



# Connecting Web Event Forecasting with Anomaly Detection: A Case Study on Enterprise Web Applications Using Self-supervised Neural Networks

Xiaoyong Yuan<sup>1</sup>(✉), Lei Ding<sup>2</sup>, Malek Ben Salem<sup>3</sup>, Xiaolin Li<sup>4</sup>,  
and Dapeng Wu<sup>5</sup>

<sup>1</sup> Michigan Technological University, Houghton, MI 49931, USA  
xyyuan@mtu.edu

<sup>2</sup> American University, Washington, DC 20016, USA  
ding@american.edu

<sup>3</sup> Accenture Labs, Arlington, VA 22209, USA  
malek.ben.salem@accenture.com

<sup>4</sup> Cognization Lab, Palo Alto, CA 94306, USA  
xiaolinli@ieee.org

<sup>5</sup> University of Florida, Gainesville, FL 32608, USA  
dpwu@ufl.edu

**Abstract.** Recently web applications have been widely used in enterprises to assist employees in providing effective and efficient business processes. Forecasting upcoming web events in enterprise web applications can be beneficial in many ways, such as efficient caching and recommendation. In this paper, we present a web event forecasting approach, DEEPEVENT, in enterprise web applications for better anomaly detection. DEEPEVENT includes three key features: web-specific neural networks to take into account the characteristics of sequential web events, self-supervised learning techniques to overcome the scarcity of labeled data, and sequence embedding techniques to integrate contextual events and capture dependencies among web events. We evaluate DEEPEVENT on web events collected from six real-world enterprise web applications. Our experimental results demonstrate that DEEPEVENT is effective in forecasting sequential web events and detecting web based anomalies. DEEPEVENT provides a context-based system for researchers and practitioners to better forecast web events with situational awareness.

**Keywords:** Anomaly detection · Event forecasting · Self-supervised learning · Neural networks

## 1 Introduction

Recently web applications play a major role in many enterprises. Enterprise web applications boost productivity by assisting employees in performing their daily

tasks. On one side, web events in the enterprise web applications provide insightful sources for analyzing employee behaviors. By forecasting web events based on user behavior, we can provide better recommendation, caching, pre-fetching, and load balancing for enterprise web applications [8,36]. Web event forecasting has been investigated for many years [2,32,35]. However, two critical challenges exist in characterizing the sequence of web events: 1) Although many works have been conducted in forecasting web events, very few works investigated enterprise web applications. Web applications significantly improve productivity in modern enterprises and have become essential components in enterprise operations. Enterprise web applications are more well-organized and closely connected to each other compared to other web applications. Specific patterns of web events exist in the employees' browsing behavior. 2) With the increasing complexity and functionality of web applications, it becomes hard to forecast web events. Browsing a web page can produce a sequence of web events, and the length of the sequence varies a lot for different web pages. For example, a web event is usually made for the HTML of the web page itself, and subsequent events are made for each image, plug-in, audio clip, and other content referenced in the HTML. The web events for each kind of content increase the complexity of the sequence. The format and semantics of web events vary significantly from application to application. It is urgent to develop techniques to automate the process of characterizing web events and representing them in the desired way.

On the other side, enterprise web applications provide extended connectivity to an organization's assets and increase the attack surface of its web-facing infrastructure. Enterprise web applications have become favorite targets of cyber-attacks due to easy access and constantly increasing vulnerabilities. Anomaly detection is a critical component to protect web applications against cyber threats. Supervised learning based anomaly detection solutions build detection systems by discovering abnormal behavioral patterns with the use of labeled training data [6,18,29,37]. Unfortunately, high quality annotated data is not easy to obtain, given the velocity, volume, and real-time nature of web events. Usually, labeled data are very imbalanced as it is hard to collect a large number of labeled anomalies as opposed to normal web events. Insufficient and imbalanced training data hinder the performance of machine learning models [12]. Limited labeled data from previous attacks undercut the ability to use supervised models, and constantly evolving attacks make such supervised models irrelevant.

To address the above challenges, we propose a novel deep neural network based approach, DEEPEVENT. By connecting web event forecasting and anomaly detection, DEEPEVENT improves the performance of web event forecasting for complicated web events, while detects anomalies by identifying the most unlikely events in the sequence. We evaluated DEEPEVENT on web events collected from six real-world web applications in a company for three months.

In this paper, we make the following contributions:

**1) Context-based web events analysis for better user behavior forecasting:** DEEPEVENT characterizes not only the individual web events, but also the relationships among the web events that co-occur within a context, *i.e.*, a




Fig. 1. Workflow of DEEPEVENT.

sequence of events that are commonly produced together for a specific user and web applications’ flow characteristics. We leverage deep neural networks to model the context of a web event in a sequence. Based on the context, DEEPEVENT predicts the web events that are likely to appear with associated probabilities.

2) Self-supervised learning to overcome the need for labeled data:

By connecting web event forecasting and anomaly detection, we can formulate anomaly detection as a self-supervised learning problem [13,30]. Self-supervised learning tasks do not require any prior knowledge of anomalous web events’ features and their sequential relationships, but leverage naturally existing evidence as labeled data for training. Therefore, self-supervised learning overcomes the need for labeled data in anomaly detection. In addition, we leverage a pre-training process that learns the representations of web events to further compensate for the lack of labeled data.

3) Quantitative measures of the anomaly: Different from traditional anomaly detection methods, DEEPEVENT goes beyond binary prediction. Upon encoding the relationship between web events within the vector space, we can quantitatively predict the web events appearing in similar contexts. Therefore, DEEPEVENT predicts if a web event belongs to normal behavior and also measures the deviation of the abnormal event with respect to expected normal events. In such a way, we provide a quantitative measure of the detected anomaly.

We empirically validate that forecasts made by DEEPEVENT are more accurate than baseline solutions. We show that with the assist of web event forecasting, DEEPEVENT can identify anomalous events (e.g., real-world exploits, widely-used web attacks) using the proposed quantitative measures. We find these results encouraging and note that they highlight the benefit of forecasting web events that co-occur within a context for anomaly detection.

## 2 Workflow of DeepEvent

We design three main components in DEEPEVENT: an event extractor to extract semantic events from web requests, a context-based model for sequential web event forecast as well as an anomaly score evaluator.

Figure 1 illustrates the workflow of DEEPEVENT. DEEPEVENT ingests a sequence of web requests from web applications as its input, e.g., with a length of  $n - 1$ . Then it extracts the contents from web requests and converts them

into a sequence of events  $\{e_1, e_2, \dots, e_{n-1}\}$ . After event extraction, DEEPEVENT performs context-based modeling where it first encodes the events into event embedding and sequence embedding. Event embedding represents the content of each event. Sequence embedding represents the order of each event in the sequence. The embedding output will then be used to build a neural network (as described in Sect. 2.2) to learn long-term dependencies of the events. The trained context-based model will calculate the probability distribution of possible events to appear as the next one given  $\{e_1, e_2, \dots, e_{n-1}\}$  and provide the forecast of the upcoming event  $e_n$ . In addition, DEEPEVENT uses the predicted probability to evaluate the anomaly score of  $e_n$ .

We describe the three main components in the following sections.

## 2.1 Event Extraction

The purpose of event extraction is to extract semantic events from web requests. We use sequences of events to characterize web requests. For well-formatted web requests (*e.g.*, following REST API design), one can easily map these requests to events by extracting their HTTP methods and well-defined endpoints. However, most web requests may not share a well-organized representation or follow a consistent format. For example, URI paths may not be named around resources. Web developers may use various URI paths for the same resource on the web server. For instance, WordPress provides five different URIs and a custom one for users to access a post<sup>1</sup>. URI paths can be generated by randomized algorithms or encoding algorithms. Diverse web requests impede the effort to extract semantic events from URIs.

We propose a three-step event extraction method, including content extracting, path uniforming, and “rare” event identifying.

- **Extract Content.** We extract three components from web requests: HTTP methods (*i.e.*, GET, POST, UPDATE, etc.), URI paths, and the number of URI parameters. Our experimental results show that using merely these three components are effective in representing user behavior.
- **Uniform Path.** We apply a two-character Markov Chain model to detect the “random” elements in URI paths. We first segment URI paths into “elements” separated by special characters such as “/” and “-”. We then investigate every character in the element from left to right. If the likelihood of the upcoming character based on the preceding two characters is lower than a certain threshold, then we consider the element as “random.”<sup>2</sup>
- **Identify “RARE” Events.** We consider the events occurring less than  $T$  times in the training data as “RARE” events. In this way, we learn the information of “RARE” events during training, which helps us to understand if such rare events are anomalous or not.

<sup>1</sup> [https://codex.wordpress.org/Using\\_Permalinks](https://codex.wordpress.org/Using_Permalinks).

<sup>2</sup> We detect randomness in URIs based on a gibberish detection tool (<https://github.com/rrenaud/Gibberish-Detector>).

The entire process of pre-processing data proceeds as follows:

1. Extract HTTP methods, URI paths, and URI queries from web requests;
2. Segment URI paths into “elements” by special characters as delimiters;
3. Flag the “elements” as “RANDOM,” if they are randomly generated or encoded;
4. Calculate the number of key-value pairs in the URI query;
5. Concatenate HTTP method, “derandomized” URI path, and the number of the URI query as an event.
6. If an event has never seen in the training set or has occurred less than  $T$  times, convert the event into a “RARE” event.

## 2.2 Context-Based Modeling

We propose context-based web request modeling, which takes a sequence of contextual events as input and outputs a sequence of corresponding events. We can mask the event of interest in the input sequence and train a model to predict it for the given sequence. Recurrent Neural Networks (RNNs), as well as their variants such as Long Short-Term Memory (LSTM) [11], have been proposed for security analytics in the sequential analysis due to their outstanding performance. Recently, self-attention neural networks [5] have been shown to be much more effective to capture long-term dependencies in a sequence compared to RNNs. Specifically, RNN passes the hidden states through the previous state while self-attention neural networks construct direct links between events within the context, which brings great merit in learning from the long-distance context.

**Self-supervised Learning.** To compensate for the lack of labeled data, we design a self-supervised learning task for event forecasting. Most existing supervised models are limited to the high-quality labeled data. In this paper, we leverage the existing event requests as the labels without any manual annotations. In practice, we randomly mask 25% events in the input sequence and replace these masked events with “mask” labels in the input sequence. We train the neural networks to predict “mask” events. In this way, the neural network learns the relationship of events and their dependency in the sequences. We use this neural network as a pretrained model for further event forecasting and anomaly detection. Our experimental results show that self-supervised pre-training significantly improves the prediction performance.

## 2.3 Anomaly Detection

We propose a way to calculate the anomaly score to quantitatively measure the likelihood of a new web request being anomalous. Given the current context, we predict a set of web requests that are likely to appear with associated probabilities. For a received web request, we rank it with the predicted set of web requests

based on their associated probabilities calculated by the trained neural network. We calculate the anomaly score for the incoming web request as follows:

$$s = 1 - \frac{1}{\tau + 1}, \quad (1)$$

where  $\tau$  denotes the rank of the newly received request based on its likelihood to appear (*i.e.*, the probability calculated by the context-based model), and  $s$  is in the range of  $(0, 1)$ . Higher anomaly scores indicate higher confidence level in classifying the event as an anomaly.

The anomaly score indicates the degree of deviation of the incoming web request from the expected normal requests. In Sect. 4.3, we show that the proposed anomaly score is able to differentiate anomalous web requests (*e.g.*, produced by various real-world web based attacks) from normal requests.

### 3 Methodology of Context-Based Modeling

In this section, we describe the detailed approach for context-based modeling (Sect. 2.2) and explain how we adapt three different types of neural networks to predict web requests. We compare the performance of the three neural networks in Sect. 4.4.

#### 3.1 Self-attention Based Modeling

The design of self-attention based modeling is shown in Fig. 2. We first embed input events using the embedding layer, then use a self-attention neural network to encode the sequence and learn the dependency between events. In the output layer, we apply a Softmax function to squash the neural network and predict future events with associated probabilities.

We introduce two critical components of our adaptation: 1) event embedding and sequence embedding; and 2) a self-attention neural network.

**Event Embedding and Sequence Embedding.** Word embedding is widely used to represent the semantic meaning of tokens in NLP tasks. We adopt the basic idea of word embedding by extracting semantic representation of web events and generate event embedding. Sequential information represents the relative positions of events in the sequence. However, self-attention neural networks do not contain sequential information of events, because the distances of events are the same. Therefore, we add sequential information into the neural networks using a sequence embedding layer. Specifically, our embedding layer maps the web event and its position in the sequence into two 128-dimension vectors:  $\{EE_1, EE_2, \dots, EE_n\}$  and  $\{SE_1, SE_2, \dots, SE_n\}$ . After generating the embedding on events and their positions, we sum up these two sequences as an embedded sequence  $\{em_1, em_2, \dots, em_n\}$ , and feed it to the encoding layers of the self-attention network.

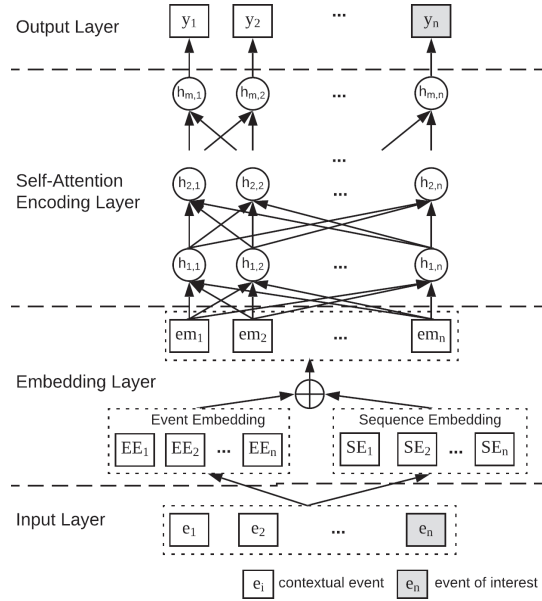


Fig. 2. Self-attention based modeling.

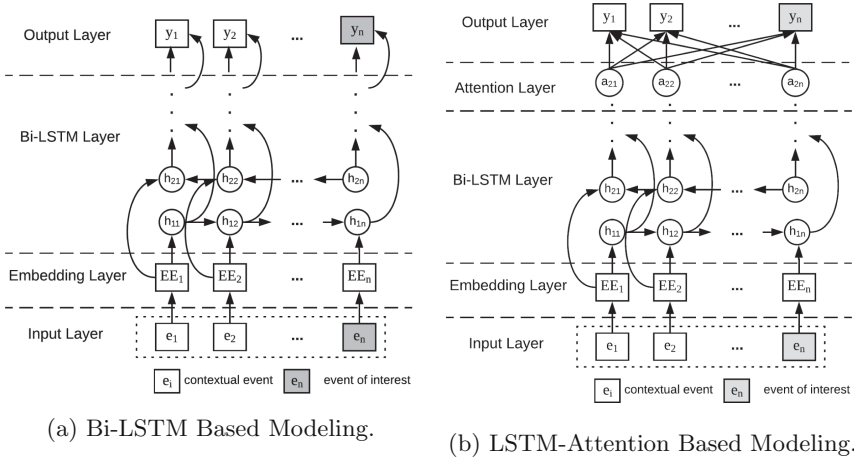
**Self-attention Neural Network.** We use a new type of neural networks, self-attention neural network [39], to solve sequential prediction problem. Specifically, we adapt BERT [5], a self-attention based neural network. BERT has been shown to outperform RNNs in almost all NLP tasks, and achieve state-of-the-art performance [5]. The success of BERT mainly comes from a scaled dot-product attention neural network:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \tag{2}$$

where  $Q, K, V$  are the query, key, and value of dimension  $d_k$ . In the self-attention neural network of DEEPEVENT,  $Q, K, V$  come from the same sequence of embedded events. The network learns to pay attention to the specific events in the sequence and captures the dependencies between events in the sequence. We use a model called Transformer, including a multi-head attention neural network (stacking several self-attention neural networks), a normalization layer, and a Softmax function. For the details of self-attention mechanism and Transformer, we refer the reader to [39].

### 3.2 Bi-LSTM Based Modeling

In addition to self-attention neural network, we implement a bidirectional LSTM neural network, called Bi-LSTM, which is widely used in sequential analysis. Similarly to Self-Attention Neural Network, we apply the event embedding



**Fig. 3.** Bi-LSTM and LSTM-attention based modeling.

here to encode events and use LSTM to represent their sequential relationship. Figure 3a illustrates the architecture of Bi-LSTM based modeling. Bi-LSTM uses an embedding layer to convert an input event into a 128-dimensional vector, then deploys multiple bidirectional LSTM layers to extract semantic information from the events. We test 1, 2, 3 layers of LSTM for each direction in the experiment. A fully connected layer with a Softmax function is stacked on top of LSTM layers to output the final prediction related to the event of interest.

### 3.3 LSTM-Attention Based Modeling

Attention mechanism is recently proposed to surpass recurrent neural networks by remembering longer sequences. The attention mechanism aims to pinpoint key events from a long sequence. We adapt an LSTM-Attention neural network using an additive attention mechanism. Figure 3b illustrates the architecture of LSTM-Attention based modeling. The additive attention mechanism sums up the outputs of LSTM with their weights and outputs the weighted sum as the prediction. The LSTM-Attention neural network consists of an embedding layer, bidirectional LSTM layers, an additive attention layer, and a fully connected layer with Softmax. The embedding layer and LSTM layers follow the setting in Bi-LSTM neural network. The attention layer learns the weight of each event and applies them to the final output. Note that the neural network in [37] detects anomaly based on the content of a single request. However, we use LSTM-Attention to predict an event in a sequence. We set the hidden number of the embedding layer in Bi-LSTM and LSTM-Attention to 128 and apply a drop-out mechanism with 20% dropout in LSTM layers to avoid overfitting.

**Table 1.** Dataset description.

Application	# of Events			# of Unique Events
	Train	Valid	Test	
Workqueue	146,489	18,026	19,004	101
DataRepo1	98,798	11,500	13,596	258
DevOpsApp	134,230	23,652	35,517	1,787
DataAnalyzer1	757,626	106,836	64,681	1,442
DataAnalyzer2	363,787	59,285	65,222	329
DataRepo2	63,862	7,235	8,307	37


## 4 Evaluation

In this section, we describe the experiments to evaluate DEEPEVENT on real-world web applications. We introduce our experimental settings in Sect. 4.1. We designed experiments to answer the following questions:

1. What is DEEPEVENT’s performance in web event forecast compared with existing methods (Sect. 4.2)?
2. How effective is DEEPEVENT in evaluating real-world threats (Sect. 4.3)?
3. How do different neural networks perform in DEEPEVENT (Sect. 4.4)?
4. How do different model settings (*e.g.*, pre-training, window size) affect prediction performance (Sect. 4.5)?

### 4.1 Experimental Setup

**Dataset.** We evaluate DEEPEVENT on six real-world enterprise web applications described: 1) Workqueue: work queue of data collections and analysis; 2) DataRepo1: meta data repository; 3) DevOpsApp: continuous integration application for DevOps; 4) DataAnalyzer1: data reporting and visualization; 5) DataAnalyzer2: data analysis and visualization; 6) DataRepo2: graph and meta data repository. We collected 84 days of HTTP requests from the above applications in a real-world enterprise network. Figure 4 shows the number of HTTP requests for each web application in the dataset. We separate the dataset based on the date and use the first 64 days of data for training, 10 days for validation, and the last 10 days for testing. We assume all the HTTP requests are legitimate. Note that we only use the HTTP requests produced by normal users, and we exclude HTTP requests that are automatically generated by machines (*e.g.*, heartbeat requests). Table 1 summarizes the statistics of training/validation/test data of those six applications. We report the number of web events used in training/validation/test data and the number of unique events observed in the training data. The number of unique events indicates how many different events after event extraction occur in the datasets.



**Fig. 4.** Number of HTTP requests of six web applications in the dataset. We annotate the days used in the training/validation/test data. The number of HTTP requests on weekdays is higher than that on weekends, indicating the activity pattern of an enterprise network.

**Evaluation Metrics.** In the experiments, we evaluate the performance of DEEPEVENT using Top-N Accuracy and Anomaly Score. **Top-N Accuracy** measures the event prediction performance. *Top-N Accuracy* calculates the percentage of the correct event occurs among the top N events predicted by the model. In the experiments, we report Top-1 and Top-10 accuracy. **Anomaly Score**, as defined in Eq. 1, evaluates the capability in terms of differentiating anomalous requests from normal requests.

**Modeling Settings.** We use self-attention based modeling in DEEPEVENT in the experiments in Sect. 4.2 and Sect. 4.3 as it performs better in general compared to Bi-LSTM, and LSTM-attention based modeling for the six enterprise

web applications. We show and discuss the comparison among the three for all six web applications in Sect. 4.4.

We set the threshold of “rare events” to 2 ( $T = 2$ ). We set the number of hidden layers of DEEPEVENT to 8, the number of attention heads in each layer to 8, the number of hidden neurons in each head to 128, the batch size to 128. We use Cross-Entropy as our loss function. We optimize the loss function using Adam [17] with L2 weight decay. We set the learning rate to 0.001 in pre-training and reduce the learning rate by 10 in training. We train the models for 100 epochs. To accelerate the training process, we adopt an early-stop strategy, which ceases training if the cross-entropy loss of the validation data does not decrease in the past 10 epochs.

## 4.2 Evaluation of DeepEvent on Web Event Forecast

In this section, we evaluate DEEPEVENT on the web event forecast. We compare DEEPEVENT with two baseline models: a Markov model and an N-gram model.

The **Markov model** predicts the upcoming event depending on the current event. It assumes the upcoming event does not depend on the previous event. It learns transition probability from the training data. In this paper, we use a first-order Markov model [16].

The **N-gram model** aims to provide maximum likelihood estimates for the last event  $e_N$  given previous  $N - 1$  contextual events  $\{e_1, e_2, \dots, e_{N-1}\}$  [36]:

$$\arg \max Pr\{e_N | e_1, e_2, \dots, e_{N-1}\}, \quad (3)$$

where  $N$  denotes the number of events considered for prediction. Following the same setting in [33], we use a 3-gram model to predict the upcoming events.

In our experiments, we observe that DEEPEVENT performs better than the Markov and the 3-gram models. We report Top-1 and Top-10 accuracy of the three methods in Table 2. We highlight all the best results among the three methods. For all six applications, DEEPEVENT increases 11.2% Top-1 accuracy and 3.48% Top-10 accuracy on average compared with the best results of the Markov model and 3-gram Model.

The 3-gram model performs better than the Markov model, and DEEPEVENT outperforms both of them. This is because the long-distance context is lost in Markov model and 3-gram model. In contrast, DEEPEVENT has longer-term memory compared with the other two models so it can capture the sequential relationships of events that are not adjacent to each other, which is especially important for web applications because task-critical web requests may be separated by “not so relevant” requests. For example, browsing a web page can produce a sequence of web requests, where one web request is made for the HTML of the web page itself, and subsequent requests are made for each image, plug-in, audio clip, and other content referenced in the HTML. The web requests for each piece of content increase the length of the sequence. These make task-critical requests located far from each other. In general Top-10 accuracy is better than Top-1 accuracy because it is more likely to provide a correct prediction with more candidate events.

**Table 2.** Model comparison. We compare the performance of Markov model, 3-gram model, and our proposed DEEPEVENT model.

Application	Method	Top-1 Accuracy (%)	Top-10 Accuracy (%)
Workqueue	Markov Model	56.30	91.24
	3 Gram Model	61.90	94.17
	<b>DeepEvent</b>	<b>75.21</b>	<b>99.27</b>
DataRepo1	Markov Model	62.59	95.73
	3 Gram Model	74.17	96.76
	<b>DeepEvent</b>	<b>79.53</b>	<b>97.82</b>
DevOpsApp	Markov Model	44.84	80.42
	3 Gram Model	48.01	82.34
	<b>DeepEvent</b>	<b>56.73</b>	<b>89.51</b>
DataAnalyzer1	Markov Model	55.10	86.53
	3 Gram Model	63.07	90.90
	<b>DeepEvent</b>	<b>70.10</b>	<b>95.44</b>
DataAnalyzer2	Markov Model	70.21	95.96
	3 Gram Model	73.21	97.10
	<b>DeepEvent</b>	<b>78.45</b>	<b>97.72</b>
DataRepo2	Markov Model	95.05	99.74
	3 Gram Model	95.56	99.79
	<b>DeepEvent</b>	<b>97.47</b>	<b>99.87</b>

### 4.3 Evaluation of DeepEvent on Anomaly Detection

In this section, we evaluate DEEPEVENT’s capability in distinguishing normal web requests from different types of anomalous requests such as real-world exploits, popular web attacks, and randomly injected HTTP requests. We compare their anomaly scores evaluated by DEEPEVENT with normal requests.

**Real-World Exploits:** We investigate five real-world exploits on DevOpsApp, which is built on top of Jenkins (jenkins.io/): CVE-2016-9299, CVE-2016-0792, CVE-2018-1999001, CVE-2018-1999002, CVE-2019-1003000. We mix the requests generated by the exploits with normal ones and test how DEEPEVENT evaluates the exploits.

**Web Attacks:** We investigate five widely-used attacks against web applications: SQL Injection [27], Cross-site Scripting (XSS) [25], Buffer Overflow [23], CRLF Injection [24], Server-Side Includes (SSI) Injection [26]. We mix the requests produced by those attacks with normal ones in the same way as real-world exploits and test how DEEPEVENT evaluates such requests. We use OWASP Zed Attack Proxy (ZAP) [28], to actively scan and attack web applications. ZAP is

**Table 3.** Anomaly score evaluation for DevOpsApp.

Request type	Average Anomaly score
Normal requests	0.316
CVE-2019-1003000	0.996
CVE-2016-9299	0.996
CVE-2016-0792	0.996
CVE-2018-1999001	0.787
CVE-2018-1999002	0.996
SQL Injection	0.976
Cross-site Scripting	0.975
Buffer Overflow	0.964
CRLF	0.964
SSI Injection	0.964
Random Injection	0.995

**Table 4.** Average anomaly score comparison between normal requests and randomly injected requests.

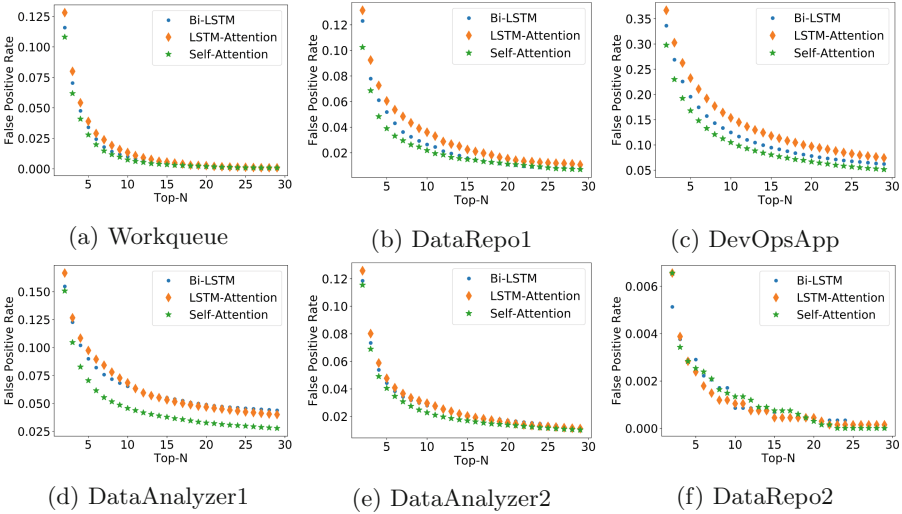
Application	Normal Request	Random Injection
Workqueue	0.151	0.948
DataRepo1	0.132	0.978
DevOpsApp	0.316	0.995
DataAnalyzer1	0.197	0.995
DataAnalyzer2	0.140	0.980
DataRepo2	0.014	0.887

one of the most popular open-source tools for web security and vulnerability assessment. ZAP accesses the web application using a normal user’s credentials. It first crawls all the URIs of the web application and then crafts malicious web requests to exploit the vulnerabilities. We collect all the malicious requests and categorize them by the type of attacks. Then we randomly inject them into normal web requests to evaluate DEEPEVENT.

**Random Injection:** We conduct experiments to test how DEEPEVENT evaluate requests produced by abnormal behaviors, such as requests are generated by normal users but at abnormal occurrences. We simulate arbitrary web requests based on normal requests and inject them into a sequence of normal requests.

Table 3 shows the performance of anomaly detection for DevOpsApp. All real exploits achieve high anomaly scores, compared to normal requests (0.316). CVE-2018-1999001 gains the lowest anomaly score (0.787). Yet, it still has a large distance from normal requests. Four real exploits (CVE-2016-9299, CVE-2016-0792, CVE-2018-1999002, CVE-2019-1003000) were identified as “RARE” events. “RARE” events gain a high anomaly score (0.996), which suggests EventExtractor performs well in extracting web events from HTTP requests.

In our experiments, anomaly scores calculated by DEEPEVENT perform well at distinguishing the requests generated by web attacks from normal requests. In general, our proposed solution can distinguish anomalous requests from the normal request. For instance, the malicious requests and normal requests can be distinguished with a threshold value of 0.9 - most malicious requests have an anomaly score higher than 0.9 while most normal requests have a score lower than 0.9. We observe a few false positives for normal requests due to the low



**Fig. 5.** False Positive Rate (FPR) comparison: we report FPR of Bi-LSTM, and LSTM-attention with different alarm thresholds (Top-N).

frequency of these normal events, which might be further improved with more training data to recognize unexpected normal requests. Some web attacks may gain a slightly high false negative rate (being classified as normal requests) than others. For example, Buffer Overflow has 2.37% false positive rate, because Buffer Overflow usually targets URI parameters as well as payload, which is not covered in this paper.

Table 4 shows the average anomaly score of normal requests and randomly injected requests for six web applications. Normal requests get low anomaly scores (0.158 on average) while randomly injected ones could raise the alarm with an extremely high anomaly score (0.964 on average). The large performance gap indicates that we can set a threshold to differentiate normal requests from randomly injected requests.

#### 4.4 Neural Network Comparison

We compare three different neural networks proposed in context-based modeling in DEEPEVENT: Bi-LSTM, LSTM-attention, and Self-attention. To make a fair and comprehensive comparison, we use the same settings for three neural networks. We evaluate Bi-LSTM and LSTM-attention networks with three different numbers of LSTM layers: 1, 2, 3, and evaluate all three neural networks with five different window sizes: 8, 16, 32, 64, 128 and report the best results.

In Table 5, we report the best results of three neural networks with different settings (window size, number of LSTM layers, with/without pre-training). From the experimental results, we observe that Self-Attention usually achieves the highest accuracy (Top-1, Top-10). In most cases, Bi-LSTM performs second best.

**Table 5.** Comparison between three neural networks.

Application	Neural network	Top-1 accuracy (%)	Top-10 accuracy (%)
Workqueue	Bi-LSTM	74.69	99.00
	LSTM-Attention	73.55	98.66
	Self-Attention	<b>75.21</b>	<b>99.27</b>
DataRepo1	Bi-LSTM	78.20	97.36
	LSTM-Attention	78.61	96.39
	Self-Attention	<b>79.53</b>	<b>97.82</b>
DevOpsApp	Bi-LSTM	53.06	87.52
	LSTM-Attention	51.58	84.63
	Self-Attention	<b>56.73</b>	<b>89.51</b>
DataAnalyzer1	Bi-LSTM	70.84	93.47
	LSTM-Attention	<b>71.51</b>	93.15
	Self-Attention	70.10	<b>95.44</b>
DataAnalyzer2	Bi-LSTM	78.23	97.21
	LSTM-Attention	77.75	97.05
	Self-Attention	<b>78.45</b>	<b>97.72</b>
DataRepo2	Bi-LSTM	97.21	<b>99.91</b>
	LSTM-Attention	97.45	99.90
	Self-Attention	<b>97.47</b>	99.87

Bi-LSTM and LSTM-attention networks need to pass the hidden states through a long path to learn from the long-distance context. The context might be lost in a long path because of gradient vanishing [10]. This may not be suitable for web applications that have task-critical requests located far from each other. On the other hand, the self-attention neural network constructs direct links between requests within the context, which brings great merit in learning from the long-distance context.

To evaluate false positive and false negative rates of the three neural networks, we flag an event as an abnormal event if it is not among the top  $K$  candidate events predicted by the neural network. Correct predictions are considered as true positives. We select the best setting of the three neural networks (Bi-LSTM, LSTM-Attention, and Self-Attention) and calculate their false positive rates based on different alarm thresholds using Top-N, *i.e.*, threshold 10 means if the event is not in Top-10 prediction, it will be labeled as an alarm. As illustrated in Fig. 5, false positive rate decreases when we use a large threshold  $K$ . For the same threshold, Self-Attention based model achieves lower false positive rates than the other two models for all web applications except for DataRepo2. For DataRepo2, all three neural networks achieve extremely low false positive rates, less than 1%.

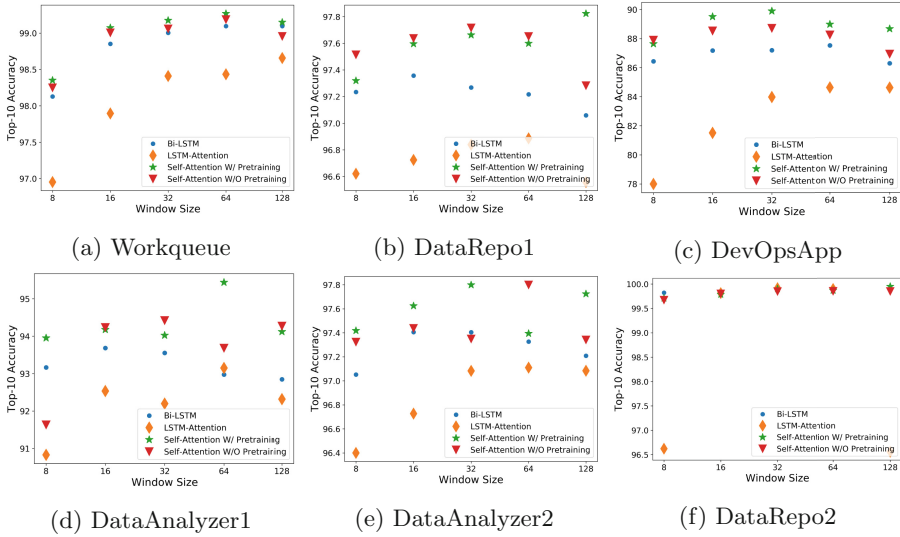


Fig. 6. Model comparison with different settings (window size and pre-training).

### 4.5 Evaluation of Different Model Settings

In this section, we evaluate the impact of different model settings on prediction performance.

**Impact of Window Size.** To evaluate the impact of window sizes, we test five different window sizes of contextual events: 8, 16, 32, 64, 128. As illustrated in Fig. 6, self-attention based model performs better than Bi-LSTM and LSTM-attention based models for longer sequences (*i.e.*, 128). The performance of Bi-LSTM and LSTM-attention based models degrade when the window size becomes 128. Bi-LSTM and LSTM-attention based models have decreasing weights on long-distance events as opposed to short-distance events, which may lead to vanishing gradients for longer-distance context.

**Effectiveness of Pre-training.** Pre-training helps to learn not only the last-event prediction but also the semantic information and relationship of all events among the sequence. We compare the performance of DEEPEVENT with and without pre-training. As shown in Fig. 6, the performance is improved with the proposed pre-training technique for all six applications in general. On average, self-attention with pre-training increases Top-10 prediction accuracy compared to self-attention without pre-training.

Especially, web applications with a large number of unique events may raise more uncertainty and lower accuracy in prediction. Pre-training largely improves the accuracy of prediction for these applications. For instance, the Top-10 accuracy of the three applications with the largest number of unique events (*i.e.*,

**Table 6.** Performance of predicting centered events.

Application	Model	Top-1 accuracy (%)	Top-10 accuracy (%)
Workqueue	Bi-LSTM	97.07	99.79
	LSTM-Attention	97.27	99.81
	Self-Attention	<b>98.33</b>	<b>99.91</b>
DataRepo1	Bi-LSTM	89.35	98.44
	LSTM-Attention	89.78	98.06
	Self-Attention	<b>90.33</b>	<b>99.16</b>
DevOpsApp	Bi-LSTM	74.00	91.72
	LSTM-Attention	71.85	89.41
	Self-Attention	<b>81.25</b>	<b>94.67</b>
DataAnalyzer1	Bi-LSTM	84.31	96.00
	LSTM-Attention	81.70	<b>96.78</b>
	Self-Attention	<b>85.01</b>	96.61
DataAnalyzer2	Bi-LSTM	86.87	97.93
	LSTM-Attention	86.83	97.71
	Self-Attention	<b>87.79</b>	<b>98.57</b>
DataRepo2	Bi-LSTM	96.87	99.87
	LSTM-Attention	97.23	99.90
	Self-Attention	<b>97.96</b>	<b>99.96</b>

DevOpsApp, DataAnalyzer1, and DataAnalyzer2) is greatly increased by applying pre-training models on the self-attention based models.

**Evaluation of Predicting Centered Events.** In the previous experiments, we predicted the last event in a sequence. For many web applications, requests are generated concurrently by a single action. The concurrent requests make it possible for us to leverage contextual events following the event of interest. We study the case where the event of interest to be predicted is centered by contextual events.

Table 6 shows the performance of DEEPEVENT predicting centered events. Comparing Table 5 and 6, we observe that the prediction performance of centered events is improved for all three models in general. For example, for Workqueue, the Top-1 accuracy achieved by self-attention based model increases from 75.21% to 98.33%. Self-attention based model achieves more improvement than Bi-LSTM and LSTM-attention based models when the event of interest is centered by contextual ones. The performance of prediction improves significantly when we predict the centered event instead of the last one. When predicting centered event, events located after the event of interest provide important information. In this way, the model incorporates context from both directions (i.e., left and

right). On average of six applications, DEEPEVENT reduces Top-1 error rate by 52.56% and Top-5 error rate by 57.84% for predicting centered events.

## 5 Related Work

### 5.1 Web Event Forecasting

Web event forecast has been investigated for many years. Su *et al.* extracted access path from server logs and used n-gram models to predict web events for web caching and prefetching [35]. Awad *et al.* analyzed various supervised machine learning approaches for forecasting web events, such as Support Vector Machine, Markov Model and its variant, All-Kth Markov Model [1, 2]. Da *et al.* summarized several clustering and Markov-based approaches for predicting web page access [4]. In this work, we target the enterprise web applications and demonstrates superior performance in forecasting web events compared to the existing approaches.

### 5.2 Web Anomaly Detection

Many statistical models have been used to detect anomaly for web applications [9, 18, 31]. Kruegel *et al.* [18, 19] leveraged statistical models for characterizing HTTP query attributes such as query attribute length, attribute character distribution, and etc. Statistical models output probability values of a query and its individual attributes. The probability values reflect the likelihood of the occurrence with respect to an established profile. Juan *et al.* conducted Kruskal-Wallis and Kolmogorov-Smirnov test on payload length and payload histogram and modeled payload of normal web requests using Markov Chain [9]. Sakib and Huang detected HTTP-based Botnet C&C traffic based on features from web request URLs and DNS responses. Three anomaly detection methods were used in the detection system: Chebyshev's Inequality, One-class SVM, and Nearest Neighbor based Local Outlier Factor. Many supervised machine learning provides have been used to detect anomaly for web applications by providing a binary prediction of normal or abnormal web requests learning from the historical data. Pham *et al.* surveyed different machine learning algorithms such as random forest, logistic regression, decision tree, AdaBoost, and SGD that are used to build Web intrusion detection systems [29]. Oprea *et al.* detected malware in enterprises based on malicious HTTP traffic [22]. They leveraged 89 features extracted from enterprise networks and applied several supervised machine learning algorithms (*e.g.*, logistic regression, decision trees, random forest, and SVM) to learn from these features. Clustering and dimension-reduction are common techniques used in unsupervised learning based solutions [14, 15, 34]. These solutions first extracted features from HTTP GET parameters and URLs, and then used Random Projection (RP), Principal Component Analysis (PCA), and Diffusion Map (DM) to reduce the dimensionality of the data. Clustering algorithms (*e.g.*, K-means) have been applied to identify abnormal behavior.

Zolotukhin *et al.* [38] used several unsupervised learning algorithms such as PCA, K-means, Density-Based Spatial Clustering (DBSCAN) to model URL and User-Agent in HTTP headers and detect anomalies in web requests.

Recently deep learning approaches, in particular RNNs, have been established as state-of-the-art approaches in anomaly detection tasks. Liang *et al.* considered URLs as natural language sequences and applied LSTM and GRU to classify URLs as normal or abnormal requests [20]. Yu *et al.* proposed a neural network consisting of Bidirectional LSTMs and an attention model to extract critical components from URI path and body [37]. Liu *et al.* proposed an attention-based deep neural network, which located the malicious regions from HTTP requests and posts, and classified the malicious HTTP requests [21]. These approaches focus on analyzing the contents in a single web request. We focus on a sequence of web requests, which involves connections among requests and represents users' normal patterns and web application flow characteristics.

### 5.3 Deep Neural Networks for Log Data Analysis

Deep neural networks have been used to analyze log data. Du *et al.* proposed to model a sequence of system logs using LSTM and identified abnormal logs from normal execution [7]. An abnormal event is flagged if such an event is not within top-K probabilities to appear next. Shen *et al.* leveraged RNNs to predict future events based on previous observations using security logs collected from an intrusion prevention system [33]. The work focuses on the prediction of the upcoming security event given a sequence of events. Recently, Recurrent Neural Networks (RNNs) and its variants, Long Short-Term Memory (LSTM) and gated recurrent neural networks [3], have been established as compelling techniques in security analytics research. The RNN based methods analyze the behavior of security event logs or system logs in a session. However, applying these models to web anomaly detection is non-trivial. Logs generated by machines (*e.g.*, heartbeat) are much easier to be detected and predicted compared to web events generated by humans due to human's unpredictable behaviors. To analyze web events, we adapt a self-attention mechanism to learn from the contextual events. With the proposed event and sequence embedding techniques, the adapted self-attention mechanism captures the dependency of long-distance events from human behaviors.

## 6 Conclusion

In this work, we proposed a self-supervised neural network based approach, DEEPEVENT for web event forecasting and anomaly detection. We evaluated DEEPEVENT on web requests collected from real-world enterprise web applications. By connecting web event forecasting with anomaly detection, DEEPEVENT outperformed baseline methods and improved the performance of web event forecasting for complicated web events, while detected anomalies by identifying the most unlikely events in the sequence. We also demonstrated DEEPEVENT's

capability in distinguishing normal web events from different types of anomalous events and measuring their anomaly scores.

## References

1. Awad, M., Khan, L., Thuraisingham, B.: Predicting www surfing using multiple evidence combination. *VLDB J.* **17**(3), 401–417 (2008)
2. Awad, M.A., Khalil, I.: Prediction of user's web-browsing behavior: application of Markov model. *IEEE Trans. Syst. Man Cybern. Part B (Cybern.)* **42**(4), 1131–1142 (2012)
3. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. In: *NIPS Deep Learning and Representation Learning Workshop* (2014)
4. Da, N.T., Hanh, T., Duy, P.H.: A survey of webpage access prediction. In: *2018 International Conference on Advanced Technologies for Communications (ATC)*, pp. 315–320. *IEEE* (2018)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186 (2019)
6. Dong, Y., Zhang, Y.: Adaptively detecting malicious queries in web attacks. *arXiv preprint [arXiv:1701.07774](https://arxiv.org/abs/1701.07774)* (2017)
7. Du, M., Li, F., Zheng, G., Srikumar, V.: DeepLog: anomaly detection and diagnosis from system logs through deep learning. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)* (2017). <https://doi.org/10.1145/3133956.3134015>
8. El-Sayed, M., Ruiz, C., Rundensteiner, E.A.: Fs-Miner: efficient and incremental mining of frequent sequence patterns in web logs. In: *Proceedings of the 6th Annual ACM International Workshop on Web Information and Data Management*, pp. 128–135 (2004)
9. Estevez-Tapiador, J.M., Garcia-Teodoro, P., Diaz-Verdejo, J.: Measuring normality in http traffic for anomaly-based intrusion detection. *Comput. Netw.* **45**, 175–193 (2004)
10. Hochreiter, S., Bengio, Y., Frasconi, P., Schmidhuber, J., et al.: Gradient flow in recurrent nets: the difficulty of learning long-term dependencies (2001)
11. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
12. Japkowicz, N., Stephen, S.: The class imbalance problem: a systematic study. *Intell. Data Anal.* **6**, 429–449 (2002)
13. Jing, L., Tian, Y.: Self-supervised visual feature learning with deep neural networks: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* (2020). <https://ieeexplore.ieee.org/abstract/document/9086055>
14. Juvonen, A., Sipola, T.: Adaptive framework for network traffic classification using dimensionality reduction and clustering. In: *2012 IV International Congress on Ultra Modern Telecommunications and Control Systems*, pp. 274–279. *IEEE* (2012)
15. Juvonen, A., Sipola, T., Häämäläinen, T.: Online anomaly detection using dimensionality reduction techniques for HTTP log analysis. *Comput. Netw.* **91**, 46–56 (2015)

16. Kemeny, J.G., Snell, J.L.: Markov Chains. Springer, New York (1976)
17. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: International Conference for Learning Representations (2015)
18. Kruegel, C., Vigna, G.: Anomaly detection of web-based attacks. In: Proceedings of the 10th ACM Conference on Computer and Communications Security, pp. 251–261. ACM (2003)
19. Kruegel, C., Vigna, G., Robertson, W.: A multi-model approach to the detection of web-based attacks. *Comput. Netw.* **48**(5), 717–738 (2005)
20. Liang, J., Zhao, W., Ye, W.: Anomaly-based web attack detection: a deep learning approach. In: Proceedings of the 2017 VI International Conference on Network, Communication and Computing, pp. 80–85. ACM (2017)
21. Liu, T., Qi, Y., Shi, L., Yan, J.: Locate-then-detect: real-time web attack detection via attention-based deep neural networks. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence, pp. 4725–4731. AAAI Press (2019)
22. Oprea, A., Li, Z., Norris, R., Bowers, K.: MADE: security analytics for enterprise threat detection. In: Proceedings of the 34th Annual Computer Security Applications Conference, pp. 124–136. ACM (2018)
23. OWASP: Buffer overflow. [https://www.owasp.org/index.php/Buffer\\_Overflow](https://www.owasp.org/index.php/Buffer_Overflow)
24. OWASP: Crlf injection. [https://www.owasp.org/index.php/CRLF\\_Injection](https://www.owasp.org/index.php/CRLF_Injection)
25. OWASP: Cross-site scripting (xss). [https://www.owasp.org/index.php/Cross-site\\_Scripting\\_\(XSS\)](https://www.owasp.org/index.php/Cross-site_Scripting_(XSS))
26. OWASP: Server-side includes (ssi) injection. [https://www.owasp.org/index.php/Server-Side\\_Includes\\_\(SSI\)\\_Injection](https://www.owasp.org/index.php/Server-Side_Includes_(SSI)_Injection)
27. OWASP: Sql injection. [https://www.owasp.org/index.php/SQL\\_Injection](https://www.owasp.org/index.php/SQL_Injection)
28. OWASP: Zed attack proxy project. [https://www.owasp.org/index.php/OWASP\\_Zed\\_Attack\\_Proxy\\_Project](https://www.owasp.org/index.php/OWASP_Zed_Attack_Proxy_Project)
29. Pham, T.S., Hoang, T.H., Vu, V.C.: Machine learning techniques for web intrusion detection - a comparison. In: Proceedings of the 8th International Conference on Knowledge and Systems Engineering, pp. 291–297. IEEE (2016)
30. Raina, R., Battle, A., Lee, H., Packer, B., Ng, A.Y.: Self-taught learning: transfer learning from unlabeled data. In: Proceedings of the 24th International Conference on Machine Learning, pp. 759–766. ACM (2007)
31. Robertson, W., Vigna, G., Kruegel, C., Kemmerer, R.A., et al.: Using generalization and characterization techniques in the anomaly-based detection of web attacks. In: NDSS (2006)
32. Sharif, M., Urakawa, J., Christin, N., Kubota, A., Yamada, A.: Predicting impending exposure to malicious content from user behavior. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 1487–1501. ACM (2018)
33. Shen, Y., Mariconti, E., Vervier, P.A., Stringhini, G.: Tiresias: predicting security events through deep learning. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS) (2018). <https://doi.org/10.1145/3243734.3243811>
34. Sipola, T., Juvonen, A., Lehtonen, J.: Anomaly detection from network logs using diffusion maps. In: Iliadis, L., Jayne, C. (eds.) AIAI/EANN -2011. IAICT, vol. 363, pp. 172–181. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-23957-1\\_20](https://doi.org/10.1007/978-3-642-23957-1_20)
35. Su, Z., Yang, Q., Lu, Y., Zhang, H.: WhatNext: a prediction system for web requests using n-gram sequence models. In: Proceedings of the First International Conference on Web Information Systems Engineering, vol. 1, pp. 214–221. IEEE (2000)

36. Yang, Q., Zhang, H.H., Li, T.: Mining web logs for prediction models in www caching and prefetching. In: Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 473–478 (2001)
37. Yu, Y., Yan, H., Guan, H., Zhou, H.: DeepHTTP: semantics-structure model with attention for anomalous HTTP traffic detection and pattern mining. arXiv preprint [arXiv:1810.12751](https://arxiv.org/abs/1810.12751) (2018)
38. Zolotukhin, M., Hämäläinen, T., Kokkonen, T., Siltanen, J.: Analysis of http requests for anomaly detection of web attacks. In: Proceedings of the 12th International Conference on Dependable, Autonomic and Secure Computing. IEEE (2014)
39. Vaswani, A., et al.: Attention is all you need. In: Advances in Neural Information Processing Systems, pp. 5998–6008 (2017)