No Free Lunch But A Cheaper Supper: A General Framework for Streaming Anomaly Detection

Ece Calikus (Corresponding author)

Center for Applied Intelligent Systems Research
Halmstad University,Sweden
ece.calikus@hh.se

Sławomir Nowaczyk

Center for Applied Intelligent Systems Research
Halmstad University,Sweden
slawomir.nowaczyk@hh.se

Anita Sant'Anna

Center for Applied Intelligent Systems Research
Halmstad University,Sweden
anita.santanna@hh.se

Onur Dikmen

Center for Applied Intelligent Systems Research
Halmstad University,Sweden
onur.dikmen@hh.se

Abstract

In recent years, research interest in detecting anomalies in temporal streaming data has increased significantly. A variety of algorithms are being developed in the data mining community. They can be broadly divided into two categories, namely general-purpose and ad hoc ones. In most cases, general approaches assume a one-size-fits-all solution model, and strive to design a single "optimal" anomaly detector which can detect all anomalies in any domain. To date, there exists no universal method that has been shown to outperform the others across different anomaly types, use cases and datasets. In this paper, we propose SA-FARI, a framework created by abstracting and unifying the fundamental tasks within the streaming anomaly detection. SAFARI provides a flexible and extensible anomaly detection procedure to overcome the limitations of one-size-fits-all solutions. Such abstraction helps to facilitate more elaborate algorithm comparisons by allowing us to isolate the effects of shared and unique characteristics of diverse algorithms on the performance. Using the framework, we have identified a research gap that motivated us to propose a novel learning strategy. We implemented twenty different anomaly detectors and conducted an extensive evaluation study, comparing their performances using real-world benchmark datasets with different properties. The results indicate that there is no single superior detector which works perfectly for every case, proving our hypothesis that "there is no free lunch" in the streaming anomaly detection world. Finally, we discuss the benefits and drawbacks of each method in-depth, drawing a set of conclusions and guidelines to guide future users of SAFARI.

1 Introduction

Anomaly detection is the problem of identifying data points, or patterns, that do not conform to the expected behavior. Anomalies correspond to (often critical) actionable information in many real-world applications, including condition monitoring, intrusion detection, fault prevention, fraud detection, and so on, across various domains such as production, finance, security, medicine, energy, and social media. In recent years, technological advances have facilitated the ability to collect large volumes of data from streams that are produced by various sensors over time. Therefore, detecting anomalies in such continuously changing temporal data has received increasing attention from both the industry and the scientific community.

However, anomaly detection in data streams is a difficult task, since it combines both the challenges associated with anomaly detection and those associated with learning from streaming data. For example, the former requires defining the exact notion of normal behavior, while the latter includes the difficulty of learning the dynamic nature of such behavior when it evolves over time. Among the many approaches currently proposed in the anomaly detection literature, one can distinguish two categories of methods: general or ad hoc. The "general" approaches aspire to detect anomalies independently of the use case and propose a single algorithm supposedly outperforming all previous ones in terms of detection accuracy. However, anomaly detection is an inherently subjective task, in which the characteristics of data and the notion of anomaly vary greatly across applications. One algorithm may perfectly capture the structure of normal behavior in one dataset but may not work at all in another dataset. Several studies show that there is no single anomaly detector that is ultimately superior in all cases (Aggarwal and Sathe, 2017; Campos et al., 2016; Emmott et al., 2015). Clearly, there is no free lunch for anomaly detection. This fact motivates the need for developing a collection of algorithms instead of seeking for the "one" that is suitable all the time.

Ad hoc approaches, on the other hand, are specifically tailored to their target application and are often designed based on complex criteria that require deep domain expertise. Even within the same domain, however, there are often different situations and circumstances where the requirements for a particular task may change. For example, carefully crafted features may become irrelevant, or the current metric to measure deviations in a specific use case may not be suitable in a new scenario. In such cases, making the necessary adaptations to the existing algorithm often amounts to redoing most of the work from scratch.

In this paper, we propose SAFARI, a meta-framework that makes it easy to create different unsupervised anomaly detectors adapted to a particular time-evolving streaming data. This framework provides a generalized procedure for streaming anomaly detection, with separate components that address the fundamental tasks of this problem as separate concerns. The proposed framework is flexible and extensible, since new methods can be easily integrated into existing framework components, to then be mixed and matched for building specific anomaly detectors. Furthermore, the "loosely coupled", modularized structure offers a higher degree of freedom for algorithm adaptations, as the properties of each component can be modified separately without the need for updating the other parts of the framework.

Additionally, the existing evaluation strategies do not provide thorough understanding and comparison of proposed algorithms. Most published experiments evaluate their algorithms by reporting performance scores on application-specific case studies or synthetic datasets. They attempt to assess the effectiveness of algorithms without characterizing the nature of the anomalies in the datasets, nor other factors that influence the performance, such as noise or concept drift. Often the methods to be compared share common properties, but it is challenging to analyze their effects on performance because those properties are hidden in the design of the algorithm and are not trivial to isolate. It is not obvious how to interpret whether an algorithm performs better because of, for example, a specific distance function, the sampling strategy, or the features that it uses. This makes it difficult to answer the question *which anomaly detection algorithm should be chosen for a specific scenario?*

By unifying and separating key concepts in existing methods, our framework allows one to study commonalities and differences of the algorithms more thoroughly, leading to more elaborate algorithm comparisons. In this work, we integrate several different methods into the framework and evaluate their performance under varying circumstances. We present how the performances of different combinations vary depending on characteristics of datasets. We also discuss the advantages and drawbacks of each method separately, to guide readers on how to effectively combine building blocks for specific scenarios. In the end, we compare the performances of our framework to those of selected state-of-the-art methods.

Finally, the vast quantity and diversity of existing approaches, as well as difficulty in having an overview of their actual properties, make it challenging to identify the gaps in the state-of-the-art. Our framework helps decomposing many existing approaches in a uniform manner and discovering their essential characteristics. By formalizing state-of-the-art anomaly detection methods within the SAFARI framework, we have determined there is no existing general approach to learning data streams for the case of anomaly detection. Therefore, we propose a novel learning strategy by generalizing the weighted reservoir-sampling schema considering the constraints of the anomaly detection problem.

Our contributions can be summarized as follows:

- We conceptualize the four high-level fundamental tasks in streaming anomaly detection problem and formulate a meta-framework that is built upon these essential concepts to provide general, flexible, and adaptable detection procedure.

- By integrating different methods into the framework's components, we implement 20 different anomaly detectors, several of which are novel approaches that have not been tried before.

- With the help of the framework, we identify a gap in existing data stream learning strategies and propose a novel anomaly-aware reservoir sampling scheme.

- We conduct an extensive comparison study on these approaches using two benchmarks (Numenta and Yahoo) that contain various real-world and synthetic time-series datasets from different domains.

The remainder of this paper is organized as follows. In Section 2, we present fundamental concepts in streaming anomaly detection and introduce our framework. We review existing work in Section 3 in the light of concepts introduced in the previous section. In Section 4, we describe in detail the methods we have integrated into the framework implementation. In Sections 5 and 6, we discuss our experiments and results, respectively. We highlight the main observations and recommendations from the results in Section 7 and conclude this study in Section 8.

2   Framework

In this section, we present our meta-framework SAFARI (Streaming Anomaly Detection Framework using Reference Instances), which combines fundamental tasks abstracted from existing anomaly detection algorithms into a united schema and provides a generic procedure for streaming anomaly detection. SAFARI is essentially based on the concept of a **reference group**, which consists a set of instances that is assumed to represent the current normal behavior of the data stream.
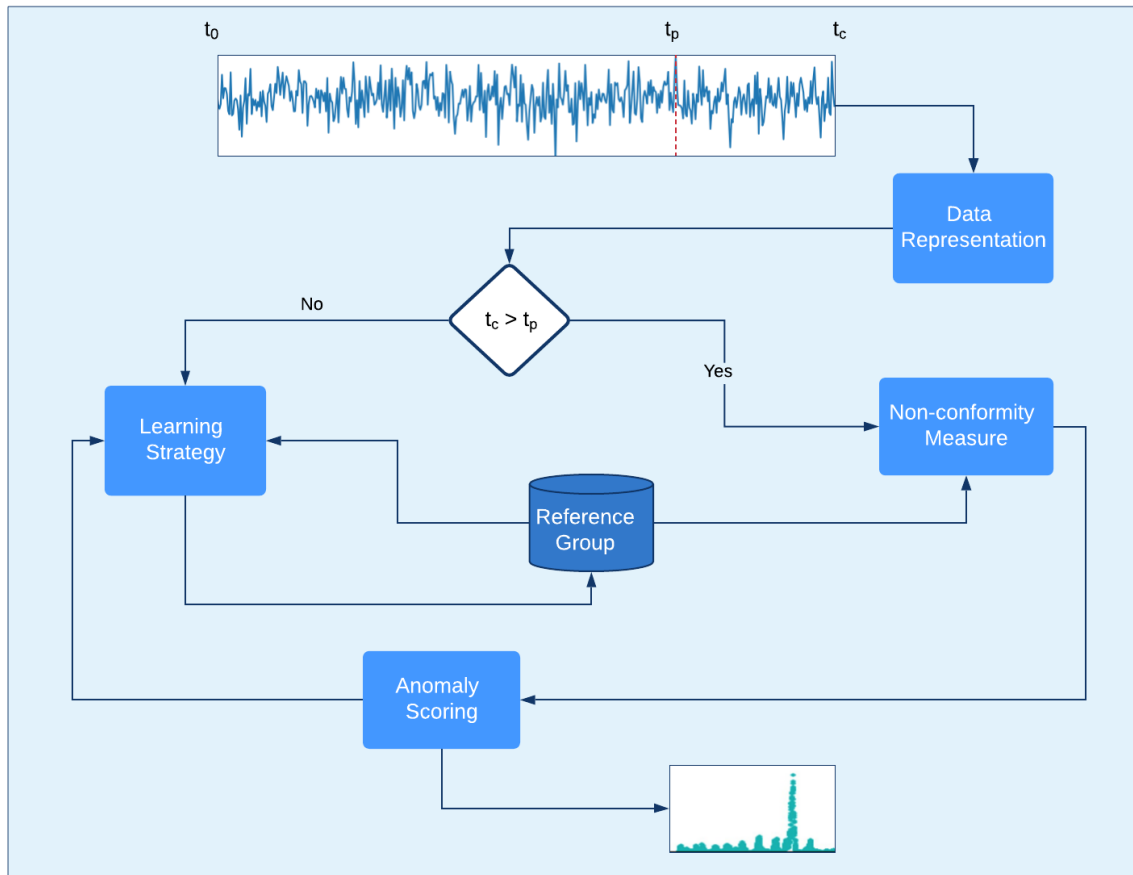
Fig. 1. An overview of the SAFARI framework that shows the flow of data and the modeling steps leading to the generation of final anomaly scores. The framework comprises four components along with a decision process where t_c represents the current time and t_p is the time at which the probationary period is ended.

2.1 Overview of SAFARI

SAFARI consists of four main components: data representation (DR), learning strategy (LS), nonconformity measure (NCM) and anomaly scoring (AS), as illustrated in Fig. 1. The first component, DR, is concerned with automatically transforming raw input data into informative representations or features, so that it can be effectively exploited in anomaly detection tasks. The second component, LS, deals with the selection of the reference group from transformed data, aiming to extract a representative sample of normal behavior from the stream over time. The third component, NCM, measures the nonconformity score of a single observation, with the goal of quantifying the "strangeness" of this observation with respect to the reference group. The last component, AS, aggregates these individual outcomes into a final anomaly score for each observation, taking the global context into account. These four components are explained in more detail in the next section.

Algorithm 1: General Procedure for Streaming Anomaly Detection

---

Input : Stream $S = s_1, s_2, ...$ at time $t = 1, 2, ...$;
         Probationary period p;
         Data representations $\mathcal{D}_1, \mathcal{D}_2, ...$;
         Learning strategies $\mathcal{L}_1, \mathcal{L}_2, ...$;
         nonconformity measures $\mathcal{A}_1, \mathcal{A}_2, ...$;
         Anomaly scorings $\mathcal{F}_1, \mathcal{F}_2, ...$

Output : Final anomaly scores for every s_i in S, where $i > p$

$R \leftarrow \{\}$;   ▷ Reference group
$\hat{A} \leftarrow \{\}$;   ▷ nonconformity scores
$\hat{F} \leftarrow \{\}$;   ▷ Final anomaly scores

Pick a data representation $\mathcal{D}$, where $\mathcal{D} \in \mathcal{D}_1, \mathcal{D}_2, ...$;
Pick a learning strategy $\mathcal{L}$, where $\mathcal{L} \in \mathcal{L}_1, \mathcal{L}_2, ...$;
Pick a nonconformity measure $\mathcal{A}$, where $\mathcal{A} \in \mathcal{A}_1, \mathcal{A}_2, ...$;
Pick an anomaly scoring $\mathcal{F}$, where $\mathcal{F} \in \mathcal{F}_1, \mathcal{F}_2, ...$;

while s_t *for time t is received* **do**
    $x_t \leftarrow \mathcal{D}(s_t)$;
    if $t < p$ **then**
        $R \leftarrow \mathcal{L}(x_t, R)$
    else
        if $\hat{A} = \emptyset$ **then**
            for $x_i \in R$ **do**
                $\hat{A} \leftarrow \hat{A} \cup A(x_i, R \setminus x_i)$ ▷ Compute the first set of nonconformity scores with the reference group R by leave-one-out fashion
            end for
        else
            $\hat{A} \leftarrow \hat{A} \cup \mathcal{A}(x_t, R)$;   ▷ Compute the nonconformity score of x_t
            $\hat{F} \leftarrow \hat{F} \cup \mathcal{F}(\hat{A})$;   ▷ Compute final anomaly score of x_t
            $R_{new} \leftarrow \mathcal{L}_i(x_t, R)$;   ▷ Update reference group with x_t
            $\hat{A} \leftarrow \{\}$;
            for $x_i \in R$ **do**
                $\hat{A} \leftarrow \hat{A} \cup A(x_i, R_{new} \setminus x_i)$;   ▷ Update the nonconformity scores with R_{new}
            end for
            $R \leftarrow R_{new}$
        end if
    end if
end while
return Φ ;

An overview of the generalized procedure of SAFARI is shown in Alg. 1. It can be seen that the overall procedure is implemented with single-pass constraint—that is, the observations in the data stream are processed one at a time without being stored. Furthermore, we define a fixed probationary period (the time that is required to initialize the framework), with the predictions starting afterwards.

The SAFARI framework is designed based on the "separation of concerns" concept, where components are self-contained, cohesive building blocks that serve different purposes in anomaly detection of data streams. This allows one to easily integrate new methods into any of the components or modify the existing components without the need for altering the rest of the framework. Implemented building blocks can be combined in various ways to obtain different and novel detectors, as instantiations of SAFARI.

This setup also provides a basis to conduct more elaborate evaluation experiments. SAFARI allows us to demonstrate, in Section 5 of this paper, the contribution of each building block separately and to conduct thorough algorithm comparisons by isolating effects of shared and unique characteristics of different streaming anomaly detection algorithms. Even though some of the existing approaches may not include all four components that we presented here or the reference group approach, most of the streaming anomaly detection algorithms can be unified using SAFARI.

2.2 The fundamental tasks

Existing algorithms solve the anomaly detection problem using various approaches and based on different assumptions. However, since the overall goal is to find the instances that do not conform to normal behavior in streaming environments, there are fundamental sub-tasks that are shared among many detectors. We have identified four core concepts that play critical roles in the unsupervised streaming anomaly detection problem and can help in distinguishing the underlying principles of different approaches. Below we present the high-level overviews of these general tasks, and in the next section, we review the state of the art from the perspective of these tasks.

Data representation: Data representation, in general, is concerned with automatically transforming raw input data into representations or features that can be effectively exploited in machine learning tasks. Useful representations can capture important clues about the past and the current state of the stream as well as the key characteristics of the object (e.g., the monitored system) that are relevant for anomaly detection. The goal of this task is to provide a more vibrant representation of a stream of data consisting of one or more time series that helps to better distinguish anomalies from normal data.

Definition 2.1. (Data representation) Let $S = \{s_1, s_2, ..., s_t\}$ be an input stream where $s_i \in \mathbb{R}^d$. A **data representation** is a function D that takes an observation (or a set of observations) in the input stream and transforms them into a feature vector $x_t = D(s_i, ..., s_t)$ such that $s_i \in S$, $x_t \in \mathbb{R}^d$.

Such features or representations can be obtained in many different ways, such as by extracting means, averages, correlations, or distributions, or by using linear/nonlinear functional relationships, domain knowledge, and so on. There are endless possible model families and hierarchies of models of increasing complexity.

Learning strategy:

This task is concerned with how to effectively learn the reference group, making sure that it represents the current normal behavior of the stream. A data stream has a continuous flow and the number of incoming observations is unbounded. Unlike static anomaly detection, algorithms that need to learn normal behavior in dynamic environments should have the ability to process new data and limit the number of data points stored in the reference group. At the same time, the reference group should be continuously updated since normal behavior changes over time in a dynamic environment. Therefore, the selection of this set is one of

6

the crucial tasks that differentiate anomaly detection in static and dynamic settings. We refer to the task of maintaining and updating the reference group, where all the observations are not available at once and arrive sequentially, as the "learning strategy."

Definition 2.2. (Learning strategy) Given that $R_0 = \emptyset$, a **learning strategy** L is a function,

$$R_t = L(x_t, R_{t-1}),$$

where x_t is the current observed feature, R_{t-1} is the reference group before observing x_t, and R_t is the new reference group at time t.

Various general windowing techniques (e.g., sliding window, damp window) or sampling algorithms (e.g., uniform sampling) can be given as different examples of learning strategies.

Nonconformity measure:

Essentially, identifying how well the samples conform to the normal behavior is a core step in all unsupervised anomaly detection algorithms. In this task, the goal is to quantify how "strange" a single observation is, using a measure which we refer to as "nonconformity measure".

Definition 2.3. (Nonconformity measure) Given the reference group R_t and the sample x_t, a **nonconformity measure** A is a function,

$$a_t = A(x_t, R_t),$$

where a_t is the nonconformity score indicating how "strange" x_t is with respect to R_t.

Various approaches with different "normality" assumptions can be used to measure nonconformity: for example, measuring the average distances to nearest neighbors, the local density, the variance in the angles, the goodness of fit to a generative model, or the difference between actual and predicted (i.e., expected) values.

Anomaly scoring: The aim of most anomaly detectors is to output a score for each sample that indicates how likely it is to be an anomaly. In some cases, the levels of "strangeness" that are measured in the previous task do not directly correspond to the desired levels of "anomalousness." The candidates with the high nonconformity scores may not be statistically significant or semantically relevant for the particular use-case. This stage is concerned with the post-processing of the nonconformity scores, which transforms "strangeness" scores of individual observations into "anomaly" scores based on the global context.

Definition 2.4. (Anomaly scoring) Let $A = \{a_1, a_2, ..., a_t\}$ be a set of nonconformity scores, an **anomaly scoring** F is a function

$$f_t = F(a_i, ..., a_t),$$

such that $a_i \in A$, that maps nonconformity scores to final anomaly scores.

For example, in some domains one is interested in collective anomalous behavior of the observations rather than the individual level of "strangeness." The final scoring task can aggregate the nonconformity scores of the samples to produce anomaly scores at the collective level. In another approach, assuming faults or anomalies do not occur suddenly and expecting a certain level of temporal "continuity" in the detection, the anomaly scoring task can track nonconformity scores and give higher scores to deviations that persist over time.

3 Related Work

To the best of our knowledge, our work is the first attempt to formalize the common tasks and properties of streaming anomaly detection. One similar work was proposed by Schubert et al. (2014), who discussed similarities and differences in local outlier detection methods, focusing on the notion of "locality" and proposed an algorithmic structure that unifies the existing methods. Therefore we review the state of the art related to each of the four core concepts that we introduced in the previous section.

3.1 Data Representation

A suitable choice of representation greatly affects the ease and efficiency of all data analysis tasks. Therefore, there is a rich literature around this subject. In this study, we review the techniques that are introduced primarily for temporal data but that are also suitable to be used in streaming fashion. The most popular techniques widely used for time series representation include the discrete Fourier transform (DFT) (Faloutsos et al., 1994), piecewise models (Geurts, 2001; Yi and Faloutsos, 2000), and singular value decomposition (Keogh et al., 2001b). Many of these techniques have already been extended to be applied in a streaming fashion. For example, Zhu and Shasha (2002) proposed a streaming version of DFT for real-time monitoring of thousands of streams. Lazaridis and Mehrotra (2003) implemented an online version of piecewise constant approximation with little loss of accuracy. Other methods also aim to summarize data streams using simple statistics (e.g., mean, standard deviation or sum) (Cohen and Strauss, 2003), wavelets (Cormode et al., 2006; Gilbert et al., 2003) or histograms (Fan et al., 2015). Bulut and Singh (2003) applied wavelets to represent data streams in a way that is biased towards more recent values.

Linear models are also widely used in many works to represent single or multiple data streams (Kargupta et al., 2004; D'Silva, 2008). Kargupta et al. (2006) and Kargupta et al. (2010) used linear correlations to monitor correlations of on-board signals for vehicles. Other studies, by contrast, incorporated methods to capture non-linear relationships such as neural gas models (Vachkov, 2006) and reservoir computing models (Chen et al., 2013; Quevedo et al., 2014). Echo state networks were applied by Fan et al. (2015) to represent air compressor sensor data.

There are also more recent techniques that use autoencoders (Li et al., 2015) or neural networks (Chen et al., 2013). For example, Rögnvaldsson et al. (2018) has formalized the "interestingness" concept to find useful data representation and include different autoencoders and histograms as representations.

3.2 Learning strategy

The large volume of data streams poses unique constraints on the computation process in terms of memory and computations. These challenges have led researchers to propose strategies to efficiently approximate streams over time. While some of these methods are general, in that they can be applied to different data mining tasks (e.g., classification or clustering), the rest are developed to be used in specific problems. The common windowing techniques that are broadly applied for different tasks can be categorized into four groups: landmark, sliding, damped and adaptive. For example, landmark windows were employed for clustering in Birch (Zhang et al., 1996) and CluStream (Aggarwal et al., 2003) and for the frequent pattern mining (Manku and Motwani, 2012; Jin and Agrawal, 2005; Li et al., 2015; Leung and Jiang, 2011). Subramaniam et al. (2006) and Angiulli and Fassetti (2007) used sliding windows for outlier detection, while Yang et al. (2009) and Rögnvaldsson et al. (2018) applied them for condition monitoring. Other methods (Cao et al., 2006; Leung and Jiang, 2011) have incorporated damped windows, in which older data points in the window get lower weights than newer points. Adaptive windows are mostly concerned with change detection where the goal is to adapt to the changes more quickly by changing the size of the window. Bifet

and Gavalda (2007) proposed ADWIN for maintaining a window of variable size, growing automatically when the data is stationary and shrinking when change is taking place.

Different sampling algorithms have also been proposed or adopted for mining data streams. Many streaming outlier detection methods exploit uniform sampling (Zimek et al., 2013; Liu et al., 2012; Sugiyama and Borgwardt, 2013), which is a simple but effective technique. Some others proposed custom sampling techniques that are tailored for specific tasks or algorithms. Kollios et al. (2003) proposed a density-biased sampling approach for clustering and outlier detection, in which the probability that a point is included in the sample is determined by the point's local density. Similarly, Zhang et al. (2018) proposed another density-biased sampling for local outlier detection. While sampling sparser regions at higher sampling rates, it also sampled dense regions, at lower sampling rates, to strengthen "outlierness" contrast.

The effectiveness of most of the general methods—for example, uniform sampling—on anomaly detection have not been well recognized and studied in the literature. On the other hand, it is difficult to identify the benefits of the custom sampling algorithms for anomaly detection, since they cannot be applied outside of the specific settings that they were designed for. For example, the density-biased sampling algorithm of Kollios et al. (2003) is only applicable with the methods that use density. We have not come across any learning strategy that is designed to consider the properties of the anomaly detection problem and is also general so that it can be combined with any anomaly detector.

3.3   Nonconformity measure

There are many ways to measure nonconformity in anomaly detection problem and there is a huge body of literature on this subject. We suggest several surveys on this topic: Aggarwal (2015), Chandola et al. (2009),Gupta et al. (2013) and Zimek et al. (2012). Here, we review some common approaches, i.e., probabilistic and statistical proximity-based and prediction-based models that are applied in the streaming setting.

In statistical-based approaches, the aim is to learn a statistical model for a normal behavior of a dataset and determine the nonconformity of new observations by measuring how well they fit into that model. Yamanishi and Takeuchi (2002) and Yamanishi et al. (2004) proposed SmartSifter, based on an online discounting learning algorithm that incrementally learns the probabilistic mixture model and calculates deviation of the incoming data from this model. Several other statistical methods (Kuncheva, 2011; Song et al., 2007) use log-likelihood criteria in order to quantify nonconformity.

The idea in proximity-based methods is to measure nonconformity of data points based on their similarity to or distance from the normal data. Angiulli and Fassetti (2007) and Kontaki et al. (2011) proposed efficient computation of nearest neighbors and use sliding windows to detect global distance-based outliers in data streams. Distance-based "local" outlier techniques that extend the local outlier factor (LOF) algorithm to the case of streaming data have been discussed by Na et al. (2018), Pokrajac et al. (2007) and Salehi et al. (2016). Many clustering-based methods use distance to the cluster centers as the measure of nonconformity, while proposing varying algorithms to effectively cluster data streams. Cao et al. (2006) used the concept of micro-clusters to distinguish between normal data and outliers based on the distance to the centers. AnyOut (Assent et al., 2012), an anytime algorithm, applied a specific tree structure that is suitable for anytime clustering and computes the nonconformity score using the distance to the nearest cluster centroid. Chenaghlou et al. (2017) has proposed a hyper-ellipsoidal clustering approach to model the normal behavior of the system, where nonconformity is determined based on the distance to the cluster boundaries.

Prediction-based methods mostly employ regression-based forecasting models, and nonconformity scores are calculated on the basis of deviations between actual observations and their expected (or forecasted) values. Some works used traditional regression methods such as autoregressive modeling, autoregressive moving average, and autoregressive integrated moving average. Since the success of the prediction process

greatly affects the final accuracy, most of the prediction-based methods focused on the prediction model rather than the anomaly detection itself. The work in Yahoo's EGADS framework (Laptev et al., 2015) has provided a set of regression methods that can be selected or integrated by the user. Another work by Ahmad et al. (2017) used hierarchical temporal memory networks as their prediction model and detected anomalies by tracking prediction error over-time. Hundman et al. (2018) and Malhotra et al. (2015) have shown that recurrent neural networks achieve high prediction performance and perform effectively across a variety of domains. Several methods have also been focused on speeding up the regression modeling in the context of a large number of data streams and real-time data (Huang et al., 2007; Jiang et al., 2011; Yi et al., 2000).

3.4 Anomaly scoring

Final scoring has not been formalized as a stand-alone process in most of the methods, and therefore it is much more difficult to abstract from the existing approaches. In some works, it appears as a global normalization step—for example Schubert et al. (2014), which transfers nonconformity scores to anomaly estimates to satisfy a clear gap between the scores of anomalies and normal samples. Kriegel et al. (2011) and Gao and Tan (2006) proposed generic scoring functions that can convert any set of the nonconformity scores into probability estimates.

Laptev et al. (2015) discussed that the prediction error (i.e., nonconformity scores) would not be suitable for time-series anomaly detection and computed relative errors in final scoring. An anomaly likelihood function was proposed by Ahmad et al. (2017) to define how anomalous the current sample is based on the prediction history of the model. A sliding window was maintained on nonconformity scores and the anomaly likelihood of each window was defined as the final anomaly score. Maurus and Plant (2017) and Rögnvaldsson et al. (2018) applied statistical tests on the nonconformity scores to produce final anomaly scores that capture only deviations that persist over time. A probabilistic approach was proposed by Olsson and Holst (2015), aggregating point outliers into group (i.e., collective) anomalies. Several other methods used martingales to convert nonconformity measures to change-point estimates (Ho, 2005; Ho and Wechsler, 2010; Volkhonskiy et al., 2017).

4   Methods

In this section, we present the details of the methods that we have implemented using SAFARI framework [1]. Experimental evaluation of these methods is presented in the next section. In total, we integrate 12 separate methods—namely, two data representations, five learning strategies, four nonconformity measures, and one anomaly scoring method. These are selected so that, using various combinations, we can build 20 different SAFARI anomaly detectors. Even though only one of these methods is new, many of the combinations themselves produce novel anomaly detectors that have not been tried before.

It is important to note that in this study we mainly focus on two out of the four tasks defined in the previous section: learning strategy and nonconformity measure. In most of the existing solutions, the procedures concerning these two essential tasks are entirely embedded in the overall solution, which makes it difficult to study their individual contributions. Furthermore, no studies address the impacts of nonconformity measures and learning strategies separately, nor the impact of combining them into different anomaly detectors. Data representation is heavily investigated in the literature, and we do not believe we can offer important contributions in that area. Furthermore, we have decided to leave the analysis of anomaly scoring to another

---

[1]The code is available at `https://github.com/caisr-hh`

work, considering that it would greatly increase the number of combinations. Ultimately, we present diverse methods for the tasks of learning strategy and nonconformity measure to investigate their behaviors in depth, while we only implement two data representation methods and one method for anomaly scoring.

4.1 *Data representations*

The first data representation implemented in the SAFARI framework is a simple approach using mean and standard deviation of the last N observations to represent a feature (Rodríguez and Alonso, 2004):

$$\mu_t = \frac{\sum_{i=0}^{N-1} s_{t-i}}{N}, \tag{1}$$

$$\sigma_t = \sqrt{\frac{\sum_{i=0}^{N-1} s_{t-i} - \mu_t}{N}}, \tag{2}$$

where s_t is the observation at time t and $x_t = \{\mu_t, \sigma_t\}$ is the feature at time t.

The second representation method is based on the SAX (symbolic aggregate approximation) (Lin et al., 2003) approach: that is, the discretization of the original data stream into symbolic strings. SAX performs this discretization by dividing a z-normalized subsequence into w equal-sized segments. For each segment, it computes a mean value (i.e., piecewise aggregate approximation, PAA, Keogh et al. 2001a) and maps it to symbols according to a user-defined set of breakpoints dividing the distribution space into α equi-probable regions, where α is the alphabet size specified by the user.

In this work, we apply SAX on overlapping subsequences in a single-pass streaming fashion. Given a data stream $S = \{s_1, s_2, ..., s_t\}$, we generate a SAX word x_t, which is the feature at time t, based on a subsequence $\hat{s}$ that comprises the last n observations: $\hat{s} = \{s_{t-n-1}, s_{t-n}.., s_t\}$.

4.2 *Learning strategies*

Here, we present five different learning strategies that are integrated into the framework.

The first strategy, fixed reference group (FR), maintains a static set of instances as a reference group—that is, it does not change over time. Clearly, this strategy is not suitable for many streaming scenarios, especially ones with concept drift. However, we include it as a benchmark and to be able to compare static and dynamic methods, showcasing how they perform under different combinations of the framework components.

The other three strategies (i.e., sliding window (SW), landmark window (LW), and uniform reservoir sampling (URES)) are popular techniques that have been widely used in many streaming applications, including classification and clustering tasks. However, a thorough analysis of these approaches in the anomaly detection problem is still missing from the literature. By integrating them, we make it possible to study their individual performances separately from the rest of the framework and identify their benefits or drawbacks in various datasets.

Finally, we propose a new learning strategy, anomaly-aware sampling (ARES), which provides a generic method that requires only anomaly scores as input and is specifically designed considering the research gap in the anomaly detection problem.

4.2.1 Fixed reference group

The fixed reference group method maintains a window that collects the observations arriving in probationary period p. This learning strategy essentially provides a static reference group that does not change over time after the probationary period is over (Fig. 2a).

$$R_t = \begin{cases} R_{t-1}, & \text{if } t > p, \\ R_{t-1} + \{x_t\}, & \text{otherwise.} \end{cases} \tag{3}$$

4.2.2 Landmark window

In this windowing technique, a fixed timestamp in the data stream is defined as a landmark, and processing is done over data points between the landmark and the present time (Fig. 2c). Landmarks are usually defined by the user where they can be chosen as the starting timestamp of the stream or a specific timestamp such as the beginning of a year. In this study, we assume the landmark is the time when we observe the first sample ($t = 0$).

The reference group at time t with landmark windowing is as follow:

$$R_t = \begin{cases} \emptyset, & \text{if } t \leq l, \\ R_{t-1} + \{x_t\}, & \text{otherwise.} \end{cases} \tag{4}$$

where l is the landmark time.

Note that learning continues by adding the new samples to the reference group unless either the query is explicitly revoked or the stream is exhausted and no additional observations are entered into the system. Therefore, the size of the reference group is not fixed over time.

4.2.3 Sliding window

In the sliding window approach, the oldest sample in the window is discarded whenever a new sample is observed (Fig. 2b). Given a window size w and the new observed sample x_t, the reference group at time t is updated as below:

$$R_t = \begin{cases} R_{t-1} - \{x_{t-w}\} + \{x_t\}, & \text{if } t > w, \\ R_{t-1} + \{x_t\}, & \text{otherwise.} \end{cases} \tag{5}$$

(a) Fixed Reference Group
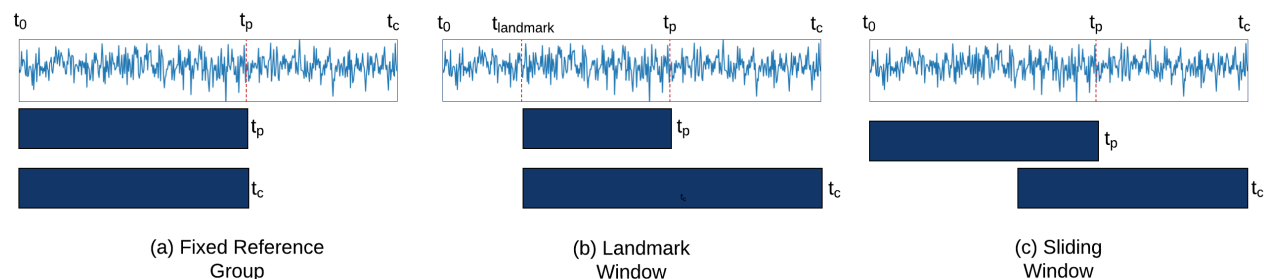
(b) Landmark Window

(c) Sliding Window

Fig. 2. Illustrations of windowing techniques, where t_c represents the current time and t_p is the time where the probationary period is ended.

4.2.4 Uniform Reservoir

The reservoir sampling algorithm (Vitter, 1985) is a classic method of sampling without replacement from a stream in a single pass when the stream is of indeterminate or unbounded length. Assume that the size of the desired sample is w. The algorithm proceeds by retaining the first w items of the stream and then sampling each subsequent element with probability $f(w,t) = \frac{w}{t}$, where t is the current time and also gives the length of the stream so far.

Given reservoir size w, the reference group with landmark at time R_t is computed as follows:

$$R_t = \begin{cases} R_{t-1} + \{x_t\}, & \text{if } t \leq w, \\ R_{t-1} - \{x^*\} + \{x_t\}, & \text{if } t > w \wedge U < \dfrac{w}{t}, \\ R_{t-1} & \text{otherwise.} \end{cases} \qquad (6)$$

where x^* is a uniformly chosen element from R_{t-1}.

4.2.5 Anomaly-aware reservoir

Most of the learning strategies in the literature focus on providing accurate approximation of the stream while processing large volumes of data efficiently. However, these algorithms, including the other learning strategies integrated into SAFARI, are not designed considering the constraints of the anomaly detection problem and do not guarantee the maintenance of a representative sample of the *normal* behavior over time. For example, the underlying assumption in uniform sampling, which is that all points are of equal importance, has a serious drawback when it is directly applied to a stream containing anomalies. Clearly, sampling the anomalies and normal samples with equal probability can cause the contamination of the reference group and leads to the phenomenon called "masking," which results in the incoming anomalies passing undetected.

To deal with this problem, we propose the anomaly-aware reservoir sampling by generalizing the weighted reservoir sampling schema for anomaly detection problem. In our method, we extend the online algorithm proposed by Efraimidis and Spirakis (2006) for the case in which data has a different anomaly score distribution. The goal here is to ensure the samples in the reservoir are more biased toward the samples with lower scores, i.e., samples more likely to be normal.

In a nutshell, the idea of the weighted sampling algorithm is to draw a sample of size k without replacement where the probability of selecting each sample at time t is equal to the sample's weight divided by the total weights of samples that are not selected before time t. Similarly, our learning strategy generates a weighted random sample in one pass over incoming streams and maintains a reservoir with a size N that constitutes the reference group R.

The process starts with assigning a "priority" to each sample using a weight function $w()$. Let x_t be the sample at time t; we define the function $w(x_t)$ which assigns the weight of x_t as follows:

$$w(x_t) = exp(-\lambda S(x_t)). \qquad (7)$$

where λ is the decay factor and $S(x_t)$ is the anomaly score of x_t. The choices for the decay factor are suggested as $0.96 \leq \lambda \leq 0.98$ by Haykin (1996), and we use $\lambda = 0.96$ in this study. Our method is generic such that $S(x_t)$ can be produced by any method chosen for anomaly scoring, as defined in Section 2.2. Our choice of a method for that task in the experiments is presented in Section 4.4.

The weight function is designed to give lower importance to instances with high anomaly scores, ensuring that anomalous points have lower probability of being represented in the reference group. Therefore, the strategy aims to avoid learning new abnormal instances while forgetting the ones that are already present. This aspect is especially important when the initial reference group is highly contaminated by the anomalies.

The learning strategy generates the "priority" $p_t = u^{\frac{1}{w(x_t)}}$ for the sample x_t, where $w(x_t)$ is the weight and u is drawn randomly from [0,1]. In the original implementation, the samples with the highest N priorities are always kept in the reservoir. In each iteration, the sample with the smallest priority is taken as a threshold T and then is replaced by sample x_t if p_t is larger than T.

However, in the presence of a nonstationary distribution, the learning strategy must incorporate some form of forgetting past and outdated information. Therefore, instead of removing the item with the lowest priority, we determine the set of candidate samples that have priorities lower than x_t and remove the oldest one among the candidates.

The goal here is to continuously update the reservoir in such a way that the older items are consistently replaced while still maintaining normal samples in the reference group. The details of the overall procedure are shown in Alg. 2

Algorithm 2: Anomaly-aware reservoir sampling

Input : Reference group R;

Reservoir size w;

New sample (x_t, t)

Output : Reference group R

$priorities \leftarrow \emptyset$;

$s_t \leftarrow$ Collect anomaly score of x_t ;

$p_t \leftarrow u^{\frac{1}{e^{-\lambda s_t}}}$;

if $t < w$ **then**

$\quad priorities \leftarrow priorities \cup (p_t, t)$;

$\quad R \leftarrow R \cup (x_t, t)$;

else

$\quad candidates \leftarrow$ Collect samples where priorities are smaller than p_t;

$\quad$ **if** $|candidates| > 0$ **then**

$\quad \quad i \leftarrow argmin(candidates)$;

$\quad \quad priorities \leftarrow priorities/(p_i, i) \cup (p_t, t)$;

$\quad \quad R \leftarrow R/(x_i, i) \cup (x_t, t)$;

$\quad$ **end if**

end if

return R ;

4.3 Nonconformity measures

The four nonconformity measures that are incorporated into the framework are as follows: (i) nearest neighbors-based (NN), (ii) density-based (DEN), (iii) clustering-based (CC) and (iv) frequency-based (FREQ). The first three approaches are based on the popular proximity-based models in which the nonconformity scores are determined by, respectively, the average k-nearest neighbor distance, local density value and distance to closest cluster centroid. Even though these are very common approaches that have been employed by many different anomaly detection algorithms, their streaming versions are not well-studied.

The fourth method is the frequency-based approach, which measures nonconformity by the number of occurrences of the pattern, with low frequencies leading to higher nonconformity scores. We have integrated

this algorithm to increase the diversity in the framework and also to provide a choice that has a significantly lower computational cost.

4.3.1 Nearest neighbors-based NCM

The average distances to the k-nearest neighbors (KNN) is used as a measure of nonconformity.

$$a_t = \frac{\sum_{i=1}^{k} d(x_t, NN_i(x_t))}{k}, \tag{8}$$

where $NN_i(x_t) \in R_t$ is ith nearest neighbour of x_t.

4.3.2 Density-based NCM

This measure quantifies the nonconformity of the samples based on their local densities, under the assumption that anomalies do not lie in dense regions. In this work, we use the LOF to estimate nonconformity scores, since it adjusts for the variations in the local densities of different regions.

Given two points x_i and $x_j \in R$, the *k-reachability distance* of x_i with respect to x_j is

$$R_k(x_i, x_j) = \max\{d(x_i, NN_k(x_i)), d(x_i, x_j)\}, \tag{9}$$

where d is the distance function and $NN_k(x_i)$ is the kth neareast neighbor of x_i.

Local reachability density, LRD_k is given by

$$LRD_k(x_t) = \left(\frac{1}{k} \sum_{i=1}^{k} R_k(x_t, NN_i(x_t)) \right)^{-1}, \tag{10}$$

where $NN_k(x_t) \in R$ is a set of the k-nearest neighbors of x_t.

Finally, the nonconformity score of x_t is equal to its local outlier factor, LOF_k given by

$$a_t = LOF_k(x_t) = \frac{1}{k} \sum_{i=1}^{k} \frac{LRD_k(x_t)}{LRD_k(NN_i(x_t))}. \tag{11}$$

In traditional LOF, the LOF scores of all data points should be updated whenever a new data point is inserted or removed from the reference group R_t, which is computationally expensive. We use iLOF (Pokrajac et al., 2007), which selectively updates the scores of only the instances affected by the change in the reference group.

4.3.3 Clustering-based NCM

In clustering-based NCM, the distance from the nearest cluster centroid is used as a measure of nonconformity. Let R_t be the reference group and x_t be the sample at time t, the nonconformity score of x_t is computed as follows:

$$a_t = min(d(x_t, C(R_t)), \tag{12}$$

where d is the distance function and $C(R_t)$ denotes all the cluster centers computed on R_t.

Here, the clustering algorithm used for partitioning the reference group into disjoint sets can be chosen freely. We use the incremental version of k-means by Ordonez (2003) to compute clusters and centroids since it can be easily adopted in the streaming scenario. In this method, every new example is added to the

cluster with the nearest centroid, and in every r steps a recomputation phase occurs, which updates both the assignment of points to clusters and the centroids. Ordonez (2003) chooses r to be the square root of the number of points seen so far, aiming to balance accuracy and computation time. However, in our case, we update the cluster centroids at each time step based on the learning strategy in which the old samples can be removed from the clusters as the new ones are added. Therefore, we follow Bifet and Gavalda (2006) for the recomputation phase, who suggests recomputing when an average point distance to centroids has changed more than an ε factor, where the ε factor is user-specified.

4.3.4 Frequency-based NCM

This nonconformity measure is motivated by the assumption that anomalies are rare items in the behavior, and samples that form infrequent patterns are more likely to be anomalous. Therefore, measuring nonconformity is directly related to measuring the surprisingness level of the sample, which is defined as the frequency of its occurrence in normal behavior.

After applying the chosen data representation method, the frequency is measured by monitoring the number of occurrences of patterns in the reference group, where a "pattern" is a subset of the feature space at any time. Together with this nonconformity measure, we specifically use SAX representation, which has been shown to be a very powerful method to capture meaningful patterns in a data stream (Keogh et al., 2001b). Nonconformity scores of the samples are determined by their "term" frequencies—that is, the number of times they occurred in the reference group.

To track term frequencies dynamically, we create a hash table using SAX words encountered in the reference group as the keys and their number of occurrences as hashed values. Given a reference group R_t, a hash table H and the current sample x_t that corresponds to a SAX word, the nonconformity score a_t of x_t is computed as

$$a_t = \frac{|R_t|}{f(x_t) + 1}.$$ \hfill (13)

where $|R_t|$ is the size of R_t, and $f(x_t)$ retrieves the frequency of x_t from the hash table H. The hash table is convenient data structure for this task since insert, update and lookup operations take $O(1)$ and the space is also bounded with $O(N)$ where N is the size of the reference group R_t.

4.4 Anomaly scoring

In this work, we incorporate only one method for final scoring. It is based on the statistics that has been used in conformal prediction (Vovk et al., 2005). The procedure of anomaly scoring can be seen in Alg. 3. To compute anomaly scores, we first estimate p-values for every new observations using nonconformity scores where p-values correspond to confidence levels for each prediction:

$$p_t = \frac{|i = 1, ..., w : a_i \geq a_t|}{w}.$$ \hfill (14)

In this case, high p-values are consistent with the definition of an outlier by Hawkins (1980), where an observation with a high p-value corresponds to the one that deviates so much from other observations as to arouse suspicion that it was generated by a different mechanism. This definition considers an anomaly as an extreme single point that occurs "individually" and "separately."

In many streaming applications, the temporal continuity plays a critical role to the notion of abnormality, since anomalies mostly occur as abnormal patterns rather than independent outlying observations, or they lead to abrupt or gradual changes exhibiting a lack of continuity with their immediate or long-term history.

Furthermore, to be able to detect anomalies in the early stages, one cannot wait for the metric to be clearly beyond the bounds (e.g., p-values) and the ability to detect subtle changes is needed.

We track p-values over time instead of reporting them directly as anomaly scores and apply statistical hypothesis testing under the null hypothesis that the p-values should be uniformly distributed (based on Theorem 1):

Theorem 4.1. (Vovk et al., 2005)

If the data samples $\{x_1, x_2, \cdots\}$ satisfy the i.i.d. assumption, the p-values $\{p_1, p_2, \cdots\}$ are independent and uniformly distributed in $[0, 1]$.

Specifically, this hypothesis is tested using the Kolmogorov-Smirnov (K-S) one-sample test (Kolmogorov, 1933), where we compare the empirical cumulative distribution function of p-values with the cumulative distribution function of the uniform.

The empirical cumulative distribution function $F_t(p)$ of the sequence of n p-values $\{p_{t-n+1}, p_{t-n+2}, \cdots, p_t\}$ is given by

$$F_t(p) = \frac{1}{n} \sum_{i=t-n+1}^{t} I(p_i \leq p),$$ (15)

where I is an indicator function such that I equals 1 if $p_i \leq p$ and 0 otherwise. Given $F(p)$ is the cumulative uniform distribution function, the one-sample Kolmogorov–Smirnov statistic for time t is

$$D_t(p) = sup_p |F_t(p) - F(p)|.$$ (16)

where sup_p denotes the supremacy of the set of distances between the curves.

The probability of observing such a D_t under the null hypothesis is evaluated. We use the significance levels obtained from the K-S tests (it should be noted that they are different than the p-values calculated in Eq. 14) as an indicator for anomaly scores. The significance levels can not be directly interpreted as anomaly scores since p-values will have very low values. Therefore, we apply a score unification step to convert these values into probability estimates by regularization, normalization and scaling steps. Following Kriegel et al. (2011), we use logarithmic inversion for regularization, a simple linear transformation for normalization and Gaussian scaling to produce final scores. The advantages of the unification of the scores is that it allows the comparison of different combinations of the framework and also makes it possible to create an ensemble of them in the future.

5 Evaluation

In this section, we conduct in-depth evaluations for the anomaly detection algorithms within our framework. We introduce the datasets and parameter configurations that we use in this study, and then report our overall evaluation methodology and results. Finally, we summarize our findings and provide intuitive recommendations on selecting appropriate settings for different scenarios.

5.1 *Datasets*

In the following, we describe the two real-world benchmark datasets—Numenta Anomaly Benchmark (NAB) and Yahoo S5 Webscope Benchmark—that were used in this work.

NAB provides a set of real-world and artificial datasets that are designed for research in streaming anomaly detection. It is composed of 58 datasets containing labeled anomalous periods of behavior. The

Algorithm 3: Anomaly Scoring

Input : Nonconformity scores of the reference group A_R;

   Nonconformity score of the current sample a_t;

   test period u

Require: Current p-values P;

if $P = \emptyset$ **then** ▷ Generate p-values of the first reference group

   for $a_i \in A_R$ **do**

   $\quad p_i \leftarrow \frac{|j=1,...,|A_R \setminus a_i|:a_j \geq a_i|}{|A_R \setminus a_i|}$;

   $\quad P \leftarrow P \cup p_i$;

   end for

end if

$p_t \leftarrow \frac{|j=1,...,|A_R|:a_j \geq a_t|}{A_R}$;   ▷ Compute p-value of the test sample x_i

$P \leftarrow P \cup p_t$;

$\sigma \leftarrow KSTEST(P,u)$;

$s_t \leftarrow UNIFICATION(\sigma)$;

return s_t ;

Output : Anomaly score s_t at time t;

---

majority of the NAB datasets are real-world from different domains and applications such as AWS server metrics, Twitter volume, advertisement click metrics, real-time traffic data from Minnesota, temperature sensor data, and so on. Each dataset exhibits different characteristics such as temporal noise, short and long-term periodicities and concept drift.

Yahoo Webscope S5 benchmark is released by Yahoo Labs for the detection of unusual traffic on Yahoo servers. It consists of 367 time-series datasets in four classes in which the ground truth anomaly information is available for all time series. In this study, we use A1 class, which consists of real datasets from Yahoo's computational services, while other classes contain synthetically generated data. A1 datasets comprise 67 time series with various seasonality, distinct change patterns, and diverse types of anomalies that are based on real measurements from various Yahoo cloud services, such as Yahoo Membership Login (YML).

5.2 Evaluation metrics

In our experiments, we adopt two metrics (i.e., ROC-AUC and NAB scoring) to evaluate the detection performances of SAFARI detectors.

The first metric, ROC-AUC, is the most popular measure for the evaluating unsupervised anomaly detection methods. It summarizes the ROC curve score with a single value that ranges between 0 and 1. According to Aggarwal (2015) given a scoring of a set of points in order of their propensity to be anomalies, the ROC AUC is equal to the probability that a randomly selected anomaly-nominal pair (a,n) is scored in a correct order where an anomaly appears before a nominal.

$$ROC-AUC = mean_{a \varepsilon A, n \varepsilon N} \begin{cases} 1, & \text{if } Score(a) > Score(n), \\ 1/2, & \text{if } Score(a) = Score(n), \\ 0 & \text{if } Score(a) < Score(n). \end{cases} \tag{17}$$

ROC-AUC is a useful measure to understand whether a method exhibits a high ratio of correctly detected anomalies (i.e., true positive rate, TPR) while providing few normal samples misidentified as anomalies (i.e.,

false positive rate, FPR). However, this metric only takes the ratio of detected anomalies to nominals into account, ignoring the positions of the samples in the time series.

The second metric that we use in this study is NAB scoring which is a measure provided by Numenta to assess the quality of streaming anomaly detection algorithms. The key aspect of NAB scoring is that it is designed to reward early detection, which is a quite useful feature for many streaming applications. To incorporate the knowledge of early or late detection into scoring, NAB Benchmark defines the concept of an "anomaly window," which consists of a sequence of data points centered on one (or more) true anomalies in a dataset. In a nutshell, NAB scoring considers detection within a window as true positives (TP), which gives positive values to the NAB score such that a TP detected at the beginning of the window has a higher value. If there are multiple detections within a particular anomaly window, the scoring considers only the earliest detection as a TP and ignores all (considered superfluous) detections that follow. This means that an anomaly detector that detects only the first point in the window as an anomaly will receive a higher score than a detector that detects as anomalies all the points in the window except the first one.

Furthermore, detections made outside the window are considered false positives (FP), and make negative contributions to the NAB score. The position of the detection is also taken into account for FPs. If an FP occurs close to a window, it gets a less negative value than if it occurs further away from the window. Missing a window completely results in a false negative (FN) and makes a strong negative contribution to the score. More details about the method can be found in (Lavin and Ahmad, 2015).

The maximum NAB score a detector can achieve in a dataset is equal to the number of anomaly windows in that dataset. To be able to compare detection performances on different datasets, we normalize the NAB scores using the number of windows such that the score of the perfect detector is 1, and the null detector is 0. It is important to note that NAB scores are not lower-bounded, since the lowest score of a detector depends on the number of FPs—that is, the number of normal samples in a dataset.

The most important drawback of the NAB scoring is defining anomaly windows efficiently. Selecting larger windows allows the rewarding of earlier detection of anomalies, but it can lead to actual FPs be counted as TPs, thus rewarding inaccurate detection. The authors of the Numenta benchmark (Lavin and Ahmad, 2015) recommend choosing the window size to be 10% of the number of instances in a dataset, divided by the number of true anomalies in the given dataset. We follow this suggestion when we generate anomaly windows for each dataset in Yahoo Benchmark to be used for the evaluation with Numenta scoring.

Contrary to the ROC-AUC score, the NAB scoring requires a threshold value on anomaly scores to cutoff between anomalies and normals. To limit the computational cost, we set a global threshold to 0.9 providing a guaranteed bound of %10 false positive rate for SAFARI detectors for all datasets, instead of optimizing the threshold for each dataset separately.

5.3   *Experimental setup*

In our experiments, all requisite parameters of the integrated methods of data representations (i.e., mean-std and SAX), nonconformity measures (i.e., NN, DEN, CC, and FREQ) and anomaly scoring (i.e. CAD) are tuned to select the best parameters for the given evaluation metric. Another parameter of SAFARI, the probationary period, p, is chosen as the first 15% of the total time series for all the datasets as was suggested by the Numenta Benchmark (Ahmad et al., 2017). Considering this, the window sizes, w, required by the learning strategies—FR, SW, URES and ARES—are also set to $w = p$.

5.4   *Evaluation on Benchmark Datasets*

In this section, we first evaluate the average detection performances of different SAFARI methods, i.e., learning strategies and nonconformity measures across all the datasets. Then, we showcase how the best performances vary among 20 SAFARI detectors (the details are described in Appendix A).

19

Table 1 presents, for each learning strategy, the NAB and ROC-AUC scores that are aggregated over all datasets combining SAFARI detectors using the same method. More precisely, e.g., the ROC-AUC score of SAFARI-SW indicates the performance of the sliding window method as the mean and the standard deviation of the ROC-AUC results from four different SAFARI detectors that have SW as their learning strategy (i.e., SAFARI-SW-NN, SAFARI-SW-DEN, SAFARI-SW-CC, and SAFARI-SW-FREQ) (see Table A1). The results show that our proposed strategy, SAFARI-ARES, outperforms other methods in both ROC-AUC and NAB scores. SAFARI-FR, as expected, results in the lowest performance.

Table 1: Detection performances of SAFARI's learning strategies presented using three different metrics: ROC-AUC, NAB and average rank. Results compare the average performances of each method reported as the mean and the standard deviation of the scores taken from all datasets and detectors using that LS. The best average scores across each row of strategies are shown in bold.

Performance	SAFARI-FR	SAFARI-LW	SAFARI-SW	SAFARI-URES	SAFARI-ARES
ROC-AUC	0.781 ± 0.15	0.810 ± 0.13	0.828 ± 0.12	0.790 ± 0.14	$\mathbf{0.835 \pm 0.12}$
NAB	0.390 ± 0.37	0.637 ± 0.31	0.629 ± 0.30	0.559 ± 0.34	$\mathbf{0.660 \pm 0.28}$
Average Rank	3.71	2.85	2.70	3.16	**2.56**

Correspondingly, Table 2 shows the performance comparisons of different nonconformity measures (e.g., SAFARI-NN), averaged over all datasets and all SAFARI detectors with the same NCMs (e.g., SAFARI-FR-NN, SAFARI-SW-NN, SAFARI-LW-NN, SAFARI-URES-NN and SAFARI-ARES-NN). It can be seen that SAFARI-CC achieves the highest performance in ROC-AUC, while SAFARI-FREQ outperforms the others in terms of NAB score. SAFARI-NN consistently leads to the lowest performance.

Table 2: Detection performances of SAFARI's nonconformity measures presented using three different metrics: ROC-AUC, NAB and average rank. Results compare the average performances of each method reported as the mean and the standard deviation of the scores taken from all datasets and detectors using that NCM. The best average scores across each row of SAFARI-NCMs are shown in bold.

Performance	SAFARI-NN	SAFARI-DEN	SAFARI-CC	SAFARI-FREQ
ROC-AUC	0.767 ± 0.15	0.820 ± 0.13	$\mathbf{0.827 \pm 0.13}$	0.822 ± 0.13
NAB	0.484 ± 0.36	0.530 ± 0.34	0.623 ± 0.32	$\mathbf{0.663 \pm 0.29}$
Average Rank	3.10	2.53	**2.13**	2.22

To determine whether there is a significant difference between the performances of the different learn-

Fig. 3. Critical difference diagram showing the streaming anomaly detection performances of the five learning strategies. Methods that are not significantly different (at p ¡ 0.05) are connected with a bar.


Fig. 4. Critical difference diagram showing the streaming anomaly detection performances of the four nonconformity measures. Methods that are not significantly different (at p ¡ 0.05) are connected with a bar.

ing strategies and nonconformity measures, we follow Demšar (2006). We first apply the Friedman test (Friedman, 1937) using the average ranks of the methods in Table 1 and Table 2 where the null hypothesis for this test assumes that there is no significant difference between the methods. The Friedman tests for learning strategies and nonconformity measures returned p-values of $2.580007E - 16$ and $5.758827E - 12$, respectively. Therefore, we reject the null hypothesis in both cases and proceed with the Nemenyi post-hoc test (Nemenyi, 1963) to compare methods pairwise and to identify the ones that differ significantly. This test identifies performances of two algorithms to be significantly different if their average ranks differ by at least the "critical difference" (CD). Fig. 3 and 4 visually represent the results of the Nemenyi tests in critical difference diagrams where methods that are not connected by a bar have significantly different performances.

For the case of learning strategies, comparing five methods combined with four nonconformity measures on 125 datasets (i.e., 67 Yahoo, 58 Numenta) using two metrics (i.e., ROC-AUC and NAB) at significance level $\alpha = 0.05$, the critical difference diagram is shown in Fig. 3. It can be seen that SAFARI-FR performs significantly worse than other learning strategies, demonstrating that a fixed reference group is not a suitable for most of the streaming environments. Furthermore, SAFARI-ARES performs significantly better than SAFARI-FR, SAFARI-LW and SAFARI-URES while the difference between SAFARI-ARES and SAFARI-SW is not statistically significant.

Similarly, in Fig. 4 we can observe that SAFARI-NN performs significantly worse than the other methods, while there is no significant difference between SAFARI-CC and SAFARI-FREQ.

In the following, we present how the best performances vary between different SAFARI detectors. Table 3 shows, for each combination, the number of datasets for which it gives the best result (in either of the performance metrics). It can be seen that all the combinations achieve the highest performance for at least one dataset, except for SAFARI-FR-NN. Another important observation is that the superiority of a method can be different in terms of average detection performance and the number of best performances. For example, although there is no significant difference among the average performances of SAFARI-CC and SAFARI-FREQ (Fig. 4), the number of best performances that SAFARI-FREQ achieves is much higher. In addition, the results show that even the detectors that use the worst methods of the two worlds according

21

to the previous results (i.e., SAFARI-NN as a nonconformity measure or SAFARI-FR as a learning strategy) can achieve the best performances in multiple datasets.

Table 3: Comparison of the SAFARI detectors based on the number of datasets for which each detector is the winner—that is, outperforms all other detectors. According to results, SAFARI-FREQ-ARES is the detector (combination) with the most wins, with 31 cases. In total, SAFARI-FREQ and SAFARI-ARES are the methods with the highest number of best performances; their results are shown in bold.

Combination	SAFARI-NN	SAFARI-DEN	SAFARI-CC	SAFARI-FREQ	Total
SAFARI-FR	0	4	11	19	34
SAFARI-LW	4	11	5	25	45
SAFARI-SW	2	14	19	26	60
SAFARI-URES	2	7	12	12	33
SAFARI-ARES	2	11	17	31	**61**
Total	10	47	64	**113**	234

The common practice in the anomaly detection literature is comparing different methods based on their "average" performances on particular datasets, which is similar to the former results shown in the first part of this section. However, the latter results presented throughout this section show that none of the detectors is able to *consistently* perform better than all the other detectors. This suggests that different combinations are appropriate for different datasets or use cases, even though some of the methods work well more often than others or achieve higher performance on average. In the next sections, we try to highlight which method is likely to be successful under which circumstances.

5.5   *Comparison based on dataset characteristics*

In this section, we discuss and compare the behavior of algorithms across a wide range of datasets with different characteristics. The datasets are delineated based on four properties—noise, concept drift, anomaly type, and anomaly rate. We specifically analyze the individual performances of different nonconformity measures and learning strategies with respect to these properties. The goal is to provide the future users of SAFARI with insights into why combining particular methods may be beneficial or which component is more important for obtaining better results under specific conditions.

We first start by characterizing the datasets and evaluation metrics that are used in this study based on the collective performances of all SAFARI detectors. For this analysis, we examine the collective performances of all 20 SAFARI detectors and measure their "difficulty" and "diversity" levels. Following Zimek et al. (2012), we define the notion of "difficulty" as the average of the scores of all anomalies in the dataset calculated, across all methods. Datasets with a low difficulty score contain anomalies that are relatively easy to detect, while a high difficulty score indicates that the majority of methods have trouble finding the anomalies. On the other hand, "diversity" reflects the (lack of) agreement among the detectors on an individual dataset. We define the diversity score of a dataset as the standard deviations of the scores reported by all 20 combinations. A high diversity score indicates a large disagreement among the detection performances.

Figs. 5a and 5b show the difficulty–diversity plots using both evaluation metrics. Results from two different benchmarks are represented with different shapes. It can be seen that difficulty and diversity levels can vary greatly between datasets and evaluation metrics. Therefore, making fair comparisons of non-equivalent groups of datasets is not straightforward. For example, suppose we would like to assess the behavior of a method (e.g., sliding window) on a property (e.g., concept drift) by comparing the performance of this method on two groups of datasets: the first group includes "drifting" datasets, while the second group

includes nondrifting ones. Directly comparing the absolute performances (i.e., the ROC-AUC and NAB score) of the method on these two groups will not be a reliable way to analyze the impact of concept drift, since there can be other factors affecting the performances; in particular, one of the groups is likely to be inherently more difficult.

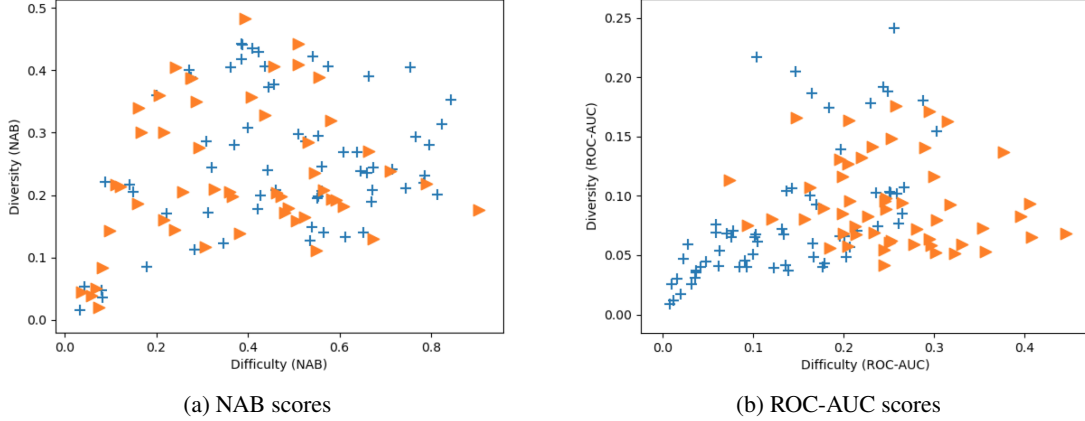(a) NAB scores   (b) ROC-AUC scores

Fig. 5. Diversity versus difficulty of the datasets based on two metrics: NAB and ROC-AUC. Numenta datasets are represented with orange triangles while Yahoo datasets are shown with blue crosses.

In this case, we try to mimic controlled experiments, while our test group (e.g., drifting) and the control group (e.g., non-drifting) have entirely different datasets and therefore, the number of independent variables (factors that are different between two groups) is unknown. To achieve this, we introduce the concept of "relative performance," where the goal is to account for the impact of uncontrollable factors while comparing algorithms performance on a specific property. The relative performance is computed by taking the average difference between the absolute performance of the method and the absolute performances of all the other methods.

It is assumed that the effects of uncontrollable factors also persist in the performance of the other methods, and computing the relative difference between two groups instead of the absolute difference will reduce the effect of this bias.

More formally, given a dataset d, let S be the list of the actual scores of a method M, and $\hat{S}$ be the list of actual scores of other methods. The relative performance score of M on d is

$$Rel_d^M = \frac{1}{|S_d|} \sum_{s \in S_d} \frac{\sum_{\hat{s} \in \hat{S}_d} s - \hat{s}}{max(S_d \cup \hat{S}_d) - min(S_d \cup \hat{S}_d)}. \tag{18}$$

Given two sets of datasets D_1 (e.g., low-noise datasets) and D_2 (e.g., high-noise datasets), the relative performance difference of M between D_1 and D_2 is

$$\Delta^M = \frac{\sum_{d_i \in D_1} Rel_{d_i}^M}{|D_1|} - \frac{\sum_{d_j \in D_2} Rel_{d_j}^M}{|D_2|}. \tag{19}$$

Table 4 reports relative performance scores of all SAFARI methods (i.e., five LS and four NCM) throughout different dataset properties. Each column represents how a method behaves under certain properties, such

as noise, concept drift, and so on. The significant score differences are marked in bold. In the following text we discuss in detail how different properties affect different SAFARI methods.

The noise effect: To compare the effect of noise in the data on the performances of different learning strategies and nonconformity measures, we divide benchmark datasets into two groups: low-noise and high-noise. However, the benchmarks do not provide information regarding the noise level of datasets. Therefore, we have determined this classification through visual analysis of each univariate time series in both Yahoo and Numenta datasets (see supplementary material).

The relative performance difference (Δ) scores in this setting reflect how the performance of a method changes from high-noise data to low-noise data, in comparison to other methods.

The first column in Table 4 shows these scores that are obtained by different SAFARI methods. It can be seen that the impact of noise is not significant in any of the given learning strategies. This result indicates that the choice of the learning strategies is not critical when the level of noise in a dataset is high. On the other hand, the performances of some of the nonconformity measures exhibit significant change under high noise. SAFARI-FREQ has the lowest (Δ) score, which reveals that its performance is the most negatively affected by the increase of noise. On the other hand SAFARI-NN and SAFARI-DEN do not show significant performance decreases between noisy and non-noisy datasets. SAFARI-CC is the most noise resilient method, achieving the highest Δ score. Considering these finding, SAFARI-CC has clear advantage when there is a clear sign of noise in a dataset, while SAFARI-FREQ should be avoided.

Drift effect: Similar to the previous case, the information about concept drift is missing, therefore we determine it by visual analysis. Following (Gama et al., 2014), we consider a dataset as drifting qualitatively if it has one of the drift types—that is, sudden, incremental, gradual, or reoccurring. The rest of the datasets are considered as non-drifting (again, see supplementary material). In this setting, a Δ score indicates how the performance of a method changes from drifting data to non-drifting data in comparison to other methods.

According to Table 4, the drift effect is quite distinct among different learning strategies. Δ scores show that SAFARI-SW and SAFARI-ARES are better than other methods at dealing with concept drift. Both of these methods have specific forgetting mechanisms, and clearly, forgetting past observations is essential when dealing with drift. The presence of drift affects SAFARI-FR the most, which is expected, considering that it is a static learning strategy that cannot adapt to changes over time.

According to Table 4, most of the nonconformity measures do not show significant performance change between drifting and non-drifting datasets. SAFARI-CC is an exception, exhibiting a clear decrease in performance when datasets are drifting. The explanation of this behavior might be our SAFARI-CC implementation. We use an incremental k-means algorithm that updates clusters over time according to the learning strategy. However, it still assigns a fixed number of clusters (k), and if a new concept emerges suddenly, the clustering structure may not adapt well enough to the new concept. This issue can be overcome using a different streaming clustering algorithm to measure nonconformity, one that can also change the number of clusters over time.

Table 4: Comparison of the SAFARI methods using relative performance scores across datasets with different characteristics: noise level, concept drift, anomaly type and anomaly rate (contamination).

DETECTOR	Δ_{noise}	Δ_{drift}	Δ_{type}	$\Delta_{contamination}$
SAFARI-FR	0.0169	$-\mathbf{0.1815}$	-0.0081	0.01817
SAFARI-LW	-0.01405	0.0178	0.0194	$-\mathbf{0.0279}$
SAFARI-SW	0.0179	$\mathbf{0.0977}$	-0.0189	-0.0184
SAFARI-URES	-0.0221	0.0279	-0.0070	$-\mathbf{0.0359}$
SAFARI-ARES	0.0012	$\mathbf{0.0362}$	0.0176	$\mathbf{0.0538}$
SAFARI-NN	0.0147	0.0129	0.0001	-0.0138
SAFARI-DEN	0.0172	0.0380	$-\mathbf{0.0391}$	0.0176
SAFARI-CC	$\mathbf{0.0216}$	-0.0624	-0.0237	0.0181
SAFARI-FREQ	$-\mathbf{0.0668}$	-0.0033	$\mathbf{0.0570}$	-0.0219

Anomaly type effect: We study the effect of two types of anomalies: clustered (pattern) anomalies and scattered anomalies (outliers). Clustered anomalies mostly occur when the same process generates anomalies multiple times, while scattered anomalies are often generated by different processes. To assess the clusteredness/scatteredness level of anomalies in each dataset, we use the normalized clusteredness measure proposed by Emmott et al. (2013). The normalized clusteredness nc is defined as $log\left(\frac{\sigma_n^2}{\sigma_a^2}\right)$, where σ_n^2 is the sample variance of the candidate normal points and σ_a^2 is the sample variance of the candidate anomalies. Then, we consider the anomaly type of a dataset as "scattered" if $nc \leq 0$ and "clustered" if $nc > 0$.

As reported in the third column of Table 4, the performances of the learning strategies do not show any significant difference when the type of the anomaly changes. However, the detection capabilities of different nonconformity measures can be influenced by anomaly type, since they mostly rely on different assumptions of the normality. The results support this argument by showing that most of the nonconformity measures integrated into SAFARI perform significantly differently on scattered and clustered anomalies. For example, the performances of SAFARI-DEN and SAFARI-CC deteriorate significantly when anomalies are clustered. Both of these methods assume that anomalies are located far away from the dense regions, and clustered anomalies can fool these methods by creating dense regions in the space. On the other hand, SAFARI-FREQ is clearly much better than the rest of the methods in handling clustered anomalies because it looks for the occurrence of the "rare" patterns rather than outlying individuals.

Anomaly rate effect: Anomaly rate reflects the contamination level of a dataset and is defined by the fraction of observations that are ground-truth anomalies. We divide the datasets into two groups as high and low contamination by considering the average contamination rate in all 112 datasets as a threshold. The datasets with higher rates than the average are categorized as high, while the rest as low contamination.

It can be observed from Table 4 that the anomaly rate profoundly affects the behavior of most of the learning strategies. The performances of SAFARI-LW, SAFARI-SW, and SAFARI-URES are significantly worsened when the contamination is high. The likely reason is that these methods learn from data instances without assessing whether they are actually normal observations. The greater the dataset contamination, the more anomalous the behavior these strategies learn. However, our proposed strategy, SAFARI-ARES, is

designed to use relative probabilities in order not to learn from potentially anomalous samples. The results show that it is clearly the best method to deal with datasets containing high anomaly rates.

SAFARI-FR also does not seem to be affected by the anomaly rate, which is understandable since it only learns during the probationary periods, which are defined in each dataset to contain only normal instances based on the ground truth. Still, we cannot recommend this strategy because the absolute performance scores of SAFARI-FR are much lower than the rest of the methods in the case of both low and high anomaly rates (see supplementary material).

Finally, no consistent performance change of nonconformity measures is observed between datasets with low and high anomaly rates.

5.6   Comparison with the baseline algorithms

In this section, we compare SAFARI with the state-of-the algorithms that are reported by Numenta benchmark. Table 5 summarizes the scores of benchmark algorithms across all application profiles (see supplementary material), including the three NAB competition winners (Ahmad et al., 2017). In addition to the various streaming anomaly detection algorithms, there are three control detectors in NAB. A "null" detector runs through the dataset passively, making no detections, accumulating all false negatives. A "perfect" detector is an oracle that outputs detections that would maximize the NAB score; that is, it outputs only true positives at the beginning of each window. The raw scores from these two detectors are used to scale the score for all other algorithms between 0 and 100. The "random" detector outputs a random anomaly probability for each data instance, which is then thresholded across the dataset for a range of random seeds. The score from this detector offers some intuition for chance-level performance on NAB.

Table 5: Comparison of SAFARI with algorithms in NAB scoreboard

Detector	Standard Profile	Reward Low FP	Reward Low FN
Perfect	100	100	100
SAFARI-Best	91.65	88.5	95.8
SAFARI-LW-CC	71.75	69.1	77.8
Numenta HTM	70.1	63.1	74.3
CAD-OSE	69.9	67	73.2
Numenta	64.6	58.8	69.6
KNN-CAD	58.0	43.4	64.8
SAFARI-Average	55.5	49.1	60.8
Relative Entropy	54.6	47.6	58.8
HTM PE	53.6	34.2	61.9
Random Cut Forest	51.7	38.4	59.7
Twitter ADVec	47.1	33.6	53.5
Etsy Skyline	35.7	27.1	44.5
Sliding Threshold	30.7	12.1	38.3
Bayesian Changepoint	17.7	3.2	32.2
EXPoSE	16.4	3.2	26.9
Random	11	1.2	19.5
Null	0	0	0

SAFARI-Best in Table 5 represents the best combination giving the highest NAB score in each dataset while SAFARI-Average reports the average NAB score of all the combinations. We have also reported

the best SAFARI detector across all NAB datasets, SAFARI-LW-CC, which combines distance to cluster centroids as a nonconformity measure and landmark window as a learning strategy.

Overall we can observe that SAFARI-Best and SAFARI-LW-CC outperform all other algorithms, while SAFARI-Average delivers competitive results. Numenta HTM, CAD-OSE, Numenta and KNN-CAD are the other detectors that perform well

6 Main Observations and Recommendations

According to the above comprehensive evaluations covering different aspects of anomaly detection, we can conclude that each approach has its own merits and weaknesses. In the following, we provide a summary of our findings and recommend for future SAFARI users potential ways to combine building blocks for specific cases.

First of all, SAFARI-ARES and SAFARI-SW as learning strategies and SAFARI-CC and SAFARI-FREQ as nonconformity measures outperform their competitors in terms of average performance across all the datasets. SAFARI-FR is the significantly worst method, which confirms the prior assumption that static learning is not suitable for streaming scenarios. On the other hand, it was unexpected to observe that SAFARI-NN performed significantly worse than the other nonconformity measures, since the nearest neighbor-based methods showed clear advantages in static datasets in the past (Aggarwal and Sathe, 2017). It is important to note that our experiments do not reflect the parameter sensitivity of the methods. We recommend users to refer to the studies by Aggarwal and Sathe (2017), Campos et al. (2016), and Goldstein and Uchida (2016) if they would like to consider the stability of the algorithms across a wide range of parameter choices.

From the perspective of different dataset properties, we observed that the choice of learning strategy is particularly important if datasets include concept drift or high anomaly rate. These properties can influence the performances of different learning strategies in different manners. While SAFARI-SW is the best method under concept drift, which shows the importance of adapting to the newest behavior, SAFARI-ARES also achieves competitive results. Furthermore, we recommend users choose SAFARI-ARES if the datasets are highly contaminated with abnormal samples or if it is difficult to obtain normal samples to initialize the model.

We have also found that the noise level and anomaly type of datasets have significant impacts on the performances of nonconformity measures, while we did not observe much effect on learning strategies. Specifically, SAFARI-CC is the most noise resilient method, while SAFARI-FREQ performs consistently worst under high noise. Regarding different types of anomalies, we recommend users consider SAFARI-CC and SAFARI-DEN for scattered anomalies and SAFARI-FREQ for anomalies that are more clustered.

7 Conclusion

In this paper, we introduced SAFARI, a framework for streaming anomaly detection based on building-blocks derived from fundamental concepts of this problem. By combining SAFARI's adaptive and extensible components, we produced 20 different anomaly detectors, a number of which are novel variants that, to the best of our knowledge, have never been tried before.

We have conducted comprehensive evaluation studies on these detectors using real-world benchmark datasets. We have discussed their merits and drawbacks thoroughly and drawn a set of interesting take-away conclusions. We have discovered that learning strategies should be chosen carefully for the cases where datasets are suspected of having concept drift or a high level of contamination. SAFARI-SW and SAFARI-

ARES are safer methods under concept drift, and SAFARI-ARES is the best option for highly contaminated datasets. Similarly, the selection of nonconformity measures is more critical if datasets include noise or different types of anomalies. Based on a detailed performance analysis, SAFARI-CC is recommended when the dataset has a high level of noise and anomalies are scattered, while SAFARI-FREQ is a better option for clustered anomalies.

The results have shown that there is no single superior detector that works well for every case and have proven our initial hypothesis that "there is no free lunch" in the streaming anomaly detection world. Furthermore, we have also showcased how SAFARI could help to ease this problem by empowering us to easily create use-case-specific detectors that are suitable for different scenarios instead of blindly relying on a single method.

Finally, we have postulated the problem of generalization and abstraction of streaming anomaly detection by considering similarities and differences in existing approaches. We believe that formally identifying core tasks as building blocks will help in understanding existing or new methods from a unified perspective and lead to identifying research gaps and unattended problems. With the help of SAFARI, we have discovered such a gap and formulated a new learning strategy specifically designed to handle high contamination while learning the normal group.

Acknowledgements

Appendix A

The details of how 20 SAFARI detectors that are used in our experiments are built as the combination of 12 SAFARI methods (i.e., 2 DR, 5 LS, 4 NCM, and 1 AS) can be found in Table A1. In this study, we mainly focus on learning strategy and nonconformity measure, as explained in Section 4. Therefore, we integrate different methods into the components of SAFARI dealing with these two tasks, while we implement only two data representation methods and one method for anomaly scoring. Furthermore, each data representation is combined with specific nonconformity measures in order to limit the number of combinations as 20. The first data representation, mean-std, is used where nonconformity measure is one of the proximity-based methods—nearest neighbors-based, density-based, and clustering-based, while SAX is found more suitable for frequency-based NCM considering its usefulness in capturing subpatterns in time-series Keogh et al. (2001b).

Table A1: The list of 20 SAFARI detectors as the combination of different SAFARI methods

Detectors	Data Representation	Learning Strategy	Non-conformity Measure	Anomaly Scoring
SAFARI-FR-NN	Mean-Std	Fixed reference group	Nearest neighbors-based	Final scoring
SAFARI-FR-DEN	Mean-Std	Fixed reference group	Density-based	Final scoring
SAFARI-FR-CC	Mean-Std	Fixed reference group	Clustering-based	Final scoring
SAFARI-FR-FREQ	SAX	Fixed reference group	Frequency-based	Final scoring
SAFARI-SW-NN	Mean-Std	Sliding Window	Nearest neighbors-based	Final scoring
SAFARI-SW-DEN	Mean-Std	Sliding Window	Density-based	Final scoring
SAFARI-SW-CC	Mean-Std	Sliding Window	Clustering-based	Final scoring
SAFARI-SW-FREQ	SAX	Sliding Window	Frequency-based	Final scoring
SAFARI-LW-NN	Mean-Std	Landmark Window	Nearest neighbors-based	Final scoring
SAFARI-LW-DEN	Mean-Std	Landmark Window	Density-based	Final scoring
SAFARI-LW-CC	Mean-Std	Landmark Window	Clustering-based	Final scoring
SAFARI-LW-FREQ	SAX	Landmark Window	Frequency-based	Final scoring
SAFARI-URES-NN	Mean-Std	Uniform Reservoir	Nearest neighbors-based	Final scoring
SAFARI-URES-DEN	Mean-Std	Uniform Reservoir	Density-based	Final scoring
SAFARI-URES-CC	Mean-Std	Uniform Reservoir	Clustering-based	Final scoring
SAFARI-URES-FREQ	SAX	Uniform Reservoir	Frequency-based	Final scoring
SAFARI-ARES-NN	Mean-Std	Anomaly-aware Reservoir	Nearest neighbors-based	Final scoring
SAFARI-ARES-DEN	Mean-Std	Anomaly-aware Reservoir	Density-based	Final scoring
SAFARI-ARES-CC	Mean-Std	Anomaly-aware Reservoir	Clustering-based	Final scoring
SAFARI-ARES-FREQ	SAX	Anomaly-aware Reservoir	Frequency-based	Final scoring

Appendix B

We empirically studied how runtime of SAFARI methods varies with the dataset size. All experiments were performed on an OSX personal computer with 16GB memory. Runtimes were averaged over five trials. To measure scalability with respect to the size of the dataset, we sampled the number of observations between [4000, 20000] in equal intervals from a single large data file in NAB, consisting of 22,695 data records. The scalability of different learning strategies and nonconformity measures can be seen in Fig. B1a and Fig. B1b, respectively.

29

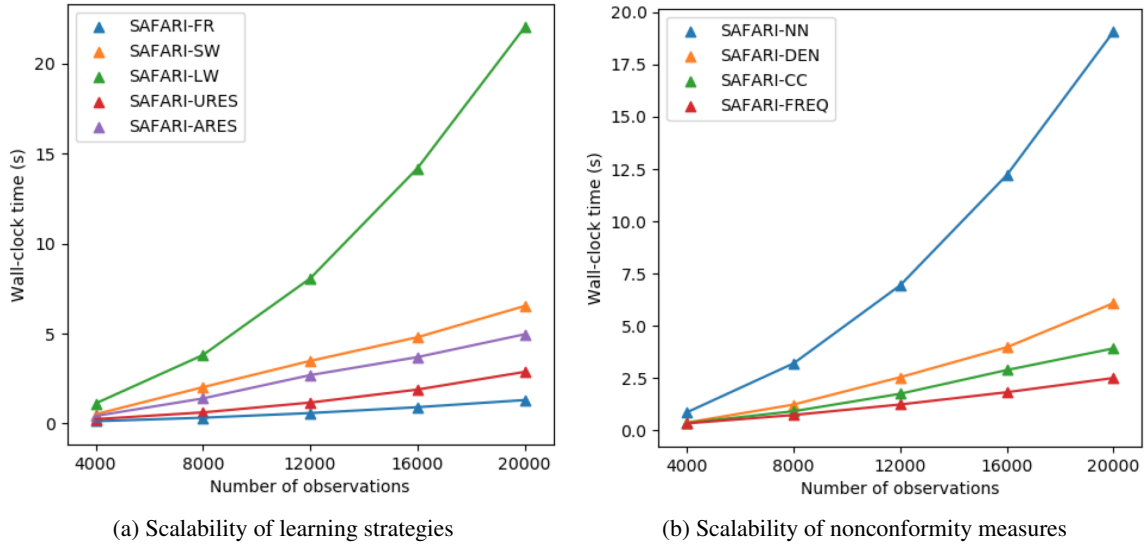(a) Scalability of learning strategies   (b) Scalability of nonconformity measures

Fig. B1. Scalability of different SAFARI methods: wall-clock time of (a) learning strategies (b) nonconformity measure against number of observations in a dataset.

Furthermore, we analyzed latency times of each SAFARI detectors used in this study. Latency measures the time taken to process a single data point for anomaly detection. Latency times of detectors are also averaged over 5 runs on the same NAB dataset, consisting of 22,695 data records and shown in Fig. B2.


Fig. B2. Comparison of 20 SAFARI detectors based on latency (ms).

References

Aggarwal, C.C., 2015. Outlier analysis, in: Data mining, Springer. pp. 237–263.

Aggarwal, C.C., Han, J., Wang, J., Yu, P.S., 2003. A framework for clustering evolving data streams, in:

Proceedings of the 29th international conference on Very large data bases-Volume 29, VLDB Endowment. pp. 81–92.

Aggarwal, C.C., Sathe, S., 2017. Outlier ensembles: An introduction. Springer.

Ahmad, S., Lavin, A., Purdy, S., Agha, Z., 2017. Unsupervised real-time anomaly detection for streaming data. Neurocomputing 262, 134–147.

Angiulli, F., Fassetti, F., 2007. Detecting distance-based outliers in streams of data, in: Proceedings of the sixteenth ACM conference on Conference on information and knowledge management, ACM. pp. 811–820.

Assent, I., Kranen, P., Baldauf, C., Seidl, T., 2012. Anyout: Anytime outlier detection on streaming data, in: International Conference on Database Systems for Advanced Applications, Springer. pp. 228–242.

Bifet, A., Gavalda, R., 2006. Kalman filters and adaptive windows for learning in data streams, in: International Conference on Discovery Science, Springer. pp. 29–40.

Bifet, A., Gavalda, R., 2007. Learning from time-changing data with adaptive windowing, in: Proceedings of the 2007 SIAM international conference on data mining, SIAM. pp. 443–448.

Bulut, A., Singh, A.K., 2003. Swat: Hierarchical stream summarization in large networks, in: Proceedings 19th International Conference on Data Engineering (Cat. No. 03CH37405), IEEE. pp. 303–314.

Campos, G.O., Zimek, A., Sander, J., Campello, R.J., Micenková, B., Schubert, E., Assent, I., Houle, M.E., 2016. On the evaluation of unsupervised outlier detection: measures, datasets, and an empirical study. Data Mining and Knowledge Discovery 30, 891–927.

Cao, F., Estert, M., Qian, W., Zhou, A., 2006. Density-based clustering over an evolving data stream with noise, in: Proceedings of the 2006 SIAM international conference on data mining, SIAM. pp. 328–339.

Chandola, V., Banerjee, A., Kumar, V., 2009. Anomaly detection: A survey. ACM computing surveys (CSUR) 41, 15.

Chen, H., Tiňo, P., Rodan, A., Yao, X., 2013. Learning in the model space for cognitive fault diagnosis. IEEE transactions on neural networks and learning systems 25, 124–136.

Chenaghlou, M., Moshtaghi, M., Leckie, C., Salehi, M., 2017. An efficient method for anomaly detection in non-stationary data streams, in: GLOBECOM 2017-2017 IEEE Global Communications Conference, IEEE. pp. 1–6.

Cohen, E., Strauss, M., 2003. Maintaining time-decaying stream aggregates, in: Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, ACM. pp. 223–233.

Cormode, G., Garofalakis, M., Sacharidis, D., 2006. Fast approximate wavelet tracking on streams, in: International Conference on Extending Database Technology, Springer. pp. 4–22.

Demšar, J., 2006. Statistical comparisons of classifiers over multiple data sets. Journal of Machine learning research 7, 1–30.

D'Silva, S.H., 2008. Diagnostics based on the statistical correlation of sensors. SAE International Journal of Passenger Cars-Electronic and Electrical Systems 1, 53–61.

Efraimidis, P.S., Spirakis, P.G., 2006. Weighted random sampling with a reservoir. Information Processing Letters 97, 181–185.

Emmott, A., Das, S., Dietterich, T., Fern, A., Wong, W.K., 2015. A meta-analysis of the anomaly detection problem. arXiv preprint arXiv:1503.01158 .

Emmott, A.F., Das, S., Dietterich, T., Fern, A., Wong, W.K., 2013. Systematic construction of anomaly detection benchmarks from real data, in: Proceedings of the ACM SIGKDD workshop on outlier detection and description, ACM. pp. 16–21.

Faloutsos, C., Ranganathan, M., Manolopoulos, Y., 1994. Fast subsequence matching in time-series databases. volume 23. ACM.

Fan, Y., Nowaczyk, S., Rögnvaldsson, T., 2015. Evaluation of self-organized approach for predicting compressor faults in a city bus fleet. Procedia Computer Science 53, 447–456.

Friedman, M., 1937. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. Journal of the american statistical association 32, 675–701.

Gama, J.a., Žliobaitė, I., Bifet, A., Pechenizkiy, M., Bouchachia, A., 2014. A survey on concept drift adaptation. ACM Comput. Surv. 46, 44:1–44:37. URL: http://doi.acm.org/10.1145/2523813, doi:10.1145/2523813.

Gao, J., Tan, P.N., 2006. Converting output scores from outlier detection algorithms into probability estimates, in: Sixth International Conference on Data Mining (ICDM'06), IEEE. pp. 212–221.

Geurts, P., 2001. Pattern extraction for time series classification, in: European Conference on Principles of Data Mining and Knowledge Discovery, Springer. pp. 115–127.

Gilbert, A.C., Kotidis, Y., Muthukrishnan, S., Strauss, M.J., 2003. One-pass wavelet decompositions of data streams. IEEE Transactions on knowledge and data engineering 15, 541–554.

Goldstein, M., Uchida, S., 2016. A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. PloS one 11, e0152173.

Gupta, M., Gao, J., Aggarwal, C.C., Han, J., 2013. Outlier detection for temporal data: A survey. IEEE Transactions on Knowledge and data Engineering 26, 2250–2267.

Hawkins, D.M., 1980. Identification of outliers. volume 11. Springer.

Haykin, S., 1996. Adaptive Filter Theory (3rd Ed.). Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

Ho, S.S., 2005. A martingale framework for concept change detection in time-varying data streams, in: Proceedings of the 22nd international conference on Machine learning, ACM. pp. 321–327.

Ho, S.S., Wechsler, H., 2010. A martingale framework for detecting changes in data streams by testing exchangeability. IEEE transactions on pattern analysis and machine intelligence 32, 2113–2127.

Huang, L., Nguyen, X., Garofalakis, M., Jordan, M.I., Joseph, A., Taft, N., 2007. In-network pca and anomaly detection, in: Advances in Neural Information Processing Systems, pp. 617–624.

Hundman, K., Constantinou, V., Laporte, C., Colwell, I., Soderstrom, T., 2018. Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding, in: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ACM. pp. 387–395.

Jiang, R., Fei, H., Huan, J., 2011. Anomaly localization for network data streams with graph joint sparse pca, in: Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM. pp. 886–894.

Jin, R., Agrawal, G., 2005. An algorithm for in-core frequent itemset mining on streaming data, in: Fifth IEEE International Conference on Data Mining (ICDM'05), IEEE. pp. 8–pp.

Kargupta, H., Bhargava, R., Liu, K., Powers, M., Blair, P., Bushra, S., Dull, J., Sarkar, K., Klein, M., Vasa, M., et al., 2004. Vedas: A mobile and distributed data stream mining system for real-time vehicle monitoring, in: Proceedings of the 2004 SIAM International Conference on Data Mining, SIAM. pp. 300–311.

Kargupta, H., Gilligan, M., Puttagunta, V., Sarkar, K., Klein, M., Lenzi, N., Johnson, D., 2010. Minefleet: The vehicle data stream mining system for ubiquitous environments, in: Ubiquitous knowledge discovery. Springer, pp. 235–254.

Kargupta, H., Puttagunta, V., Klein, M., Sarkar, K., 2006. On-board vehicle data stream monitoring using minefleet and fast resource constrained monitoring of correlation matrices. New Generation Computing 25, 5–32.

Keogh, E., Chakrabarti, K., Pazzani, M., Mehrotra, S., 2001a. Dimensionality reduction for fast similarity search in large time series databases. Knowledge and information Systems 3, 263–286.

Keogh, E., Chakrabarti, K., Pazzani, M., Mehrotra, S., 2001b. Locally adaptive dimensionality reduction for indexing large time series databases. ACM Sigmod Record 30, 151–162.

Kollios, G., Gunopulos, D., Koudas, N., Berchtold, S., 2003. Efficient biased sampling for approximate clustering and outlier detection in large data sets. IEEE Transactions on Knowledge and Data Engineering 15, 1170–1187.

Kolmogorov, A.L., 1933. Sulla determinazione empirica di una legge di distribuzione. G. Ist. Ital. Attuari 4, 83–91. URL: https://ci.nii.ac.jp/naid/10030673552/en/.

Kontaki, M., Gounaris, A., Papadopoulos, A.N., Tsichlas, K., Manolopoulos, Y., 2011. Continuous monitoring of distance-based outliers over data streams, in: 2011 IEEE 27th International Conference on Data Engineering, IEEE. pp. 135–146.

Kriegel, H.P., Kroger, P., Schubert, E., Zimek, A., 2011. Interpreting and unifying outlier scores, in: Proceedings of the 2011 SIAM International Conference on Data Mining, SIAM. pp. 13–24.

Kuncheva, L.I., 2011. Change detection in streaming multivariate data using likelihood detectors. IEEE transactions on knowledge and data engineering 25, 1175–1180.

Laptev, N., Amizadeh, S., Flint, I., 2015. Generic and scalable framework for automated time-series anomaly detection, in: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM. pp. 1939–1947.

Lavin, A., Ahmad, S., 2015. Evaluating real-time anomaly detection algorithms–the numenta anomaly benchmark, in: 2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA), IEEE. pp. 38–44.

Lazaridis, I., Mehrotra, S., 2003. Capturing sensor-generated time series with quality guarantees, in: Proceedings 19th International Conference on Data Engineering (Cat. No. 03CH37405), IEEE. pp. 429–440.

Leung, C.K.S., Jiang, F., 2011. Frequent pattern mining from time-fading streams of uncertain data, in: International Conference on Data Warehousing and Knowledge Discovery, Springer. pp. 252–264.

Li, J., Struzik, Z., Zhang, L., Cichocki, A., 2015. Feature learning from incomplete eeg with denoising autoencoder. Neurocomputing 165, 23–31.

Lin, J., Keogh, E., Lonardi, S., Chiu, B., 2003. A symbolic representation of time series, with implications for streaming algorithms, in: Proceedings of the 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery, ACM. pp. 2–11.

Liu, F.T., Ting, K.M., Zhou, Z.H., 2012. Isolation-based anomaly detection. ACM Transactions on Knowledge Discovery from Data (TKDD) 6, 3.

Malhotra, P., Vig, L., Shroff, G., Agarwal, P., 2015. Long short term memory networks for anomaly detection in time series, in: Proceedings, Presses universitaires de Louvain. p. 89.

Manku, G.S., Motwani, R., 2012. Approximate frequency counts over data streams. Proceedings of the VLDB Endowment 5, 1699–1699. doi:`10.14778/2367502.2367508`.

Maurus, S., Plant, C., 2017. Let's see your digits: Anomalous-state detection using benford's law, in: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM. pp. 977–986.

Na, G.S., Kim, D., Yu, H., 2018. Dilof: Effective and memory efficient local outlier detection in data streams, in: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ACM. pp. 1993–2002.

Nemenyi, P., 1963. Distribution-free Multiple Comparisons. Princeton University. URL: `https://books.google.com.tr/books?id=nhDMtgAACAAJ`.

Olsson, T., Holst, A., 2015. A probabilistic approach to aggregating anomalies for unsupervised anomaly detection with industrial applications, in: The Twenty-Eighth International Flairs Conference.

Ordonez, C., 2003. Clustering binary data streams with k-means, in: Proceedings of the 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery, ACM. pp. 12–19.

Pokrajac, D., Lazarevic, A., Latecki, L.J., 2007. Incremental local outlier detection for data streams, in: 2007 IEEE symposium on computational intelligence and data mining, IEEE. pp. 504–515.

Quevedo, J., Chen, H., Cugueró, M.À., Tino, P., Puig, V., Garciá, D., Sarrate, R., Yao, X., 2014. Combining learning in model space fault diagnosis with data validation/reconstruction: Application to the barcelona water network. Engineering Applications of Artificial Intelligence 30, 18–29.

Rodríguez, J.J., Alonso, C.J., 2004. Support vector machines of interval-based features for time series classification, in: International Conference on Innovative Techniques and Applications of Artificial Intelligence, Springer. pp. 244–257.

Rögnvaldsson, T., Nowaczyk, S., Byttner, S., Prytz, R., Svensson, M., 2018. Self-monitoring for maintenance of vehicle fleets. Data mining and knowledge discovery 32, 344–384.

Salehi, M., Leckie, C., Bezdek, J.C., Vaithianathan, T., Zhang, X., 2016. Fast memory efficient local outlier detection in data streams. IEEE Transactions on Knowledge and Data Engineering 28, 3246–3260.

Schubert, E., Zimek, A., Kriegel, H.P., 2014. Local outlier detection reconsidered: a generalized view on locality with applications to spatial, video, and network outlier detection. Data Mining and Knowledge Discovery 28, 190–237.

Song, X., Wu, M., Jermaine, C., Ranka, S., 2007. Statistical change detection for multi-dimensional data, in: Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM. pp. 667–676.

Subramaniam, S., Palpanas, T., Papadopoulos, D., Kalogeraki, V., Gunopulos, D., 2006. Online outlier detection in sensor data using non-parametric models, in: Proceedings of the 32nd international conference on Very large data bases, VLDB Endowment. pp. 187–198.

Sugiyama, M., Borgwardt, K., 2013. Rapid distance-based outlier detection via sampling, in: Advances in Neural Information Processing Systems, pp. 467–475.

Vachkov, G., 2006. Intelligent data analysis for performance evaluation and fault diagnosis in complex systems, in: 2006 IEEE International Conference on Fuzzy Systems, IEEE. pp. 1213–1220.

Vitter, J.S., 1985. Random sampling with a reservoir. ACM Transactions on Mathematical Software (TOMS) 11, 37–57.

Volkhonskiy, D., Nouretdinov, I., Gammerman, A., Vovk, V., Burnaev, E., 2017. Inductive conformal martingales for change-point detection. arXiv preprint arXiv:1706.03415 .

Vovk, V., Gammerman, A., Shafer, G., 2005. Algorithmic learning in a random world. Springer Science & Business Media.

Yamanishi, K., Takeuchi, J.i., 2002. A unifying framework for detecting outliers and change points from non-stationary time series data, in: Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, ACM. pp. 676–681.

Yamanishi, K., Takeuchi, J.I., Williams, G., Milne, P., 2004. On-line unsupervised outlier detection using finite mixtures with discounting learning algorithms. Data Mining and Knowledge Discovery 8, 275–300.

Yang, W., Tavner, P.J., Crabtree, C.J., Wilkinson, M., 2009. Cost-effective condition monitoring for wind turbines. IEEE Transactions on industrial electronics 57, 263–271.

Yi, B.K., Faloutsos, C., 2000. Fast time sequence indexing for arbitrary lp norms, in: VLDB, Citeseer. p. 99.

Yi, B.K., Sidiropoulos, N.D., Johnson, T., Jagadish, H., Faloutsos, C., Biliris, A., 2000. Online data mining for co-evolving time sequences, in: Proceedings of 16th International Conference on Data Engineering (Cat. No. 00CB37073), IEEE. pp. 13–22.

Zhang, T., Ramakrishnan, R., Livny, M., 1996. Birch: an efficient data clustering method for very large databases, in: ACM Sigmod Record, ACM. pp. 103–114.

Zhang, X., Salehi, M., Leckie, C., Luo, Y., He, Q., Zhou, R., Kotagiri, R., 2018. Density biased sampling with locality sensitive hashing for outlier detection, in: International Conference on Web Information Systems Engineering, Springer. pp. 269–284.

Zhu, Y., Shasha, D., 2002. Statstream: Statistical monitoring of thousands of data streams in real time, in: VLDB'02: Proceedings of the 28th International Conference on Very Large Databases, Elsevier. pp. 358–369.

Zimek, A., Gaudet, M., Campello, R.J., Sander, J., 2013. Subsampling for efficient and effective unsupervised outlier detection ensembles, in: Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM. pp. 428–436.

Zimek, A., Schubert, E., Kriegel, H.P., 2012. A survey on unsupervised outlier detection in high-dimensional numerical data. Statistical Analysis and Data Mining: The ASA Data Science Journal 5, 363–387.