

SALAD: Self-Adaptive Lightweight Anomaly Detection for Real-time Recurrent Time Series

Ming-Chang Lee

*Dept. of Information Security
and Communication Technology,
Norwegian University of Science
and Technology
Gjøvik, Norway
ming-chang.lee@ntnu.no*

Jia-Chun Lin

*Dept. of Information Security
and Communication Technology,
Norwegian University of Science
and Technology
Gjøvik, Norway
jia-chun.lin@ntnu.no*

Ernst Gunnar Gran

*Dept. of Information Security
and Communication Technology,
Norwegian University of Science
and Technology
Simula Research Laboratory
Gjøvik, Norway
ernst.g.gran@ntnu.no*

Abstract—¹ Real-world time series data often present recurrent or repetitive patterns and it is often generated in real time, such as transportation passenger volume, network traffic, system resource consumption, energy usage, and human gait. Detecting anomalous events based on machine learning approaches in such time series data has been an active research topic in many different areas. However, most machine learning approaches require labeled datasets, offline training, and may suffer from high computation complexity, consequently hindering their applicability. Providing a lightweight self-adaptive approach that does not need offline training in advance and meanwhile is able to detect anomalies in real time could be highly beneficial. Such an approach could be immediately applied and deployed on any commodity machine to provide timely anomaly alerts. To facilitate such an approach, this paper introduces SALAD, which is a **Self-Adaptive Lightweight Anomaly Detection** approach based on a special type of recurrent neural networks called Long Short-Term Memory (LSTM). Instead of using offline training, SALAD converts a target time series into a series of average absolute relative error (AARE) values on the fly and predicts an AARE value for every upcoming data point based on short-term historical AARE values. If the difference between a calculated AARE value and its corresponding forecast AARE value is higher than a self-adaptive detection threshold, the corresponding data point is considered anomalous. Otherwise, the data point is considered normal. Experiments based on two real-world open-source time series datasets demonstrate that SALAD outperforms five other state-of-the-art anomaly detection approaches in terms of detection accuracy. In addition, the results also show that SALAD is lightweight and can be deployed on a commodity machine.

Index Terms—Recurrent time-series anomaly detection, lightweight LSTM, unsupervised learning, real time, self-adaptive threshold

I. INTRODUCTION

A time series is a series of data points that are evenly indexed in time order. Such time series are generated by

various devices such as smart meters, system monitoring tools, IoT devices, etc. It is often the case that time series in the real world reflect a certain pattern that might occur periodically, such as transportation passenger volume, network traffic, click-through rate, energy usage, system resource consumption, and human gait. This kind of time series is referred as recurrent time series in this paper. Finding an anomalous event or abnormal behavior in time series has attracted great attention in recent years since it helps to trigger immediate reaction, prompt troubleshooting and enable appropriate countermeasures to be soon conducted.

During the last decade, a number of methods and approaches based on statistics or machine learning have been proposed for time series anomaly detection in different areas, including intrusion detection, system health monitoring, power demand monitoring, electrocardiogram monitoring, etc. However, statistical-based approaches require to obtain the target time series beforehand so as to generate a statistical model that fits the time series for conducting further anomaly detection. Due to this characteristic, statistical-based approaches might not be suitable to provide real-time anomaly detection on the fly (i.e., along with data collection). One might argue that statistical-based approaches can still be used on the fly. However, these approaches might generate inconsistent anomaly detection results as time goes by, meaning that they might produce different detection results over time (i.e., overwriting and denying some of anomalies that they have claimed previously). On the other hand, most approaches based on machine learning require humans to manually label training data, conduct an offline training process, or determine an appropriate detection threshold in advance. These requirements considerably limit the applicability and usefulness of these approaches in practice.

It is therefore highly valuable if an anomaly detection approach is able to learn the target time series without an offline training/learning process and able to accurately detect anomalies without the above-mentioned requirements. Furthermore, it might also be highly valuable to provide a lightweight anomaly detection approach that is able to be deployed on any commodity machine and able to detect anomalies in real time

¹This is a draft preprint of a paper to be published in the 45th IEEE Computer Society Signature Conference on Computers, Software, and Applications (COMPSAC 2021), IEEE. The final paper may be different from this version. Please use the following citation for this paper: Ming-Chang Lee, Jia-Chun Lin, and Ernst Gunnar Gran, "SALAD: Self-Adaptive Lightweight Anomaly Detection for Real-time Recurrent Time Series," In Proceedings of the 45th IEEE Computer Society Signature Conference on Computers, Software, and Applications (COMPSAC 2021), IEEE.

so that appropriate countermeasures can be undertaken as soon as possible.

To facilitate the abovementioned anomaly detection approach, this paper introduces SALAD, a Self-Adaptive Lightweight Anomaly Detection approach based on a special type of recurrent neural networks called Long Short-Term Memory (LSTM) [1]. SALAD is designed for recurrent time series data, especially for those that is generated in real time. In order to keep computational cost as low as possible, SALAD employs LSTM with a simple network structure (e.g., one hidden layer with only 10 hidden nodes). This kind of LSTM is referred to as *simple LSTM* in this paper. However, according to our experience, using a simple LSTM might not be able to effectively learn a recurrent time series or detect anomalies in the time series if the data pattern of the recurrent time series is complicated.

Therefore, instead of directly training a simple LSTM with the target time series to predict future data points and detect anomalies, SALAD converts the time series into a series of average absolute relative error (AARE) values on the fly. More specifically, SALAD keeps training a simple LSTM with historical data points in a sliding window manner to predict every upcoming data point and calculate the corresponding AARE value. Through the conversion, the resulting AARE series would be much smooth than the original time series, which helps both training and prediction.

When SALAD derives a sufficient number of AARE values, it starts detecting anomalies by training a simple LSTM with short-term historical AARE values to predict an AARE value for every upcoming data point. In the meanwhile, SALAD keeps updating its detection threshold based on all previously derived AARE values and the Three-Sigma Rule [3]. If the difference between a predicted AARE value and its corresponding calculated AARE value is higher than the self-adaptive detection threshold, the corresponding data point is considered anomalous. Otherwise, the data point is considered normal.

To evaluate SALAD, we used two real-world open-source time series datasets to compare SALAD with five state-of-the-art anomaly detection approaches. The results demonstrate that SALAD provides good detection accuracy in terms of precision, recall, and F-score. In addition, SALAD is time-efficient and cost-efficient, and it is able to directly run on a commodity machine. The contributions of this paper are as follows:

- 1) The proposed SALAD is a completely unsupervised anomaly detection approach without a requirement for offline training. It can be easily and immediately applied to a recurrent time series and detect anomalies on the fly without requiring human experts to pre-train a learning model, pre-build a data model, tune parameters, or to determine a detection threshold beforehand.
- 2) SALAD is a self-adaptive and flexible approach. It automatically re-trains its LSTM and adjusts its detection threshold to adapt to pattern changes in the target time

series. These properties make SALAD an ideal anomaly detection approach for real-time recurrent time series.

- 3) SALAD is lightweight and cost-effective. By converting a complex time series into a much simpler AARE series, SALAD successfully avoids using a deep LSTM structure, reduces computation complexity, and makes itself being able to deploy and execute on a commodity machine.

The rest of the paper is organized as follows: Section II presents related work. In Section III, we introduce the details of SALAD. Section IV presents and discusses the experiments and the corresponding results. In Section V, we conclude this paper and outline future work.

II. RELATED WORK

As stated earlier, anomaly detection approaches proposed in recent years could be statistical-based or machine learning-based. Statistical-based approaches work by establishing a statistical model for a given normal dataset and then use the model to check whether a new data point fits this model or not. If the data point has a low probability to be generated from the model, it is considered anomalous. For example, Twitter introduced two anomaly detection algorithms, called AnomalyDetectionTs (ADT for short) and AnomalyDetection-Vec (ADV for short) [3]. ADT is designed to detect statistically significant anomalies in a given time series, whereas ADV is designed to detect statistically significant anomalies in a given vector of observations without timestamp information. Since both ADT and ADV are statistical based, they need to obtain the target time series in advance for establishing the corresponding data model. Therefore, they might not be suitable to detect anomalies in a time series that is generated in real time. Furthermore, both ADT and ADV require human experts to set appropriate values to their parameters in order to achieve good detection performance. For example, it is required to set parameter *max_anoms* (which indicates the maximum number of anomalies to detect as a percentage of the time series) for both ADT and ADV. In addition, for ADV, it is also required to set parameter *period* (which defines the number of data points in a single period, and used during seasonal decomposition).

luminol [5] is proposed by LinkedIn and implemented as an open-source Python library for identifying anomalies in real user monitoring (RUM) data for LinkedIn pages and applications. Given a time series dataset, luminol calculates an anomaly score for each data point in the time series based on statistical analysis. If a data point has a high score, this data point is likely to be anomalous as compared with other data points in the time series. However, luminol relies on human experts to set a threshold to further determine which data points are anomalies.

GrammarViz 3.0 is an anomaly detection approach for recurrent time series developed by Senin et al. [8]. Given a recurrent time series, GrammarViz 3.0 discretizes the time series into a symbolic form, infers a context free grammar, and exploits its hierarchical structure to discover anomalies.

Since GrammarViz 3.0 is statistical-based, it needs to obtain the whole time series in advance for offline analysis and consequently might not be suitable for detecting anomalies in a time series on the fly, i.e., along with the data acquisition. Although the authors have mentioned that GrammarViz 3.0 can process every upcoming data point in time, this approach announces anomalies until all data points in the time series have been processed. Therefore, it might not be able to trigger immediate reaction or enable appropriate countermeasures to be soon conducted. Furthermore, GrammarViz 3.0 requires human experts to set several parameters (e.g., *window size*, *Piecewise Aggregate Approximation size*, and *alphabet size*) for each dataset.

Matrix Profile is a data structure and associated algorithms for time series analysis proposed by Yeh et al. [9]. Matrix Profile is able to discover motifs and find the top- K number of discords where a motif is a repeated pattern in a time series and a discord is an anomaly. In order to use Matrix Profile, users are required to determine a window size, i.e., a user-defined subsequence length as mentioned in [10]. However, users need to decide an appropriate window size since it will significantly impact the execution result of Matrix Profile [10]. Furthermore, users also need to determine the value of K such that Matrix Profile can report the top- K anomalies that it discovers in a given time series. It means that this approach will always report K anomalies it finds regardless of the actual number of anomalies in the time series. However, it would be difficult for users to configure this parameter since users cannot foresee the number of anomalies appear in a time series. Different from Matrix Profile, our proposed (i.e., SALAD) does not require users to determine a value for K since SALAD is not based on the top- K discord approach. The detection threshold used by SALAD is automatically calculated and updated at each time point to adapt to pattern changes in a given time series.

Machine learning approaches are another category of anomaly detection. Most approaches belonging to this category require labelled datasets or/and offline training processes. For examples, Yahoo proposed EGADS [15] to detect anomalies on time series based on a collection of anomaly detection and forecasting models. However, EGADS requires to go on an offline process to model the target time series before it can predict a data value later used by its anomaly detection module and its altering module. Lavin and Ahmad [16] introduced Hierarchical Temporal Memory (HTM) for time series anomaly detection. However, HTM requires 15% of a training dataset to be non-anomalous for training its neural network. Different from EGADS and HTM, our approach SALAD does not have these requirements.

Siffer et al. [17] proposed a time series anomaly detection approach based on Extreme Value Theory. This approach makes no assumption on the distribution of time series and requires no threshold manually set by humans. However, this approach needs a long offline time period to do necessary calibration before it can conduct anomaly detection. Greenhouse [18] is a zero-positive anomaly detection algorithm for time

series based on LSTM. Greenhouse requires all data points in its training datasets to be non-anomalous, making Greenhouse a kind of supervised learning approach. During the training phase, Greenhouse adopts a Look-Back and Predict-Forward strategy to detect anomalies. For a given time point t , a window of most recently observed values of length B is used as *Look-Back* to predict a subsequent window of values of length F as *Predict-Forward*. However, if the training data is not representative, Greenhouse might not be able to capture and accommodate pattern changes in the target time series.

RePAD [2] is a real-time anomaly detection approach for time series. RePAD utilizes a single LSTM model trained with short-term historical data points to predict and detect if any upcoming data point is anomalous or not. According to the experiment results shown in [2], RePAD is able to detect anomalies either proactively or on time, but RePAD suffers from some undesirable false positives. To further reduce the false positive problem, Lee et al. [19] introduce a lightweight real-time ready-to-go anomaly detection approach called ReRe. ReRe employs two lightweight LSTM models with two levels of detection sensitivity to jointly detect if any upcoming data point is anomalous based on short-term historical data points. However, both RePAD and ReRe do not have ability to learn the repetitive pattern in a recurrent time series. In other words, they might not be able to recall the data pattern that they have seen before and might repeatedly generate false anomalies when processing subsequent recurrent time series. This phenomenon will be presented in Section IV. Different from RePAD and ReRe, SALAD does not have this issue since SALAD is designed to learn recurrent time series. SALAD has capability to remember recurrent patterns and detect anomalous data points that are significantly different from the data patterns.

III. METHODOLOGY OF SALAD

When SALAD is launched to detect anomalies in a recurrent time series, SALAD firstly converts the target time series into a series of average absolute relative error (AARE) values. Such a process is called *conversion*. Note that AARE is a well-known measure for determining the prediction accuracy of a forecast approach [21]. Fig. 1 shows the conversion algorithm. Let t be the current time point, and t starts from 0 (i.e., the time point at which SALAD is launched). Let b be the number of data points used to train a simple LSTM model in the conversion process. For example, $b = 100$.

After SALAD is launched at time point 0, it keeps collecting every data point v_t observed at time point t (see line 3 of Fig. 1). When t has advanced to $b-1$, SALAD uses all collected data points (i.e., $v_{t-b+1}, v_{t-b+2}, \dots, v_t$) to train a simple LSTM model (called M_1) and then uses M_1 to predict the value of the next data point, denoted by $\widehat{v_{t+1}}$ (see lines 5 and 6 of Fig. 1). When t further advances to $2b-1$ (see line 7) and SALAD has collected v_b , SALAD is able to calculate an AARE value,

The conversion algorithm	
Input:	Every data point in the target recurrent time series
Output:	A series of AARE values
Procedure:	
1:	While time has advanced {
2:	Let t be the current time point;
3:	Collect data point v_t ;
4:	if $t \geq b - 1$ & $t < 2b - 1$ { // i.e., $99 \leq t < 199$, if $b = 100$
5:	Train a simple LSTM model by using data points $v_{t-b+1}, v_{t-b+2}, \dots, v_t$ as the training data;
6:	Let M_1 be the resulting LSTM model and use M_1 to predict \widehat{v}_{t+1} ;
7:	else if $t \geq 2b - 1$ & $t < 2b + 1$ { // i.e., $199 \leq t < 201$ if $b = 100$
8:	Calculate $AARE_t$ based on Equation 1;
9:	Train another simple LSTM model by using $v_{t-b+1}, v_{t-b+2}, \dots, v_t$ as the training data;
10:	Let M_1 be the resulting LSTM model and use M_1 to predict \widehat{v}_{t+1} ;
11:	else if $t \geq 2b + 1$ { // i.e., $t \geq 201$ if $b = 100$
12:	Calculate $AARE_t$ and thd_{AARE} based on Equations 1 and 2, respectively;
13:	if $AARE_t \leq thd_{AARE}$ {
14:	Keep using M_1 to predict \widehat{v}_{t+1} ; Let A_t be $AARE_t$ and output A_t ;
15:	else {
16:	Train another simple LSTM model by taking $v_{t-b}, v_{t-b+1}, \dots, v_{t-1}$ as the training data;
17:	Replace M_1 with the new trained LSTM model to re-predict \widehat{v}_t ;
18:	Re-calculate $AARE_t$ based on Equation 1; Let A_t be $AARE_t$ and output A_t ;

Fig. 1. The conversion algorithm of SALAD.

denoted by $AARE_t$, based on the following equation:

$$AARE_t = \frac{1}{b} \cdot \sum_{y=t-b+1}^t \frac{|v_y - \widehat{v}_y|}{v_y}, t \geq 2b - 1 \quad (1)$$

where v_y is the observed data point at time point y , and \widehat{v}_y is the forecast data value at y . For instance, if $b = 100$ and $t = 199$, $AARE_{199} = \frac{1}{100} \cdot \sum_{y=100}^{199} \frac{|v_y - \widehat{v}_y|}{v_y}$. A lower AARE value indicates that the predicted value is close to the observed value.

After calculating $AARE_t$, SALAD trains an simple LSTM model with the most recent b data points (i.e., $v_{t-b+1}, v_{t-b+2}, \dots, v_t$), and then uses this model to predict the value of the next data point, denoted by \widehat{v}_{t+1} (see lines 9 and 10). The abovementioned procedure will repeat while $2b - 1 \leq t < 2b + 1$.

As t further advances to $2b + 1$ (see line 11), SALAD calculates $AARE_t$ based on Equation 1 and then calculates a detection threshold called thd_{AARE} based on Equation 2 (see line 12) by considering all previously derived AARE values and following the Three-Sigma Rule [3], commonly used rule for anomaly detection.

$$thd_{AARE} = \mu_{AARE} + 3 \cdot \sigma \quad (2)$$

In Equation 2, $\mu_{AARE} = \frac{1}{t-b-1} \cdot \sum_{x=2b-1}^t AARE_x$, and σ is the standard deviation derived as below.

$$\sigma = \sqrt{\frac{\sum_{x=2b-1}^t (AARE_x - \mu_{AARE})^2}{t - b - 1}} \quad (3)$$

As shown from lines 13 to 14, if $AARE_t$ is smaller than or equal to thd_{AARE} , implying that v_t is similar to previous data points, SALAD uses the same LSTM model (i.e., M_1) to predict the next data point \widehat{v}_{t+1} . In addition, SALAD renames $AARE_t$ as A_t and outputs A_t immediately.

However, if $AARE_t$ is larger than thd_{AARE} (see line 15), implying that the data pattern of the target time series has changed, SALAD trains another simple LSTM model to replace M_1 by taking the most recent b data points (i.e., $v_{t-b}, v_{t-b+1}, \dots, v_{t-1}$), as the training data. The new trained LSTM model is immediately used to re-predict \widehat{v}_t . The purpose

is to adapt to the pattern change and to make further prediction as accurate as possible. After that, the corresponding $AARE_t$ is re-calculated, renamed as A_t , and then outputted (see lines 16 and 18).

Fig. 2 illustrates the detection algorithm of SALAD. When t has advanced to time point $2b+1$, SALAD starts the detection algorithm by inputting every single calibrated AARE value (denoted by A_t) from the conversion algorithm. Before SALAD can officially detect anomalies in the target time series, it needs to obtain three historical AARE values from the conversion algorithm so as to predict the next AARE value. As shown in Fig. 2, the detection algorithm requires to firstly collect A_{2b+1}, A_{2b+2} , and A_{2b+3} at time points $2b+1, 2b+2$, and $2b+3$, respectively. At time point $2b+3$, SALAD is able to train a simple LSTM model by taking A_{2b+1}, A_{2b+2} , and A_{2b+3} as the training data. The trained LSTM model, denoted by M_2 , is then immediately used to predict the AARE value of the next time point, denoted by \widehat{A}_{2b+4} (see lines 3 to 5 of Fig. 2).

The detection algorithm	
Input:	A_t where $t \geq 2b + 1$
Output:	Anomaly notifications
Procedure:	
1:	While $t \geq 2b + 1$ {
2:	Collect A_t from the conversion algorithm;
3:	if $t \geq 2b + 3$ and $t < 2b + 5$ { // i.e., $203 \leq t < 205$, if $b = 100$
4:	Train a simple LSTM model by A_{t-2}, A_{t-1}, A_t as the training data;
5:	Let M_2 be the resulting LSTM model and use M_2 to predict \widehat{A}_{t+1} ;
6:	else if $t \geq 2b + 5$ and $t < 2b + 7$ { // i.e., $t = 205$ or 206 if $b = 100$
7:	Calculate $AARE_t$ based on Equation 4;
8:	Train another simple LSTM model by using A_{t-2}, A_{t-1}, A_t as the training data;
9:	Let M_2 be the resulting LSTM model and use M_2 to predict \widehat{A}_{t+1} ;
10:	else if $t \geq 2b + 7$ { // i.e., $t \geq 207$ if $b = 100$
11:	Calculate $AARE_t$ and thd_{AARE} based on Equations 4 and 5, respectively;
12:	if $AARE_t \leq thd_{AARE}$ {
13:	Report that v_t is <u>not</u> anomalous; Keep using M_2 to predict \widehat{A}_{t+1} ;
14:	else {
15:	Train a simple LSTM model by taking $[A_{t-3}, A_{t-2}, A_{t-1}]$ as the training data;
16:	Use the new trained LSTM model to re-predict \widehat{A}_t ; Re-calculate $AARE_t$ based on Equation 4;
17:	if $AARE_t \leq thd_{AARE}$ {
18:	Report that v_t is <u>not</u> anomalous; Replace M_2 with the new trained LSTM model;
19:	else {
20:	Report that v_t is anomalous; } } }

Fig. 2. The detection algorithm of SALAD.

When time has advanced to $2b + 5$ (see line 6), SALAD calculates $AARE_{2b+5}$ according to the following equation:

$$AARE_t = \frac{1}{3} \cdot \sum_{z=t-2}^t \frac{|A_z - \widehat{A}_z|}{A_z}, t \geq 2b + 5 \quad (4)$$

In Equation 4, z is the index of time where $z \geq 2b + 3$, A_z is the AARE value calibrated and output by the conversion algorithm for time point z , and \widehat{A}_z is the AARE value predicted by M_2 for time point z . After calculating $AARE_{2b+5}$, SALAD trains an LSTM model with the three most recent AARE values, i.e., $A_{2b+3}, A_{2b+4}, A_{2b+5}$ to replace M_2 , and then uses this model to predict the next AARE value, denoted by \widehat{A}_{2b+6} (see lines 7 to 9). When time advances to $2b + 6$, the same procedure (i.e., lines 7 to 9) repeats again.

At time point $2b + 7$ (i.e., $t = 2b + 7$), SALAD is ready to detect anomalies since it has calculated three AARE values, which is sufficient to calculate the detection threshold thd_{AARE} based on Equation 5.

$$thd_{AARE} = \mu_{AARE} + 3 \cdot \sigma_{AARE} \quad (5)$$

where $\mu_{AARE} = \frac{1}{t-2b-4} \cdot \sum_{w=2b+5}^t AARE_w$, and standard deviation σ_{AARE} can be derived as below.

$$\sigma_{AARE} = \sqrt{\frac{\sum_{w=2b+5}^t (AARE_w - \mu_{AARE})^2}{t - 2b - 4}} \quad (6)$$

If $AARE_{2b+7}$ is smaller than or equal to thd_{AARE} (i.e., line 12 holds), it means that A_{2b+7} is similar to previous AARE values. In this case, SALAD does not consider the corresponding data point (i.e., v_{2b+7}) anomalous, and it keeps using the current LSTM model (i.e., M_2) to predict the next AARE value, i.e., A_{2b+8} (Please see lines 12 and 13 of Fig. 2).

However, if $AARE_{2b+7}$ is larger than thd_{AARE} (see line 14), implying that either the data pattern of the AARE series has changed or an anomaly might happen, SALAD trains an LSTM model by taking the three most recent AARE points (i.e., A_{2b+4} , A_{2b+5} , A_{2b+6}) as the training data. After that, SALAD uses the new trained LSTM model to re-predict A_{2b+7} and then re-calculate $AARE_{2b+7}$ (see lines 15 and 16).

If the new $AARE_{2b+7}$ is smaller than or equal to thd_{AARE} (see line 17), SALAD considers that the data pattern of the AARE series has slightly changed and that v_{2b+7} is not anomalous. In this case, SALAD replaces M_2 with this newly trained LSTM model so as to adapt to the pattern change (see line 18). This double-check approach (i.e., retraining LSTM and re-comparing the newly calculated $AARE_{2b+7}$ with thd_{AARE}) enables SALAD to adapt to minor pattern change without being too sensitive.

On the contrary, if the new $AARE_{2b+7}$ is still larger than thd_{AARE} (see line 19), SALAD immediately reports data point v_{2b+7} as an anomaly since the LSTM trained with the most recent AARE values is still unable to well predict A_{2b+7} . The same detection procedure will repeat over and over again to detect any possible anomaly as time further advances.

IV. EXPERIMENT RESULTS

To evaluate SALAD, we compared its performance against the performance of RePAD, ReRe, GrammarViz 3.0, ADT, and ADV by conducting two experiments. Recall that RePAD [2] and ReRe [19] are two state-of-the-art real-time anomaly detection approaches. The reason why they were chosen is that, similar to SALAD, they require neither offline training nor manually determined detection thresholds. GrammarViz 3.0 [8] is a statistical open-source anomaly detection approach for anomalous pattern discovery in recurrent time series. ADT and ADV are two open-source statistical-based approaches introduced by Twitter. Note that we did not evaluate luminol [5] and Matrix Profile [9] since both of them rely on human experts to further set a detection threshold or top- K discords for determining which data points are anomalies. It would be difficult or subjective to select an appropriate detection threshold or an appropriate value for K without trial and error since users have no knowledge of the exact number of anomalies in the target time series.

In the first experiment, a real-world open-source time series dataset called New York City Taxi demand (NYC for short)

TABLE I
TWO REAL-WORLD TIME SERIES DATASETS.

Name	Time Period	The total number of data points	Time Interval
NYC	From 2014-07-01, 00:00 to 2015-01-31, 23:30	10320	30 min
MRT	From 2016-02-01, 05:00 to 2016-02-29, 01:00	609	1 hour

TABLE II
PARAMETER SETTING OF THE TWO EXPERIMENTS.

Approach	Experiment 1 (NYC)	Experiment 2 (MRT)
SALAD	$b=288$	$b=42$
RePAD	$LB = 3$	$LB = 3$
ReRe	$LB = 3$	$LB = 3$
GrammarViz 3.0	$PS=4, AS=4, WS=49$	$PS=4, AS=4, WS=22$
ADT	$max_anoms=0.02$	$max_anoms=0.02$
ADV	$max_anoms=0.02,$ $period=48$	$max_anoms=0.02,$ $period=21$

was chosen from the Numenta Anomaly Benchmark (NAB) [11]. This dataset consists of aggregating the total number of taxi passengers in New York city into 30 minute buckets from July 2014 to January 2015. This dataset has been also used by several prior works, e.g., [12] [13] [14], to evaluate their prediction accuracy. In the second experiment, another real-world time series dataset was chosen from Taipei Mass Rapid Transit (MRT) [22] [23]. This dataset records passenger volume at the Taipei Xindian District Office metro station in February 2016. This dataset is abbreviated as MRT in this paper.

TABLE I provides key information about the two datasets. The interval time between data points in the NYC dataset is 30 minutes, whereas the interval time in the MRT dataset is one hour. Note that there is no record between 2AM and 5AM every day in the MRT dataset. Hence, there are only 21 data points per day, and the total number of data points is therefore 609 ($=21*29$).

TABLE II shows all parameter settings for each of the six abovementioned approaches. When SALAD is employed, we only need to set a value for parameter b . Recall that b indicates the number of data points used to train a simple LSTM model in the conversion algorithm. Since the data pattern of NYC is more complicated than that of MRT, the value of b in the first experiment is set to be larger than that in the second experiment. Recall that both RePAD and ReRe are based on short-term historic data points to detect anomalies, we followed the setting used in [2] [19] [20] (i.e., value 3 for the Look-back parameter, abbreviated as LB in TABLE II) for these two approaches in both experiments.

On the other hand, three parameters are required to set before we can use GrammarViz 3.0. These three parameters are PS (which indicates PAA size), AS (i.e., $alphabet$ size), and WS (i.e., $window$ size). We followed the default settings suggested by [6] to set both PS and AS to 4 in both experiments. Furthermore, parameter WS was intentionally set to 49 and 22 in the first and second experiments, respectively, since

these two values enable GrammarViz 3.0 to achieve its best detection accuracy.

For both ADT and ADV, it is required to set parameter max_anoms (which indicates the maximum number of anomalies to detect as a percentage of the time series). We followed the setting mentioned in [4] and set max_anoms as 0.02. Furthermore, it is also a requirement to set parameter $period$ for ADV. In the first experiment, we set this parameter to 48 (because totally 48 data points were collected per day). In the second experiment, we set this parameter to 21 since totally 21 data points were collected per day. Note that these values were intentionally set since they enable ADV to achieve its best detection accuracy.

TABLE III lists all LSTM hyperparameters used by SALAD, RePAD, and ReRe. Note that all LSTM models utilized by these approaches always have a simple structure: only one hidden layer with only 10 hidden units. With respect to epoch (which is defined as one forward pass and one backward pass of all the training data), it is clear that too many epochs might overfit the training data, whereas too few epochs may underfit the training data. To address this issue, whenever SALAD needs to train/retrain LSTM model, just like RePAD and ReRe, Early Stopping [7] is employed, which is an approach to automatically determine the number of epochs for preventing LSTM from overfitting the training data. In this paper, Early Stopping always chooses a number between 1 and 100 for LSTM M_1 , and a number between 1 and 50 for LSTM M_2 . We can see that the domain range for M_1 is bigger than that for M_2 since the number of data points used to train M_1 is more than the number of data points used to train M_2 . On the other hand, for RePAD and ReRe, Early Stopping always chooses a number between 1 and 50 since LB is both set to 3 in RePAD [2] and ReRe [19].

We used a laptop running MacOS 10.15.6 with 2.6 GHz 6-Core Intel Core i7 and 16GB DDR4 SDRAM to conduct the two experiments. The purpose is to evaluate the performance efficiency of all the six approaches when they are separately deployed on a commodity machine like the abovementioned laptop. In order to evaluate the prediction accuracy for each of these approaches, we calculate precision, recall, and F-score for each of the six approaches in both experiments. Recall that $precision = \frac{TP}{TP+FP}$ and $recall = \frac{TP}{TP+FN}$ where TP is the total number of true positives, FP is the total number of false positives, and FN is the total number of false negatives. F-score is defined as the weighted harmonic mean of the precision as follows:

$$F\text{-score} = 2 \cdot \frac{precision \cdot recall}{precision + recall} \quad (7)$$

The F-score reaches the best value, meaning perfect precision and recall, at a value of 1. The worst F-score would be a value of 0, implying the lowest precision and the lowest recall.

Note that in this paper we did not employ traditional point-wise metrics to measure precision, recall, and F-score since it is acceptable for an algorithm to trigger an alert if the delay of detecting an anomaly is not too long. It is also nice if

TABLE III
LSTM HYPERPARAMETERS.

Approach	SALAD	RePAD	ReRe
The number of hidden layers	1	1	1
The number of hidden units in a hidden layer	10	10	10
The number of epochs	1~100 for LSTM M_1 , 1~50 for LSTM M_2	1~50	1~50

an algorithm is able to proactively detect an anomaly before the anomaly happens. Therefore, we adopted and revised the evaluation method proposed by [24] to provide appropriate and fair comparison.

More specifically, if any anomaly lasting for a period from time point s to time point e can be detected by an approach within the time period from time point $s-X$ to time point $e+X$ (where X is a small number of time points), we say that this approach is able to detect this anomaly, and the time period from $s-X$ to $e+X$ is called Valid Detection Period in the rest of this paper, where $X < s$ and e . According to [24], value 3 is suggested for X if the data points in a dataset are collected every hour. Since the data points in the MRT dataset were collected every hour, so we set X to be 3 in the second experiment. On the other hand, the data points in the NYC dataset were collected every half hour, we set X to be 6 (i.e., $2*3$) in the first experiment.

A. Experiment 1

As mention earlier, dataset NYC was used to evaluate all the six approaches in the first experiment. This dataset contains five anomalies identified by [12] [13]:

- 1) The NYC marathon and Daylight Saving Time Ended (Nov. 2, 2014) [25] [26]
- 2) Thanksgiving (Nov. 27, 2014) [27]
- 3) Christmas (Dec. 25, 2014)
- 4) New Year's Day (Jan. 1, 2014)
- 5) Blizzard (Jan. 26-27, 2015) [28]

In addition, we also noticed other anomalies (i.e., obvious taxi demand change) that were caused by the following holidays and events/activities in NYC on Sundays:

- 1) Independence Day (July 4, 2014)
- 2) Labor Day (Sept. 1, 2014)
- 3) Climate March (Sept. 21, 2014) [29]
- 4) New York Comic Con (Oct. 12, 2014) [30]
- 5) Football match (Nov. 23, 2014) [31]
- 6) Millions March NYC (Dec. 13-14, 2014) [32]
- 7) MLK day (Jan. 19, 2015) [33]

Fig. 3 illustrates each data point in the NYC dataset and the detection results of all the six approaches. All the abovementioned anomalies were highlighted in red. Apparently, each of these anomalies lasted for a period of time. TABLE IV lists the detection accuracy of all the approaches on the dataset.

As we can see from Fig. 3, both RePAD and ReRe are able to detect several anomalies, including the NYC marathon and daylight savings time ended day, the blizzard, and some

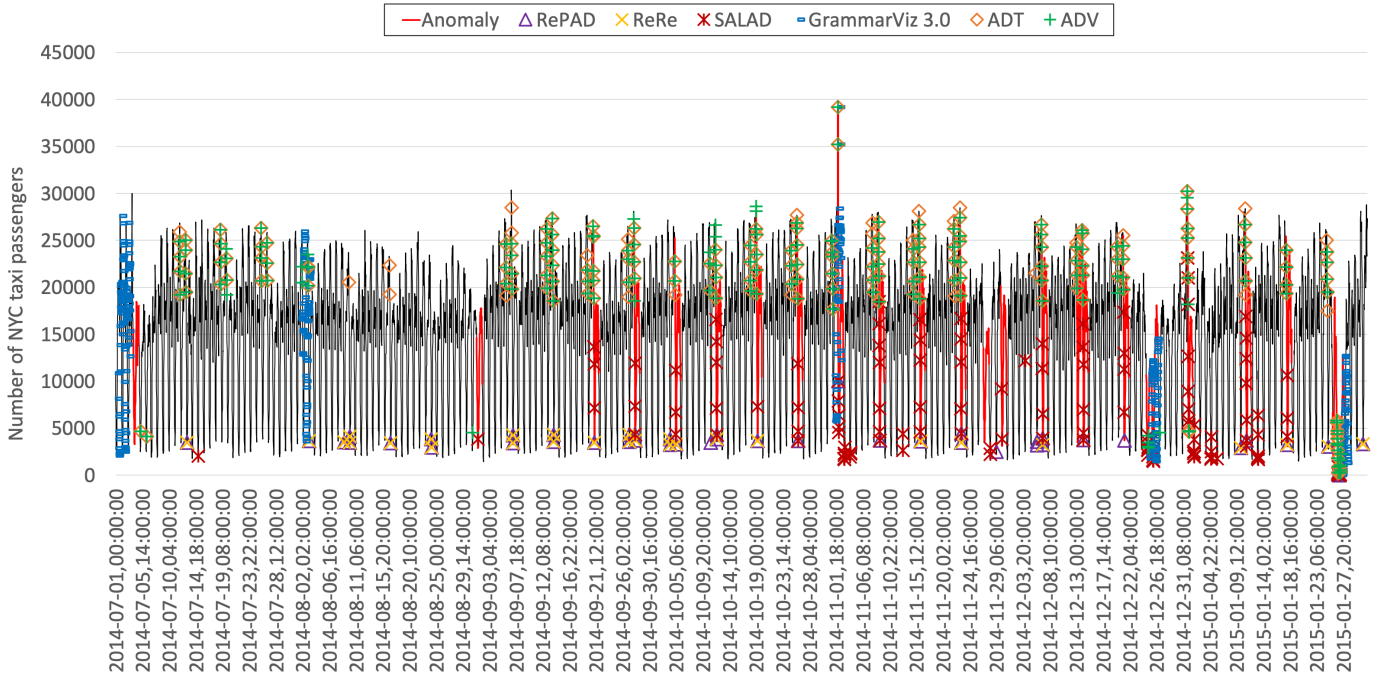


Fig. 3. The detection results of all the six approaches on the NYC dataset in the first experiment.

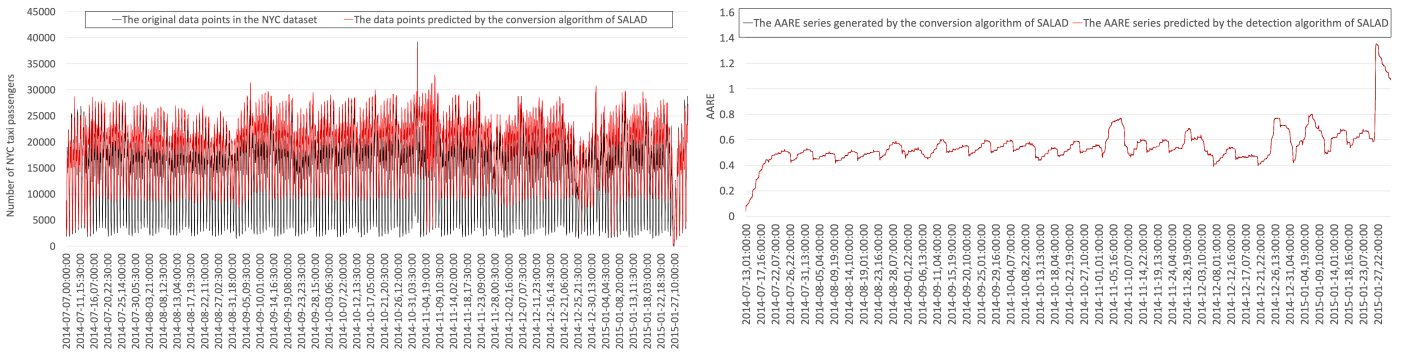


Fig. 4. The original data points in the NYC dataset versus the data points predicted by the conversion algorithm of SALAD.

Fig. 5. The AARE series generated by the conversion algorithm of SALAD versus the AARE series predicted by the detection algorithm of SALAD in the first experiment.

Sunday events. The key reason is that both RePAD and ReRe utilize short-term historical time series data points to predict future data-point values and utilize long-term self-adaptive thresholds to detect anomalies. However, both of them repeatedly made a number of false positives since, as mentioned earlier, they are unable to learn any recurrent data pattern in a recurrent time series. Even though ReRe was designed to reduce the false positive problem of RePAD, ReRe fails to detect more anomalies than RePAD did in this dataset. This is why ReRe has a slightly lower F-score than RePAD.

Recall that we intentionally set three required parameters to GrammarViz 3.0 in order to enable this approach for achieving its best prediction accuracy on the NYC dataset. Even so, GrammarViz 3.0 can only detect three anomalies: The NYC marathon and daylight savings time ended on Nov.

2nd 2014, Christmas on Dec. 25th 2014, and the blizzard between Jan. 26th and 27th 2015. As compared with all the other approaches, GrammarViz 3.0 has the lowest F-score. Recall that both ADT and ADV are statistical-based approaches and they require to analyze the entire NYC dataset before producing their detection results. It is clear that both approaches are able to detect a number of anomalies, but they also introduce some false positives and false negatives.

Among all the approaches, it is clear that SALAD achieves the highest precision, recall, and F-score, implying that SALAD provides the best prediction accuracy. The reason behind this is mainly due to the employed conversion and detection algorithms. Recall that the conversion algorithm of SALAD aims to convert a complex time series into a less complex AARE series by predicting a value for each

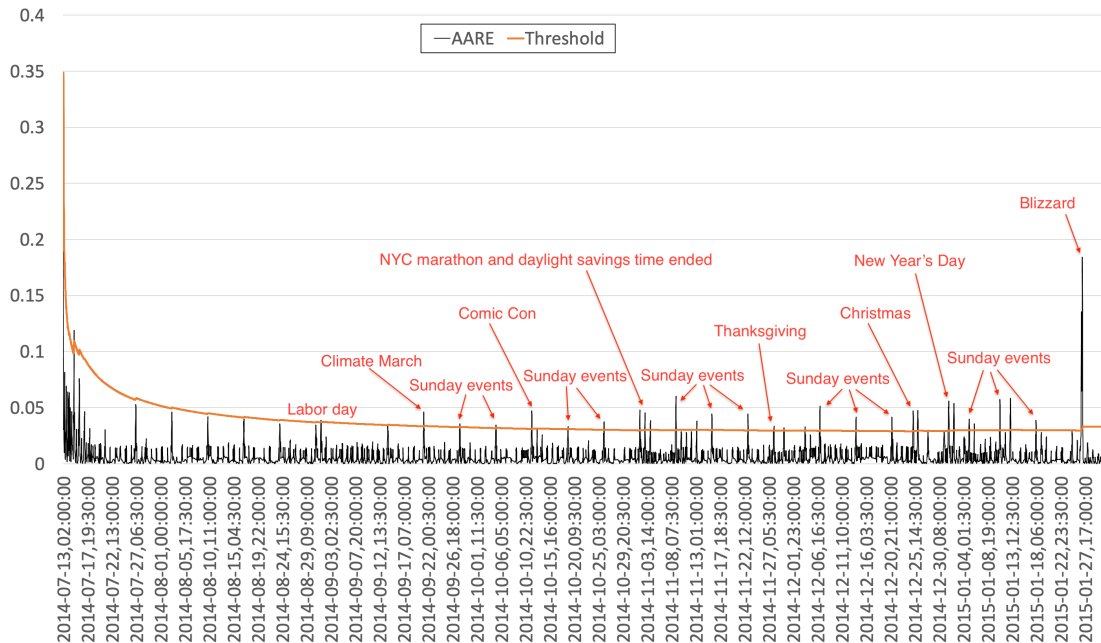


Fig. 6. The AARE values generated by the detection algorithm of SALAD versus the self-adaptive detection threshold in the first experiment.

future data point, measuring the difference between each pair of predicted and actual data points, and calculating the corresponding AARE values. As we can see from Fig. 4 that not all data points predicted by the conversion algorithm of SALAD perfectly match the corresponding real data points but we can see that the AARE series generated by the conversion algorithm (see the black line in Fig. 5) is much less complex than the original NYC time series (see the black line in Fig. 4). Such a conversion facilitates the detection algorithm of SALAD to learn the dataset and make correct predictions with a simple LSTM model.

Due to the above reason, we can see from Fig. 5 that the AARE series predicted by the detection algorithm of SALAD is similar to the AARE series generated by the conversion algorithm. In fact, all the resulting AARE values between the two AARE series are presented in Fig. 6. Apparently, most values are quite low, and those AARE values that are higher than the dynamically self-adaptive detection threshold are considered anomalies by SALAD. Since the threshold considers a good overview of the all previous AARE values, it is able to capture all anomalies (except the thanksgiving one on Nov. 27th, 2014) without generating many false positives or false negatives.

Since SALAD, RePAD, and ReRe are designed to be real-time anomaly detection approaches, it is also important and necessary to compare their individual time efficiency of detecting anomalies. Note that GrammarViz 3.0, ADT, and ADV are statistical-based and offline-based approaches, such an evaluation and comparison cannot be applied to them. TABLE V shows the average time taken by SALAD, RePAD, and ReRe to determine if a data point is anomalous or not. Note that the detection time for each data point includes

both the corresponding LSTM retraining time (if retraining is necessary) and the corresponding detection time. Apparently, RePAD consumes the shortest average detection time with the lowest standard deviation. This is because that RePAD is based on a simple LSTM model trained with short-term historical data points. The average detection time of ReRe is slighter higher than RePAD since ReRe employs two simple LSTM models for anomaly detection.

The average detection time required by SALAD is around two seconds longer than those of RePAD and ReRe. The key reason is that SALAD needs to convert the NYC time series into the corresponding AARE series and then conducts the detection algorithm to detect anomalies. During these two phases, SALAD needs to train its simple LSTM model and might need to re-train its LSTM model in order to adapt to pattern changes. Nevertheless, we can see that SALAD is still time-efficient and lightweight even though it is just deployed on a regular laptop. The short response time helps trigger immediate reaction or countermeasures for handling detected anomalies.

TABLE IV
DETECTION ACCURACY ON THE NYC DATASET.

Approach	Precision	Recall	F-score
SALAD	0.978	0.957	0.967
RePAD	0.973	0.739	0.840
ReRe	0.976	0.727	0.833
GrammarViz 3.0	0.548	0.222	0.316
ADT	0.897	0.826	0.860
ADV	0.911	0.870	0.890

TABLE V
AVERAGE DETECTION TIME ON THE NYC DATASET.

Approach	Average Detection Time	Standard Deviation
SALAD	2.072 sec	4.145 sec
RePAD	0.021 sec	0.025 sec
ReRe	0.041 sec	0.047 sec

B. Experiment 2

As stated earlier, we used the MRT dataset in the second experiment. This dataset is only one month long, and all its data points are illustrated in Fig. 7 (see the black line). This dataset contains only one anomaly. The anomaly is from Feb. 6th to Feb. 14th (i.e., 9 days) during which is Chinese New Year in 2016. The reason why this period was considered as an anomaly is due to the fact that the passenger volume taking Taipei MRT in this time period was apparently less than the rest days in the same month.

Fig. 7 plots the detection results of all the approaches on this dataset except ADT. In the second experiment, ADT was unable to execute since it requires a dataset to be about four times longer than the length of MRT. TABLE VI also lists the detection accuracy of all the five approaches on the dataset. According to the defined Valid Detection Period, all of the approaches are able to detect the Chinese-New-Year anomaly. However, as we can see from Fig. 7, it is clear that SALAD is the first one that detects the anomaly, followed by ADV, GrammarViz 3.0, RePAD, and ReRe.

In addition, SALAD outperforms the other four approaches when it comes to precision, recall, and of course F-score. SALAD did not make any false positives or false negatives. Therefore, its F-score reaches the best value of 1. The reason for such a great detection performance is as follows: SALAD successfully converts the MRT time series (as shown in Fig. 8) into a less complex AARE series (as shown in Fig. 9). Such an AARE series enables SALAD to easily predict using its simple LSTM structure without having to use a more complicated LSTM structure. As we can see from Fig. 10, the AARE values calculated by the detection algorithm of SALAD together with the self-adaptive detection threshold enable SALAD to easily detect the Chinese-New-Year anomaly. It is also clear that the AARE values associated with all other data points at different time points did not exceed the detection threshold, so those data points were not considered as anomalous by SALAD.

From Fig. 7 and TABLE VI, we can see that RePAD and ReRe have the same results in F-score, and that they generated the same number of false positives. Apparently, RePAD and ReRe are unable to perfectly learn the recurrent data pattern of the MRT dataset by using short-term historical data points (i.e., $LB = 3$ [20]) to determine the anomalousness of each data point.

Recall that the Chinese-New-Year anomaly lasted for 9 days in the entire MRT dataset, meaning that around 31% of data points are anomalous. This kind of the dataset introduces some detection difficulty for GrammarViz 3.0 since this approach

tends to consider infrequent and irregular data points as anomalies. Due to this reason, GrammarViz 3.0 made some false positives even though it was able to detect the Chinese-New-Year anomaly in the valid detection period.

When ADV was employed, we can see from Fig. 7 that it also considered some weekends (e.g., 21th, 28th, and 29th of the February) as anomalies since this approach is statistical based. It seems unable to differentiate the data pattern in these weekends and that in the Chinese-New-Year period.

TABLE VII shows the average detection time of SALAD, RePAD, and ReRe on the MRT dataset. Due to the same reason mentioned earlier, the average detection time of SALAD is slightly higher than those of RePAD and ReRe. Nevertheless, SALAD still offers time-efficient and lightweight detection performance due to its ability of converting the MRT dataset into a less complex AARE series and its ability to predict AARE values by using a simple LSTM model, rather than employing deep network structures.

TABLE VI
DETECTION ACCURACY ON THE MRT DATASET.

Approach	Precision	Recall	F-score
SALAD	1	1	1
RePAD	0.913	1	0.955
ReRe	0.913	1	0.955
GrammarViz 3.0	0.762	1	0.865
ADT	N/A	N/A	N/A
ADV	0.979	1	0.989

TABLE VII
AVERAGE DETECTION TIME ON THE MRT DATASET.

Approach	Average Detection Time	Standard Deviation
SALAD	0.109 sec	0.174 sec
RePAD	0.031 sec	0.064 sec
ReRe	0.065 sec	0.100 sec

V. CONCLUSION AND FUTURE WORK

In this paper, we have introduced SALAD, which is designed to detect anomalies in a recurrent time series on the fly in real time. SALAD does not need to go through an offline training process to pre-train a learning model, nor require human experts to tune several parameters or determine a detection threshold in advance. By converting a complex time series into a less complex AARE series, SALAD is able to effectively and efficiently conduct anomaly detection using a simple network structure without having to employ a deep network structure, thereby reducing computation complexity and enabling itself to be deployed and executed on a commodity machine. Based on the self-adaptive detection threshold that is automatically derived and updated at each time point, SALAD does not suffer the issue of top- K discords or requires users to decide an appropriate detection threshold beforehand.

The experiments based on two real-world open-source time series datasets demonstrate that SALAD provides better detection accuracy than the other five state-of-the-art anomaly

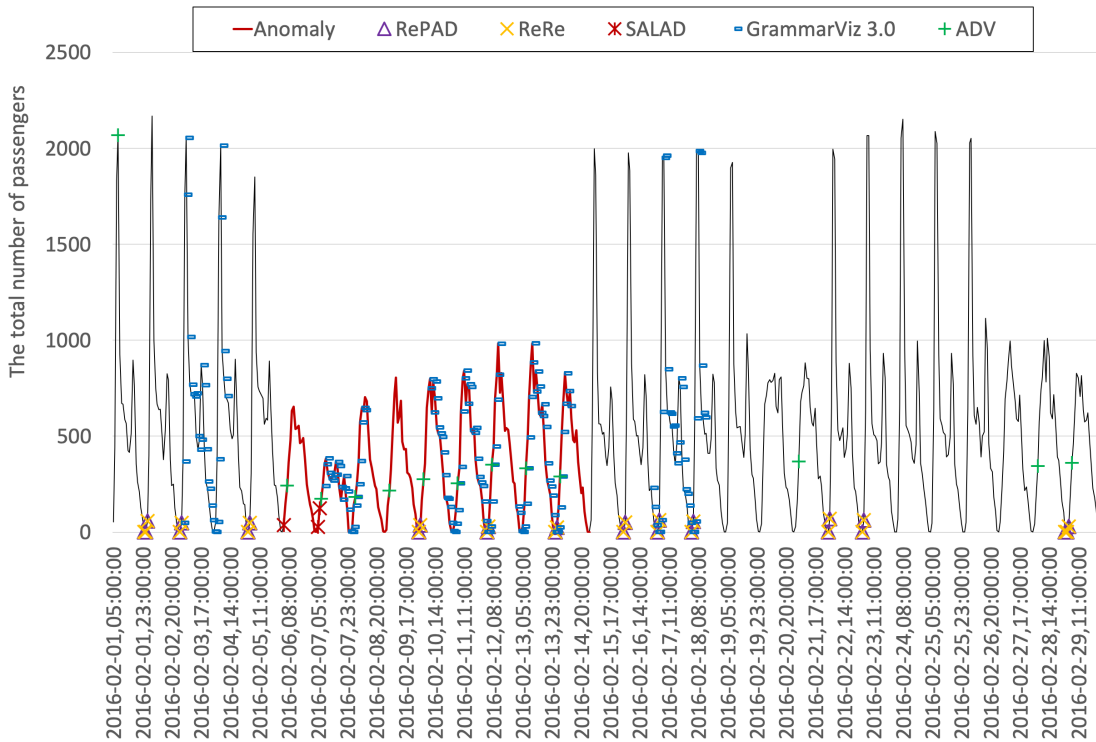


Fig. 7. The detection results of all the approaches (except ADT) on the MRT dataset in the second experiment.

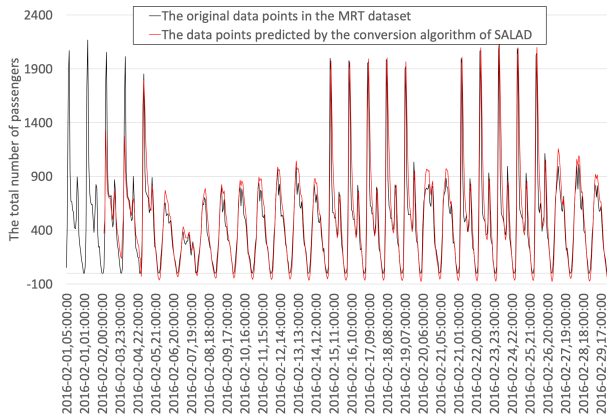


Fig. 8. The original data points in the MRT dataset versus the data points predicted by the conversion algorithm of SALAD.

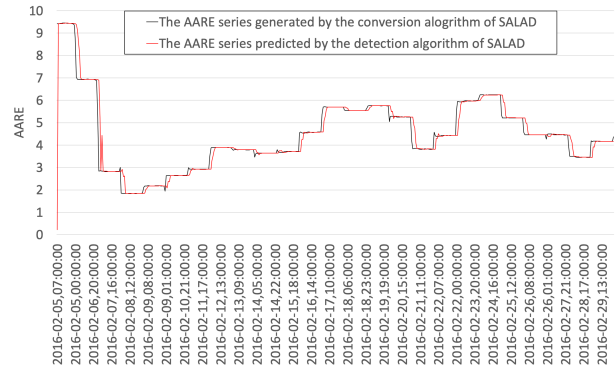


Fig. 9. The AARE series generated by the conversion algorithm of SALAD versus the AARE series predicted by the detection algorithm of SALAD in the second experiment.

detection approaches. In addition, SALAD does not introduce significant computation overhead, which makes it a practical solution to be deployed on a commodity machine in real scenarios.

As future work, we would like to study an approach that is able to automatically determine an appropriate value for parameter b such that it suits any given recurrent time series. In addition, we plan to introduce an approach that is able to automatically recognize whether a given time series is recurrent or not on the fly and then use an appropriate anomaly detection algorithm. Furthermore, we would also like to extend and deploy SALAD on the eX³ HPC cluster [35] to further

handle large-scale time series in a parallel and distributed way, by referring to [36] [37].

ACKNOWLEDGMENT

This work was supported by the project eX³ - *Experimental Infra-structure for Exploration of Exascale Computing* funded by the Research Council of Norway under contract 270053 and the scholarship under project number 80430060 supported by Norwegian University of Science and Technology.

REFERENCES

- [1] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

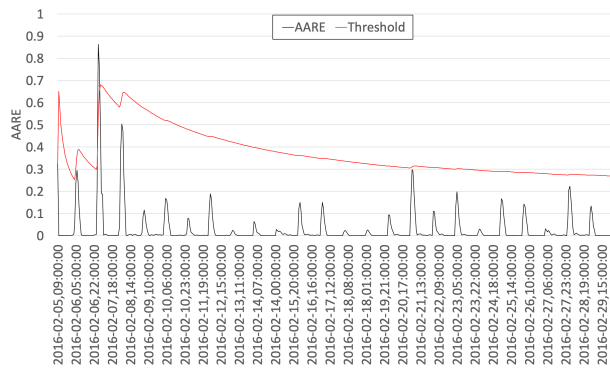


Fig. 10. The AARE values generated by the detection algorithm of SALAD versus the self-adaptive detection threshold in the second experiment.

- [2] M.-C. Lee, J.-C. Lin, and E. G. Gran, "RePAD: real-time proactive anomaly detection for time series," In Proceedings of the 34th International Conference on Advanced Information Networking and Applications (AINA 2020), pp. 1291–1302, 2020. arXiv preprint arXiv:2001.08922
- [3] J. Hochenbaum, O.S. Vallis, and A. Kejariwal, "Automatic anomaly detection in the cloud via statistical learning," 2017. arXiv preprint arXiv:1704.07706
- [4] Twitter/AnomalyDetection [Online code repository], Available: <https://github.com/twitter/AnomalyDetection> (Last Accessed: May 15th 2021)
- [5] linkedin/luminol [Online code repository], Available: <https://github.com/linkedin/luminol> (Last Accessed: May 15th 2021)
- [6] GrammarViz 3.0: Interactive discovery of variable-length time series patterns, Available: https://github.com/GrammarViz2/grammarviz2_src (Last Accessed: May 15th 2021)
- [7] What is early stopping?, Available: <https://deeplearning4j.konduit.ai/tuning-and-training/early-stopping> (Last Accessed: May 15th 2021)
- [8] P. Senin, et al., "Grammarviz 3.0: Interactive discovery of variable-length time series patterns," ACM Transactions on Knowledge Discovery from Data (TKDD), vol. 12, no. 1, pp. 1–28, 2018.
- [9] C.-C. M. Yeh, et al., "Matrix profile I: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets," IEEE 16th international conference on data mining (ICDM 2016), 2016, pp. 1317–1322.
- [10] Introduction to Matrix Profiles, Available: <https://towardsdatascience.com/introduction-to-matrix-profiles-5568f3375d90> (Last Accessed: May 15th 2021)
- [11] numenta/NAB [Online code repository], Available: <https://github.com/numenta/NAB> (Last Accessed: May 15th 2021)
- [12] A Tour of AI Technologies in Time Series Prediction, Available: <https://www.soa.org/resources/research-reports/2019/tour-ai-technologies/> (Last Accessed: May 15th 2021)
- [13] NYC Taxi Analysis - Anomalies, Available: https://matrixprofile.docs.matrixprofile.org/examples/NYC_Taxis.html (Last Accessed: May 15th 2021)
- [14] I.V. Farahani, A. Chien, R.E. King, M.G. Kay, and B. Klensz, "Time Series Anomaly Detection from a Markov Chain Perspective," IEEE 18th International Conference On Machine Learning And Applications (ICMLA 2019), 2019, pp. 1000–1007.
- [15] N. Laptev, S. Amizadeh, and I. Flint, "Generic and scalable framework for automated time-series anomaly detection," Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining, 2015, pp. 1939–1947.
- [16] A. Iavin and S. Ahmad, "Evaluating Real-Time Anomaly Detection Algorithms—The Numenta Anomaly Benchmark," 2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA 2015), 2015, pp. 38–44.
- [17] A. Siffer, P.-A. Fouque, A. Termier, and C. LARGOUET, "Anomaly detection in streams with extreme value theory," Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2017, pp. 1067–1075.
- [18] T.J. Lee, J.Gottschlich, N. Tatbul, E. Metcalf, and S. Zdonik, "Greenhouse: a zero-positive machine learning system for time-series anomaly detection," 2018. arXiv preprint arXiv:1801.03168
- [19] M.-C. Lee, J.-C. Lin, and E. G. Gran, "ReRe: A Lightweight Real-time Ready-to-Go Anomaly Detection Approach for Time Series," In Proceedings of the 44th IEEE Computer Society Signature Conference on Computers, Software, and Applications (COMPSAC 2020), pp. 322–327, 2020. arXiv preprint arXiv:2004.02319
- [20] M.-C. Lee, J.-C. Lin, and E. G. Gran, "How Far Should We Look Back to Achieve Effective Real-Time Time-Series Anomaly Detection?" In Proceedings of the 35th International Conference on Advanced Information Networking and Applications (AINA 2021), pp. 1–13, 2021. arXiv preprint arXiv:2102.06560
- [21] M.-C. Lee, J.-C. Lin, and E. G. Gran, "Distributed Fine-Grained Traffic Speed Prediction for Large-Scale Transportation Networks based on Automatic LSTM Customization and Sharing," In Proceedings of the 26th International European Conference on Parallel and Distributed Computing (Euro-PAR 2020), vol. 12247, pp. 234–247, 2020. arXiv preprint arXiv:2005.04788
- [22] T. Nakamura, M. Imamura, R. Mercer, and E. Keogh, "MERLIN: Parameter-Free Discovery of Arbitrary Length Anomalies in Massive Time Series Archives," Proceedings of the 20th IEEE International Conference on Data Mining, 2020.
- [23] C.-C. M. Yeh, et al., "Online Amnesic Dynamic Time Warping to Allow Real-Time Golden Batch Monitoring," 2019.
- [24] H. Ren, et al., "Time-Series Anomaly Detection Service at Microsoft," Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 3009–3017, 2019.
- [25] 2014 New York City Marathon, Available: https://en.wikipedia.org/wiki/2014_New_York_City_Marathon (Last Accessed: May 15th 2021)
- [26] 2014 Daylight Saving Time Ended, Available: <https://www.timeanddate.com/time/change/usa/new-york?year=2014> (Last Accessed: May 15th 2021)
- [27] Thanksgiving in 2014 is on Thursday, the 27th of November. Available: <https://www.when-is.com/thanksgiving-2014.asp> (Last Accessed: May 15th 2021)
- [28] The Blizzard of January 26-27, 2015. Available: https://www.weather.gov/okx/Blizzard_01262715 (Last Accessed: May 15th 2021)
- [29] People's Climate March (PCM) on September 21, 2014, Available: https://en.wikipedia.org/wiki/2014_People%27s_Climate_March (Last Accessed: May 15th 2021)
- [30] New York Comic Con on October 9-12, 2014, Available: https://en.wikipedia.org/wiki/New_York_Comic_Con (Last Accessed: May 15th 2021)
- [31] Dallas Cowboys at New York Giants - November 23rd, 2014, Available: <https://www.pro-football-reference.com/boxscores/201411230nyg.htm> (Last Accessed: May 15th 2021)
- [32] Millions March NYC, Available: <https://www.businessinsider.com/millions-march-nyc-2014-12?r=US&IR=T> (Last Accessed: May 15th 2021)
- [33] Is Martin Luther King Jr. Day a Public Holiday?, Available: <https://www.timeanddate.com/holidays/us/martin-luther-king-day> (Last Accessed: May 15th 2021)
- [34] Intelligence Processing Unit, Available: <https://www.graphcore.ai/products/ipu> (Last Accessed: May 15th 2021)
- [35] The eX³ Research Infrastructure, Available: <https://www.ex3.simula.no> (Last Accessed: May 15th 2021)
- [36] M.-C. Lee, J.-C. Lin, and R. Yahyapour, "Hybrid job-driven scheduling for virtual MapReduce clusters," IEEE Transactions on Parallel and Distributed Systems, vol. 27, no. 6, pp.1687–1699, June 2016. arXiv preprint arXiv:1808.08040
- [37] J.-C. Lin and M.-C. Lee, "Performance evaluation of job schedulers on Hadoop YARN," Concurrency and Computation: Practice and Experience, vol. 28, no. 9, pp.2711–2728, 2016. arXiv preprint arXiv:1808.08289