker  $a_{j+1}$  and the timing  $t_{j+1}$  to calculate the respective loss functions (Figs. 1 and 2).

## 5 Experimental analysis

As discussed in Morris et al. (2015), currently there are very few data logs related to ICS well suited for intrusion and anomaly detection research. Most of the literature is based on data logs obtained by attacking locally owned systems. The logs resulting from those attacks are usually not shared publicly, thus making hard any comparison between different approaches. The most used dataset for IDS research is the 1999 DARPA dataset, however, despite its historic importance, this dataset is not a great choice for training and testing classifiers, since it was found to contain unintended



Fig. 3 The six different features of the dataset

Table 1 Macro categories

Label	Attack category
0	Normal traffic
1	Naïve malicious response injection (NMRI)
2	Complex malicious response injection (CMRI)
3	Reconnaissance (Recon)
4	Denial of service (DOS)
5	Malicious state command injection (MSCI)
6	Malicious parameter command injection (MPCI)
7	Malicious function command injection (MFCI)

patterns which made easier the task of learning differences among scenarios (McHugh 2000).

The dataset used in our experimental evaluation is a set of data logs, described in Morris et al. (2015), obtained from a virtual gas pipeline, including normal activities and 35 different types of cyber-attacks. The authors of the dataset previously captured data from a laboratory scale gas pipeline system discussed in Morris et al. (2011) obtaining the dataset discussed in Morris and Gao (2014). The dataset has been lately restructured after finding unintended patterns related to unrealistic behaviours that lead to overly optimistic accuracy in the classification tasks.

Every line of the dataset contains six different features, as shown in Fig. 3:

- a Modbus<sup>1</sup> Frame;
- two integers representing the categorization (see Table 1) and the specific attack (see Table 2);
- the Source and the Destination of the frame;
- a time stamp.

Table 2 Cyber attacks

Attack category	Attack name	Label	
NMRI	Random value attacks	29–31	
	Negative pressure attacks	32	
CMRI	Rise/fall attacks	25–26	
	Slope attacks	27–28	
Recon	Fast attacks	33–34	
	Slow attack	35	
	Device scan attack	20	
	Read id attack	23	
	Function code scan attack	24	
	DOS	Bad CRC attack	18
		Pump attack	13
		Solenoid attack	14
	MSCI	System mode attack	15
		Critical condition attacks	16–17
MPCI		Setpoint attacks	1–2
		PID gain attacks	3–4
		PID reset rate attacks	5–6
MFCI	PID rate attacks	7–8	
	PID deadband attacks	9–10	
	PID cycle time attacks	11–12	
	Clean registers attack	19	
	Force listen attack	21	
	Restart attack	22	

Table 3 Performance results macro categories

	Time error	Precision (%)	Recall (%)	F1-score
RMTTP	7.899	0.938	0.958	0.948
ERPP	6.100	0.955	0.963	0.959

Table 4 Performance results detailed cyber attacks

	Time error	Precision (%)	Recall (%)	F1-score
RMTTP	7.721	0.871	0.907	0.888
ERPP	6.875	0.935	0.954	0.945

<sup>1</sup> <https://modbus.org/>.

0	0.99	0	0	0	0	0	0	0
1	0.04	0.95	0.01	0	0	0	0	0
2	0.04	0	0.95	0	0	0	0	0
3	0.03	0	0	0.96	0.01	0	0	0
4	0.03	0	0	0	0.97	0	0	0
5	0.08	0	0.01	0	0.01	0.89	0	0
6	0.06	0	0	0.02	0	0	0.91	0
7	0.1	0	0	0	0.01	0	0	0.89
	0	1	2	3	4	5	6	7

**Fig. 4** RMTTP considering only the eight macro categories

## 5.1 Performance evaluation

The evaluation of the performances of our approach has been conducted by asking to the algorithm when the next event is going to happen and to which cyberattack class it belongs. It is straightforward to note that basically the same approach can be adopted for both identifying and classifying

past unlabeled activities and to predict future ones. This is because we can simply think about past activities as future ones by simply shifting in the past our notion of present and trying to predict the cyberattack class of the following events. The evaluation has been performed using both RMTTP and ERPP on the dataset, trying first to classify and predict the eight attack categories (7 attack classes,

0	0.99	0	0	0	0	0	0	0
1	0.03	0.95	0.02	0	0	0	0	0
2	0.04	0.01	0.95	0	0	0	0	0
3	0.02	0	0	0.97	0	0	0	0
4	0.02	0	0	0	0.98	0	0	0
5	0.03	0	0	0	0.01	0.94	0.01	0.02
6	0.02	0.01	0	0.01	0	0	0.96	0
7	0.08	0.01	0	0	0.01	0	0	0.9
	0	1	2	3	4	5	6	7

**Fig. 5** ERPP considering only the eight macro categories



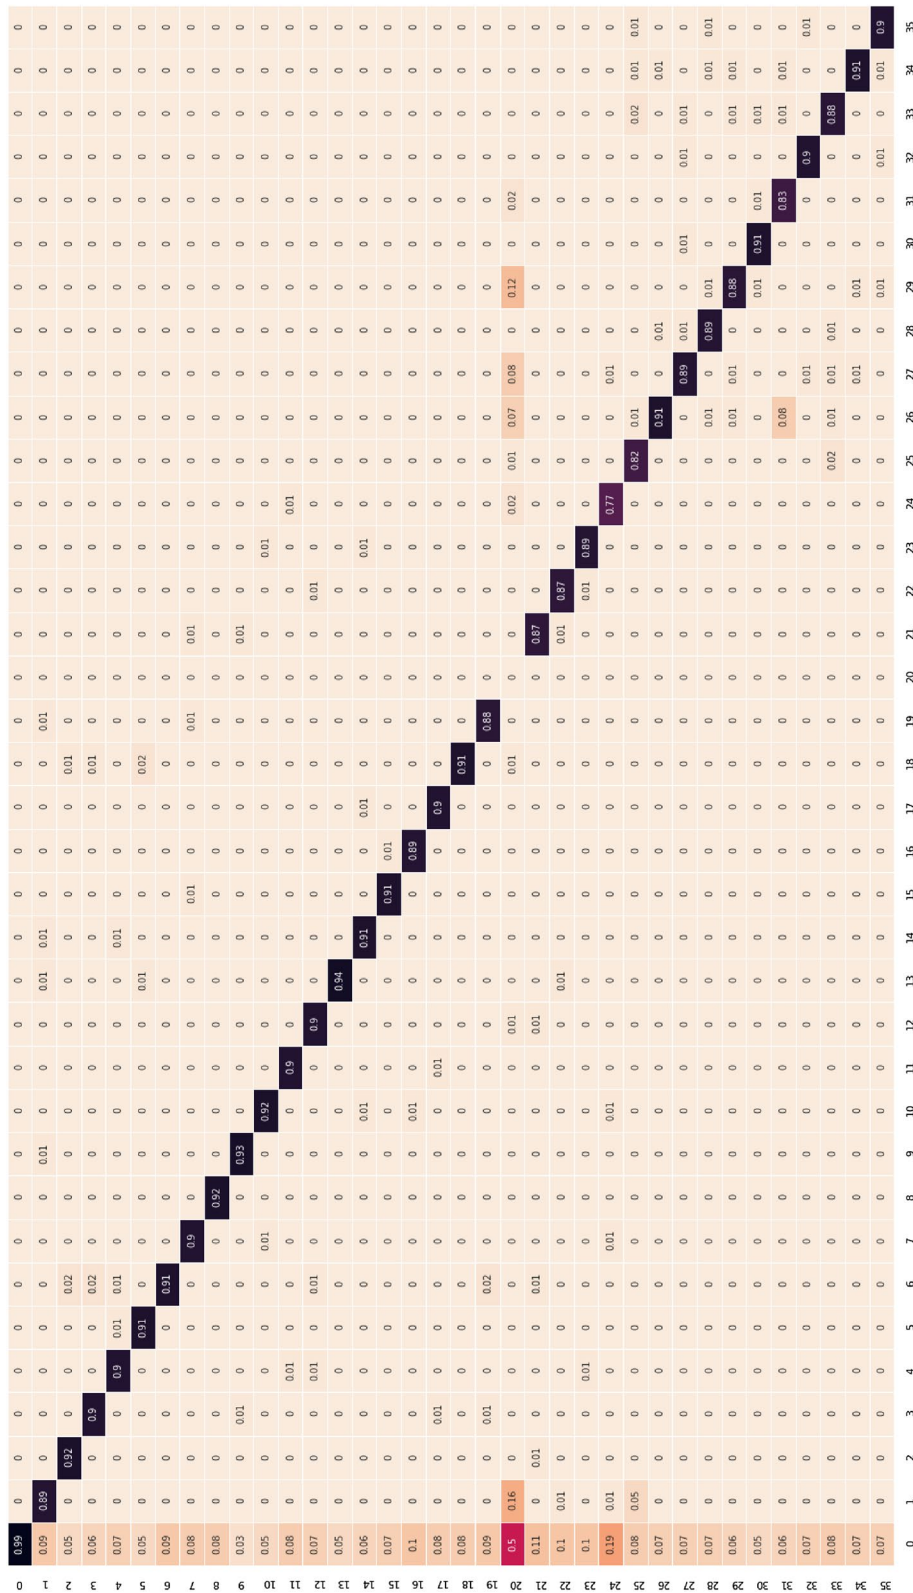


Fig. 6 RMTPP considering the full set of cyber-attacks



plus normal traffic). The very good results obtained with the attack categories identification and prediction motivated a new set of experiments, performed taking into account the full set of 35 cyberattacks (plus normal traffic) listed in Table 2. For a detailed description of each attack please refer to Morris et al. (2015). The results of both sets of experiments are shown in Tables 3 and 4.

We based our implementation of RMTTP on the paper: Du et al. (2016). In the case identification of the attack categories, RMTTP obtains an F1 score of 0.948 whereas, considering the full set of cyber attacks the F1-score drops to 0.888, still a consistent result that proves the effectiveness of the proposed approach.

The implementation of ERPP is based on the article (Xiao et al. 2017b). We use a single layer LSTM of size 32 with Sigmoid gate activations and tanh activation for hidden representation. It is important to underline that, as previously pointed out, for both RMTTP and ERPP, no prior knowledge about the hidden functional forms of the latent temporal dynamics is needed.

As shown in Table 3 and in Table 4, ERPP achieves better results w.r.t. RMTTP in all the metrics taken into account.

The results shown in Tables 3 and 4 clearly show the effectiveness of our approach. We can observe that the two proposed architectures achieve very similar performances on both sets of experiments. This lead to the consideration that using two separate neural networks for time and cyber attacks does not necessarily lead to an improvement in accuracy metrics. Our view is that this phenomena can be explained by the peculiarity of the application domain. In ICS attacks, in fact, the timestamp of a malicious activity and the activity itself are strictly correlated, suggesting that the strategy of learning them separately in the ERPP has marginal effect. The quality of the results can also be appreciated by observing the confusion matrices. In Figs. 4 and 5 we can analyze the results of our performance evaluation related to the attack categories, where the labels are the ones listed in Tables 3 and 4. The results of the detailed experimental analysis, related to the full set of attacks, can be observed in Figs. 6 and 7. In all four scenarios the confusion matrices are very dense on the diagonal, thus highlighting the quality of the classification obtained by means of both the proposed approaches.

The reason that the use of deep learning algorithms turns out to be outperforming for solving prediction problems in different application domains is that it manages to capture and express complex dependencies between data better than traditional approaches.

## 6 Conclusions and future work

In this work, we proposed the use of MTPP, and, more in detail ERPP and RMTTP in order to express the sequence of actions contained in logs or live stream data from Industrial

Control Systems. The approach proposed has proved to be an effective and efficient choice in the task of identifying malicious activities and even predict future threats. The techniques discussed have been applied to a dataset coming from a gas pipeline, identifying and predicting a small set of attack categories and even a detailed list of specific cyber attacks. Results obtained clearly show the effectiveness of the approach, with very good values of accuracy and recall.

Note that, although the effectiveness of the method is evaluated on the log of a gas pipeline system, the formulation of the MTPP problem is general and, eventually, can be applied to different types of ICSs, when the logs extracted from these systems are mapped as sequences of temporal events, including a label identifying them as normal traffic or as attack traffic with the designated attack class.

As future work, we are going to include additional features in our architecture taking into account different measurements and values gathered by the ICS. By using this approach, we trust that the temporal dependency will be modeled in a more precise way and the dynamics that govern the generation process of the activity events would be better embedded.

**Funding** Open access funding provided by Università della Calabria within the CRUI-CARE Agreement.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Ahmad Z, Khan AS, Nisar K, Haider I, Hassan R, Haque MR, Tarmizi S, Rodrigues JJPC (2021) Anomaly detection using deep neural network for iot architecture. *Appl Sci* 11(15):7050
- Aleskerov E, Freisleben B, Rao B (1997) Cardwatch: a neural network based database mining system for credit card fraud detection. In: *Proc. of the IEEE/IAFE 1997 CIFE*. IEEE, pp 220–226
- Alsoufi MA, Razak S, Siraj MM, Nafea I, Ghaleb FA, Saeed F, Nasser M (2021) Anomaly-based intrusion detection systems in iot using deep learning: a systematic literature review. *Appl Sci* 11(18):8383
- Apa L, Penagos CM (2013) Compromising industrial facilities from 40 miles away. IOActive Technical White Paper
- Aven T (2007) A unified framework for risk and vulnerability analysis covering both safety and security. *Reliab Eng Syst Saf* 92(6):745–754



- Babbin J (2006) Security log management: identifying patterns in the chaos. Elsevier, Amsterdam
- Barbará D, Li Y, Couto J, Lin J-L, Jajodia S (2003) Bootstrapping a data mining intrusion detection system. In: Proceedings of the 2003 ACM symposium on applied computing, pp 421–425
- Bieber P, Blanquart J-P, Descargues G, Dulucq M, Fourastier Y, Hazane E, Julien M, Léonardon L, Sarouille G (2012) Security and safety assurance for aerospace embedded systems. In: Embedded real time software and systems (ERTS2012)
- Brockett PL, Xia X, Derrig RA (1998) Using Kohonen's self-organizing feature map to uncover automobile bodily injury claims fraud. *J Risk Insurance* 65:245–274
- Checkoway S, McCoy D, Kantor B, Anderson D, Shacham H, Savage S, Koscher K, Czeskis A, Roesner F, Kohno T et al. (2011) Comprehensive experimental analyses of automotive attack surfaces. In: USENIX security symposium, vol 4, San Francisco, p 2021
- Chen K, Lu SC, Teng HS (1990) Adaptive real-time anomaly detection using inductively generated sequential patterns. In: Fifth intrusion detection workshop. SRI International, Menlo Park, CA
- Cockram TJ, Lautieri SR (2007) Combining security and safety principles in practice. In: 2007 2nd Institution of engineering and technology international conference on system safety. IET, pp 159–164
- Daley DJ, Vere-Jones D (2003) An introduction to the theory of point processes. Springer-Verlag, Berlin (ISBN 978-0-387-21564-8)
- Dilek S, Çakır H, Aydın M (2015) Applications of artificial intelligence techniques to combating cyber crimes: a review. arXiv preprint [arXiv:1502.03552](https://arxiv.org/abs/1502.03552)
- Djenna A, Harous S, Saidouni DE (2021) Internet of things meet internet of threats: New concern cyber security issues of critical cyber infrastructure. *Appl Sci* 11(10):4580
- Du N, Dai H, Trivedi RS, Upadhyay U, Gomez-Rodriguez M, Song L (2016) Recurrent marked temporal point processes: Embedding event history to vector. In: Proceedings of the 22nd ACM SIGKDD
- Edgeworth FY (1887) Xli on discordant observations. *Lond Edinb Dublin Philos Mag J Sci* 23(143):364–375
- Emamian V, Kaveh M, Tewfik AH (2000) Robust clustering of acoustic emission signals using the Kohonen network. In: 2000 IEEE ICASSP, vol 6. IEEE, pp 3891–3894
- Ertöz L, Steinbach M, Kumar V (2002) A new shared nearest neighbor clustering algorithm and its applications. In: Workshop on clustering high dimensional data and its applications at 2nd SIAM international conference on data mining, pp 105–115
- Ester M, Kriegel H-P, Sander J, Xiaowei X et al (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd* 96:226–231
- Fortino G, Guzzo A, Ianni M, Leotta F, Mecella M (2020) Exploiting marked temporal point processes for predicting activities of daily living. In: 2020 IEEE international conference on human-machine systems (ICHMS). IEEE, pp 1–6
- Fortino G, Guzzo A, Ianni M, Leotta F, Mecella M (2021) Predicting activities of daily living via temporal point processes: approaches and experimental results. *Comput Electr Eng* 96:107567
- Frustaci M, Pace P, Aloï G, Fortino G (2018) Evaluating critical security issues of the iot world: present and future challenges. *IEEE Internet Things J* 5:2483–2495
- Greensmith J (2015) Securing the internet of things with responsive artificial immune systems. In: Proceedings of the 2015 annual conference on genetic and evolutionary computation, pp 113–120
- Grøtan TO, Jaatun MG, Knut Øien, Onshus T (2007) The sesa method for assessing secure remote access to safety instrumented systems. SINTEF Report A, 1626
- Guha S, Rastogi R, Shim K (2000) Rock: a robust clustering algorithm for categorical attributes. *Inf Syst* 25(5):345–366
- Gumaei A, Hassan MM, Shamsul Huda M, Rafiul Hassan M, Camacho D, Del Ser J, Fortino G (2020) A robust cyberattack detection approach using optimal features of SCADA power systems in smart grids. *Appl Soft Comput* 96:106658
- Guzzo A, Ianni M, Pugliese A, Saccà D (2020) Modeling and efficiently detecting security-critical sequences of actions. *Futur Gener Comput Syst* 113:196–206
- Hawkes AG (1971) Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58(1):83–90
- Hawkes AG, Oakes D (1974) A cluster process representation of a self-exciting process. *J Appl Probab* 11(3):493–503
- Ianni M, Masciari E (2022) Some experiments on high performance anomaly detection. In: 2022 30th Euromicro international conference on parallel, distributed and network-based processing (PDP), pp 226–229. [10.1109/PDP55904.2022.00042](https://doi.org/10.1109/PDP55904.2022.00042)
- Ianni M, Masciari E, Mazzeo GM, Mezzanzanica M, Zaniolo C (2020) Fast and effective big data exploration by clustering. *Futur Gener Comp Syst* 102:84–94
- Isham V, Westcott M (1979) A self-correcting point process. *Stochastic Process Appl* 8(3):335–347
- Jain AK, Dubes RC (1988) Algorithms for clustering data. Prentice-Hall Inc, New Jersey
- Jang-Jaccard J, Nepal S (2014) A survey of emerging threats in cybersecurity. *J Comput Syst Sci* 80(5):973–993
- Jiang M-F, Tseng S-S, Chih-Ming S (2001) Two-phase clustering process for outliers detection. *Pattern Recogn Lett* 22(6–7):691–700
- Johnsen SO (2012) Resilience at interfaces: improvement of safety and security in distributed control systems by web of influence. In: Information management & computer security
- Johnson CW (2012) Cybersafety: on the interactions between cybersecurity and the software engineering of safety-critical systems. In: Achieving system safety, pp 85–96
- Kent K, Souppaya M (2006) Guide to computer security log management. NIST Spec Publ 92:1–72
- Kingman JFC (1993) Poisson processes, vole 3 of Oxford Studies in Probability. Clarendon Press, Oxford Science Publications
- Kohonen T (1990) The self-organizing map. *Proc IEEE* 78(9):1464–1480
- Kornecki AJ, Zalewski J (2010) Safety and security in industrial control. In: Proceedings of the sixth annual workshop on cyber security and information intelligence research, pp 1–4
- Koscher K, Czeskis A, Roesner F, Patel S, Kohno T, Checkoway S, McCoy D, Kantor B, Anderson D, Shacham H et al (2010) Experimental security analysis of a modern automobile. In: 2010 IEEE symposium on security and privacy. IEEE, pp 447–462
- Kumar V (2005) Parallel and distributed computing for cybersecurity. *IEEE Distrib Syst Online* 6(10)
- Lee A, Brewer T (2009) Smart grid cyber security strategy and requirements. Draft Interagency Report NISTIR, 7628
- Mahbub M (2020) Progressive researches on iot security: an exhaustive analysis from the perspective of protocols, vulnerabilities, and preemptive architectonics. *J Netw Comput Appl* 168:102761
- Mahoney MV, Chan PK (2003) Learning rules for anomaly detection of hostile network traffic. In: Third IEEE international conference on data mining. IEEE, pp 601–604
- Manadhata PK, Wing JM (2010) An attack surface metric. *IEEE Trans Softw Eng* 37(3):371–386
- Markou M, Singh S (2003a) Novelty detection: a review-part 1: statistical approaches. *Signal Process* 83(12):2481–2497
- Markou M, Singh S (2003b) Novelty detection: a review-part 2: neural network based approaches. *Signal Process* 83(12):2499–2521
- McHugh J (2000) Testing intrusion detection systems: a critique of the 1998 and 1999 darpa intrusion detection system evaluations as performed by Lincoln laboratory. *ACM Trans Inf Syst Security (TISSEC)* 3(4):262–294

- Meixell B, Forner E (2013) Out of control: demonstrating scada exploitation. Black Hat, p 2013
- Morgan S (2020) Cybercrime to cost the world \$10.5 trillion annually by 2025, Nov 2020. <https://cybersecurityventures.com/hackerpocalypse-cybercrime-report-2016/>
- Morris T, Gao W (2014) Industrial control system traffic data sets for intrusion detection research. In: International conference on critical infrastructure protection. Springer, pp 65–78
- Morris T, Srivastava A, Reaves B, Gao W, Pavurapu K, Reddi R (2011) A control system testbed to validate critical infrastructure protection concepts. *Int J Crit Infrastruct Prot* 4(2):88–103
- Morris TH, Thornton Z, Turnipseed I (2015) Industrial control system simulation and data logging for intrusion detection system research. In: 7th annual southeastern cyber security summit, pp 3–4
- Otey M, Parthasarathy S, Ghoting A, Li G, Narravula S, Panda D (2003) Towards nic-based intrusion detection. In: Proc. of ACM SIGKDD, pp 723–728
- Pires AM, Santos-Pereira C (2005) Using clustering and robust estimators to detect outliers in multivariate data. In: Proceedings of the international conference on robust statistics
- Ramadas M, Ostermann S, Tjaden B (2003) Detecting anomalous network traffic with self-organizing maps. In: Int. workshop on recent advances in intrusion detection. Springer, pp 36–54
- Schultz EE (2002) A framework for understanding and predicting insider attacks. *Comput Secur* 21(6):526–531
- Smith KA, Woo F, Ciesielski V, Ibrahim R (2002) Matching data mining algorithm suitability to data characteristics using a self-organizing map. In: Hybrid information systems. Springer, pp 169–179
- Tahsien SM, Karimipour H, Spachos P (2020) Machine learning based solutions for security of internet of things (iot): a survey. *J Netw Comput Appl* 161:102630
- Xiao S, Yan J, Farajtabar M, Song L, Yang X, Zha H (2017a) Joint modeling of event sequence and time series with attentional twin recurrent neural networks. ArXiv, abs/1703.08524
- Xiao S, Yan J, Yang X, Zha H, Chu SM (2017b) Modeling the intensity function of point process via recurrent neural networks. In: Proceedings of the 31st AAAI conference on artificial intelligence, pp 1597–1603
- Yan J (2019) Recent advance in temporal point process: from machine learning perspective. SJTU technical report
- Yan J, Xu H, Li L (2019) Modeling and applications for temporal point processes. In: Proceedings of the 25th ACM SIGKDD, pp 3227–3228
- Zikopoulos P, Eaton C et al (2011) Understanding big data: analytics for enterprise class hadoop and streaming data. McGraw-Hill Osborne Media, New York

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.