TRACKING CHANGES USING KULLBACK-LEIBLER DIVERGENCE FOR THE CONTINUAL LEARNING *

Ⓘ **Sebastián Basterrech**
Faculty of Electrical Engineering and Computer Science
VŠB-Technical University of Ostrava
Ostrava, Czech Republic
Sebastian.Basterrech@vsb.cz

Ⓘ **Michał Woźniak**
Department of Systems and Computer Networks
Wroclaw University of Science and Technology
Wroclaw, Poland
Michal.Wozniak@pwr.edu.pl

ABSTRACT

Recently, continual learning has received a lot of attention. One of the significant problems is the occurrence of *concept drift*, which consists of changing probabilistic characteristics of the incoming data. In the case of the classification task, this phenomenon destabilizes the model's performance and negatively affects the achieved prediction quality. Most current methods apply statistical learning and similarity analysis over the raw data. However, similarity analysis in streaming data remains a complex problem due to time limitation, non-precise values, fast decision speed, scalability, etc. This article introduces a novel method for monitoring changes in the probabilistic distribution of multi-dimensional data streams. As a measure of the rapidity of changes, we analyze the popular Kullback-Leibler divergence. During the experimental study, we show how to use this metric to predict the concept drift occurrence and understand its nature. The obtained results encourage further work on the proposed methods and its application in the real tasks where the prediction of the future appearance of concept drift plays a crucial role, such as predictive maintenance.

Keywords Drift detection · Continual Learning · Relative Entropy · Data Stream Learning · Lifelong Learning

1 Introduction

One of the critical problems in the analysis of streaming data is the possibility of changing the probabilistic characteristics of the task during the prediction model run. This phenomenon, called *concept drift* [29], may deteriorate the classification quality of the using predictor, and usually, its occurrence is unpredictable. Its appearance is typical in many daily decision-making tasks, as fraudsters may change the content of the e-mail to get past spam filters.

We may propose several *concept drift* taxonomies. The first one focuses on how the drift impacts the probability characteristics of the learning task. If it changes decision boundary shapes, i.e., the *posterior* probabilities have been changed [24] then we face with so-called *real concept drift*. *Virtual drift* does not change the decision boundary shape. However, it changes the unconditional probability density function [28]. Nevertheless, Oliveira et al. [20] pointed out

---

that while *virtual concept drift* does not change the shape of decision boundaries, it can reduce the usefulness of the decision boundaries used by classifiers.

Another taxonomy considers the change speed. Mainly, we may distinguish sudden *concept drift*, when a new concept abruptly replaces an old one; and incremental *concept drift*, when we may observe a steady progression from an old concept toward a new one. We should also mention gradual *concept drift*, when during the transition between an old and a new concept, the two concepts may occur with different intensities. An interesting phenomenon is periodic changes, referred to as recurring *concept drift* [27], when previously occurring concepts may recur, which is very typical for seasonal phenomena (cyclic concept [14]).

Currently, most approaches are reactive, i.e., they focus on the problem of concept drift detection or continuous model adaptation to emerging changes. These approaches include concept drift detectors, which act as triggers for significant shifts in the probability distribution, and methods that continuously try to adapt to changes in the probability distribution. The second group of the mentioned methods mainly use the classifier ensemble [15], where the model consistency with the current probability distribution is ensured either by changing the composition of the classifier ensemble or updating individual base classifiers.

This paper focuses on the *concept drift* monitoring problem and proposes a novel drift detection method. The significant advantage of the proposed approach is that it can compare distributions of unknown types to detect a drift without any assumptions about distribution parameters or distribution types.

This work offers the following contributions:

- Using Kullback-Leibler divergence for ongoing monitoring of changes in probability distributions.
- Introducing a novel *concept drift* detector based on the mentioned metric over the raw data, so-called KL-divergence- based concept drift detector (KLD).
- Fast and robust decision rule with a control real-value parameter that can be dynamically adjusted for improving performance of the matching matrix.
- Initial experimental evaluation of the proposed method for non-stationary data streams.

The remainder of the paper is organized as follows. The next section presents a background over the main concepts of this article. Section 3 fully describes our contribution. Experimental study is presented in Section 4. We close with a summary of our contributions and discussing directions for future work.

2 Preliminaries

The following section presents the key information about data stream, concept drift, and Kullback-Leibler divergence necessary to explain the proposed approach.

2.1 Streaming data

Continual learning of streaming data is usually associated to problems, where datasets become available in chunks. A data stream often is defined as an ordered sequence of chunks $\{S_1, S_2, \ldots, S_k, \ldots\}$, where the chunk S_i consists in a finite data sequence $S_i = \{\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}, \ldots, \mathbf{z}_i^{(K)}\}$. Since our goal is to analyze continual learning over a data stream in a supervised context, then we assume data points $\mathbf{z}_i^{(k)}$ defined as input-output pairs $(\mathbf{u}_i^{(k)}, \mathbf{y}_i^{(k)}) \in \mathbb{U} \times \mathbb{Y}$. The input vectors are in a p-dimensional space $\mathbb{U}$, and the outputs belongs to a d-dimensional space $\mathbb{Y}$. In case of a classification problem, the space $\mathbb{Y}$ describes a small number of classes. Otherwise, it is a regression problem where $\mathbb{Y}$ is typically a subset in a d-dimensional real space.

Tracking changes in a data stream often are made using drift detection methods. Most of the drift detection techniques use a base predictor to classify incoming instances. In general, the methods work as follows. For each input instance

the base predictor outputs a class label which is compared to the true class label. Then, the accuracy is evaluated and used as tool for deciding whether a drift has occurred or not. There are several metrics to evaluate the accuracy of such prediction. The choice of the selected metric, conventionally called loss function, depends on both the domain of $\mathbb{Y}$ and the optimization method applied for adjusting the base predictor. Some commonly used functions are $0/1$ loss function and cross-entropy in the case of classification problems and quadratic errors in the case of regression problems [25].

2.2 Drift detection

A concept drift detector is an algorithm that can inform on data distribution changes. To detect a real drift, the labeled data or a classifier's performance are usually required, but also some proposition employing unlabeled data only, usually based on statistical tests as Smirnov-Kolmogorov are also proposed [27]. Nevertheless, we should be aware of its limitations. Such detectors are good at detecting virtual concept drift and are not guaranteed to detect real drift. However, changes in the unconditional probability density distributions are rarely not accompanied by changes in the conditional distributions. Moreover, as mentioned in [20], virtual drift often affects the utility of the prediction models used.

Some drift detectors can also return a warning signal, i.e., the distribution changing, but the change is insignificant. A warning signal could trigger collecting new data for upcoming updates or rebuilding the current predictor.

Detecting the drift should be done as quickly as possible to replace an outdated model, thus minimizing the restoration time. On the other hand, false positives are also unacceptable because they can cause the model to be corrected unnecessarily, resulting in model adjustment to unrepresentative samples and increased, unwarranted consumption of computational resources [13].

Nevertheless, according to the authors' best knowledge, the recommended drift detector's testing method should focus on the predictive performance of the classifier that employs the examined detectors. It is recommended to use the following framework [12].

Let us shortly review the most important drift detection algorithms. CUSUM *(Cumulative sum)* [21] is a simple sequential analysis technique based on the measurement of the mean value of the input data. If the mean is significantly bigger than zero drift is. *Exponentially Weighted Moving Average (*EWMA*)* [23] combines current and historical observations that allow detecting changes in the mean value quickly by using aggregated charting statistics. A weight factor is introduced that promotes the most current observations. DDM *(Drift Detection Method)* [11] incrementally estimates a classifier performance, which (assuming the convergence of the learning algorithm) has to diminish along with new instances from the given distribution being continuously presented to the learner [22]. If the reverse behavior is observed, we may suspect a change of probability distributions. EDDM *(Early Drift Detection Methods)* [1] extends the DDM. It proposed the heuristic window size selection procedure, used different performance metric to follow, and implemented new warning and drift levels . Blanco et al. [5] propose detectors that use the non-parametric estimation of classifier error employing Hoeffding's and McDiarmid's inequalities. ADWIN [4] employs an adapting sliding window, being highly suitable for handling sudden drifts. When no change is apparent, it automatically grows the window size and shrinks it when data stream distribution changes. It tests the hypothesis about the average equality in two subwindows obtained using Hoeffding bounds. Nishida and Yamauchi [19] developed STEPD that employs statistics, equivalent to the chi-square test with Yates's continuity correction to infer whether the classification quality of the current observations had changed since the current model was used. We have also mention the compound detection models which employ ensemble approach, that could be found, e.g., in [18], [10], [17]. Concept drift detection has also been developed over high-dimensional and sparse spaces (e.g. sparse time series) [30, 26].

2.3 Kullback-Leibler divergence

A function for assessing the similarity between two distributions is the relative entropy, so-called, *Kullback-Leibler (KL)* divergence. The KL divergence between two probability mass functions (pmf) $p(x)$ and $q(x)$ over a set $\mathbb{X}$ is defined as follows [9]:

$$\mathrm{KL}(p(x) \parallel q(x)) = \sum_{x \in \mathbb{X}} p(x) \log \frac{p(x)}{q(x)}, \tag{1}$$

where the $\log(\cdot)$ is logarithmic function with an arbitrary selected log base (most commonly used 2 or e).

Expression (1) is zero if and only if both distributions have identical quantities of information. An interpretation coming from *Information Theory* is that $KL(p(x) \parallel q(x))$ gives the information lost when $q(x)$ is used to approximate $p(x)$. When both distributions are strictly positive, then KL divergence value is always a positive real number. However, expression (1) is problematic when either $p(x)$ or $q(x)$ has zero values. The expression goes to an infinity value when $p(x)$ tends to zero, and it may takes undefined value when $q(x) = 0$. A common convention is to define $0 \log(0) = 0 \log(0/0) = 0$ and $x \log(x/0) = \infty$ if $x > 0$ [8]. A common strategy for avoiding numerical problems consists in applying a smoothing approach of the empirical distribution, which is done by adding a stable fixed noise to each element with a probability mass equal to zero. Another important property of the KL divergence is that is not symmetric. Neither satisfies triangle inequality for any pair of distributions $p(x)$ and $q(x)$ [7]. Therefore, it is not strictly a metric. As a consequence, it can be inconvenient to be used in some applications [9].

In spite of the described drawbacks, the KL divergence has also several advantages that makes expression (1) suitable for being used as a similarity measure between two distributions. Here, we further discuss some of those helpful properties [7, 9].

- Relationship with likelihood ratio: KL divergence is often presented as the expected value of the log likelihood ratio,

$$\mathrm{KL}(p(x) \parallel q(x)) = \mathbb{E}_p \left(\log \frac{p(x)}{q(x)} \right). \tag{2}$$

  This definition can be specially useful when we know the distribution from the data, and it is easy to compute the log likelihood ratio.

- Relationship with Maximum Likelihood Estimation (MLE): MLE can be viewed as minimizing the KL divergence between the empirical distribution data and the proposed model distribution.

- Hypothesis tests: several hypothesis tests are equivalent to KL divergence. The statistic of t-test is the KL distance between two Gaussian distributions. The first term of the Taylor expansion of KL divergences is equivalent to the statistic of chi-square test. Furthermore, the Kulldorff spatial test used for clustering data has statistic that is the KL-distance between the empirical data distribution and the Poisson distribution.

- Pinsker's inequality: total variation distance between two densities, so called statistical distance, it is computed as a function of KL divergence of those two densities.

3 KL-divergence-based concept drift detector (KLD)

This section presents a pipeline for a KL-divergence-based concept drift detector (KLD). Firstly, we show how to determine the empirical distribution from streaming data, calculate KL-divergence between two data chunks, and detect the drift using a simple threshold scheme.

3.1 Creation of empirical distribution from a data stream

Each chunk defines an empirical distribution, and we compute the similarity between two adjacent chunks using KL divergence. Following the previous notation, let's define the data presented in chunk S_i as a sequence of input-output

pairs $\mathbf{z}_i^{(k)} = (\mathbf{u}_i^{(k)}, \mathbf{y}_i^{(k)})$, where $\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}, \ldots, \mathbf{z}_i^{(K)}$ is a random sample from a discrete distribution with an unknown probability mass function over a discrete set $\mathbb{U} \times \mathbb{Y}$.

We assume unidimensional output space. For the probability mass function estimation, we divide the input space $\mathbb{U}$ in a multidimensional regular grid composed by J bins. Let $\mathbf{b}_i^j$ be the bin j in the chunk i. A bin can be formalized defining a parameterized range constraint, and each point $\mathbf{u} \in \mathbb{U}$ could be assigned to a specific bin according to the parameterized range constraint. In a chunk S_i the conditional probability of having a class l in bin $\mathbf{b}_i^j$ is estimated by

$$\Pr\left(\mathbf{y}_i = l | \{\mathbf{u}_i : \mathbf{u}_i^k \in \mathbf{b}_i^j, \forall k\}\right) = \frac{\sum_{k=1}^K \mathbb{1}_{\{\mathbf{u}_i^k \in \mathbf{b}_i^j\} \cap \{\mathbf{y}_i^k = l\}}}{\sum_{k=1}^K \mathbb{1}_{\{\mathbf{u}_i^k \in \mathbf{b}_i^j\}}}, \tag{3}$$

where $\mathbb{1}_{\{A\}}$ is 1 if predicate A is true, 0 otherwise. Then, by applying expression (3) for each output class l is computed the probability mass function $\Pr(\mathbf{y}_i | \{\mathbf{u}_i : \mathbf{u}_i^k \in \mathbf{b}_i^j, \forall k\})$.

3.2 Proposed KL-divergence-based similarity metric

Hence, each bin has associated a probability mass function, then we may see changes through any two chunks S_i and S_t computing

$$d_{i,t}^j = KL\left(\Pr(\mathbf{y}_i | \{\mathbf{u}_i : \mathbf{u}_i^k \in \mathbf{b}_i^j, \forall k\}) \parallel \Pr(\mathbf{y}_t | \{\mathbf{u}_t : \mathbf{u}_t^k \in \mathbf{b}_t^j, \forall k\})\right). \tag{4}$$

Finally, similarity between any two chunks S_i and S_t is defined

$$D(S_i, S_t) = \frac{1}{J} \sum_{j=1}^J d_{i,t}^j. \tag{5}$$

Note that, in expression (5) each bin has the same relevance in the total sum. We present a slight variation that considers the probability of sampling points in each bin. Then, the proposed similarity measure for comparing two chunks is given by the weighted sum

$$D(S_i, S_t) = \frac{1}{J} \sum_{j=1}^J \gamma_i^j d_{i,t}^j, \tag{6}$$

where the weight γ_i^j is the probability estimation of sampling in a specific region $\mathbf{b}_i^j$ of the reference chunk S_i

$$\gamma_i^j = \frac{1}{K} \sum_{k=1}^K \mathbb{1}_{\{\mathbf{u}_i^k \in \mathbf{b}_i^j\}}.$$

3.3 Decision rule using a threshold value

In this first work, we analyze similarity between consecutive chunks. Therefore, we apply expression (6) over $\{S_1, S_2, \ldots, S_k, \ldots\}$, then we obtain a sequence of positive reals $\{D(S_1, S_2), D(S_2, S_3) \ldots, D(S_k, S_{k+1}) \ldots\}$. Intuitively, when $D(S_k, S_{k+1})$ value is *large* (much larger than previous ones), then the proposed metric has identified a concept drift. On the other hand, low values of $D(S_k, S_{k+1})$ represents unchanged distribution of the data. However, important changes between two consecutive chunks does not necessarily mean that has occurred a concept drift. For instance, it can be just due to the occurrence of outliers in one chunk and not in the next one. Even, it can also occurs because the data stream is provided in small buffers (chunk size too small). As a consequence, we smooth out the sequence $\{D(S_1, S_2), D(S_2, S_3) \ldots, D(S_k, S_{k+1}) \ldots\}$ for mitigating the impact of outliers in sequential data applying a smoother [6]. In the experimental part, we analyzed two classic operations: moving average (with a time windows of 5 chunks) and a Locally Weighted Scatterplot Smoothing [6]. By definition, the smoothing mapping provides a curve with continuous first derivative (at least first derivative is continuous). We created a really fast and simple decision rule based on a threshold over the first derivative of the smoothed curve. Let's denote the points of the first derivative of the smoothed curve by $\{l_1, l_2, \ldots\}$. We identify a *critical* point location (i.e., timestamps when a concept drift occurs)

when a gradient point

$$l_i \notin [\bar{l} - \alpha\sigma(l), \bar{l} + \alpha\sigma(l)], \tag{7}$$

where $\bar{l}$ denotes the mean of the sequence, $\sigma(l)$ is the standard deviation, and α is a real-value control parameter.

3.4 Main parameters

- **Control of KL-divergence numerical stability.** In cases of having pmf values equal to zero, then is necessary to make a smooth correction (using an arbitrarily small $\varepsilon > 0$ parameter) as described in section 2.3.

- **Estimation of probability mass function.** The estimation is based on the partition of the input space. The proposed partition is done using a regular grid with J bins. This parameter may impact on the probability mass estimation and in the KL-divergence numerical stability. If J is large, then it may exist bins without any data instance. As a consequence, J large can increase the instability of the KL divergence computation. In addition, it impacts in the computational cost of the pmf estimation. To estimate a pmf over a regular grid with many bins is much more costly than to split the input space in few sub-regions. After carrying out multiple runs, on our experiments we fixed the number of bins to $5 \times p$, where p is the dimension of $\mathbb{U}$.

- **Control of the smoothing function.** Smoothing methods have associated several parameters such as the number of iterations, smoothing, relaxation and so on [6].

- **Tuning of the decision rule.** The α-threshold value presented in the decision rule may produce impact in the result performance (accuracy, confusion matrix, etc.). A lower value of α makes a larger interval, then it is possible to fall in the error of detecting false positive drifts. On the other, a large value of α means that the method may increase the false negative error. That decision rule is commonly used for outlier detection and artifact removal over signals [3, 2]. Note that the mean and standard deviation can be done over a segment of the sequence (for instance last n points), it does not need to be over the full sequence. In addition, in practice the parameter can be dynamically corrected according to the performance of the matching matrix (in cases when the information about changes on the distribution arrives at certain moment).

3.5 Algorithm overview

We present a pseudo-code of the KLD detector in Algorithm 1. The presented version is online, it means given a new chunk performs an evaluation (it checks if there is a change in the distribution). Although, it is enough few intuitive changes to transform the online version to an offline version. Both lines 4 and 6 apply expression (3) for estimating the probability mass function. In this estimation is necessary to define a regular grid of the input space. This grid can be specifically created for the data on the consecutive two chunks (S_i and S_{i+1}) (partitioning a region of $\mathbb{U}$) or it can be made just once for the whole data (partitioning $\mathbb{U}$). Numerical computation of KL divergence is more stable when a regular grid is created specifically for each pair of consecutive chunks. In line 10 is applied the smoothing mapping. In our experiments, we didn't evaluate several smoothing methods, we applied moving averages and LOWESS. Further research can be done for analysing the impact of the smoothing mapping computation. Figure 1 presents a high-level pipeline of algorithm 1. It shows the sequence of the main operations. Each arrow represents a mathematical expression. For simplification, some operations such as the domain partition and smoothness of the dissimilarity metric are not presented in the diagram.

4 Experimental study

4.1 Goal of the experiment

We designed a set of experiments to answer if we can identify concept drift (critical points) in incrementally drifting data stream only based on topological characteristics of the raw data?

Algorithm 1: KL-divergence-based concept drift detector (KLD).

```
// Main global parameters:  number of bins (J), α (threshold in decision rule)
// Initial sets
```
1 $D = \emptyset$ (similarity distances) ;
2 $C = \emptyset$ (critical points - concept drift appearances) ;
```
// Initial chunk
```
3 Given chunk S_0;
4 Apply (3) for data in S_0 ;
5 **while** *(it arrives a new chunk (S_i))* **do**
```
       // Estimation of pmf
```
6   Apply (3) for data in S_i ;
```
       // KL divergence for each bin
```
7   Apply (4) for data in S_{i-1} and S_i;
```
       // Metric aggrregation
```
8   Apply (6) to compute $D(S_{i-1}, S_i)$;
9      $D \leftarrow D \cup D(S_{i-1}, S_i)$;
10   To apply a smoothing mapping over sequence D;
11   To compute first derivative (sequence l);
```
       // Decision rule
```
12   **if** *(expression (7) is satisfied)* **then**
13         $C \leftarrow C \cup l_i$;
14   **end if**
15 **end while**
16 Return the set of timestamps of detected drifts (critical points) C;
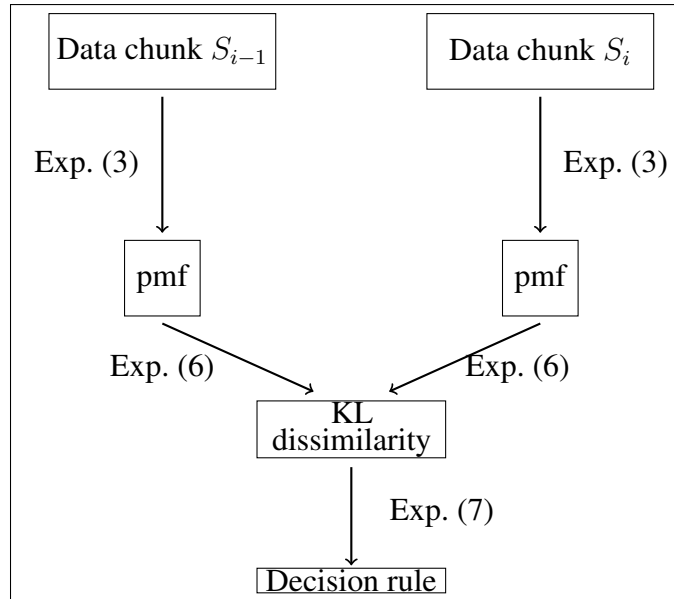

Figure 1: High-level flowchart of KLD algorithm 1.

4.2 Experimental setup

Choice of benchmark data streams. As there is a lack of a large diverse collection of real data streams where concept drift appearances are additionally marked. Therefore, to properly evaluate the proposed methods, a set of simulated data was generated using *Stream-learn* library [16]. Each data stream has 10000 chunks with 250 instances. We worked over two types of binary datasets with four and six features. A total of eight simulated data were evaluated. Drifts were generated using a sigmoid function that determines how sudden the change of drift concept it is. We used a concepts

Table 1: Applied parameters of the simulated data using *Stream-learn* library.

Description	Values
Seeds	$\{1410, 6543, 2345, 9876, 3946\}$
Number of concept drifts	20
Input features	$\{4, 6\}$
Number of classes	2
Number of chunks	10000
Chunk size	250
Concept sigmoid spacing	99
Class flip	0.01

sigmoid parameter of 99 (maximum value is 999 - it is the value for simulating a sudden drift). Each dataset has 20 induced concepts drifts. We present in Table 1 the used parameters for generating the data. More details about each of the parameters are specified in [16].

Experimental protocol and implementation. The implementation of the proposed method and the experimental environment are done using the *Python* v3.9 programming language and a few libraries: *NumPy* v1.19.5, *tsmoothie* v1.0.4 and *stream-learn* v0.8.16. We used the implementation of CART Decision Tree (DT) and Gaussian Naïve Bayes Classifier form sk-learn v1.0.2, and learning protocol was based on the *Test-Than-Train* [15] evaluation protocol.
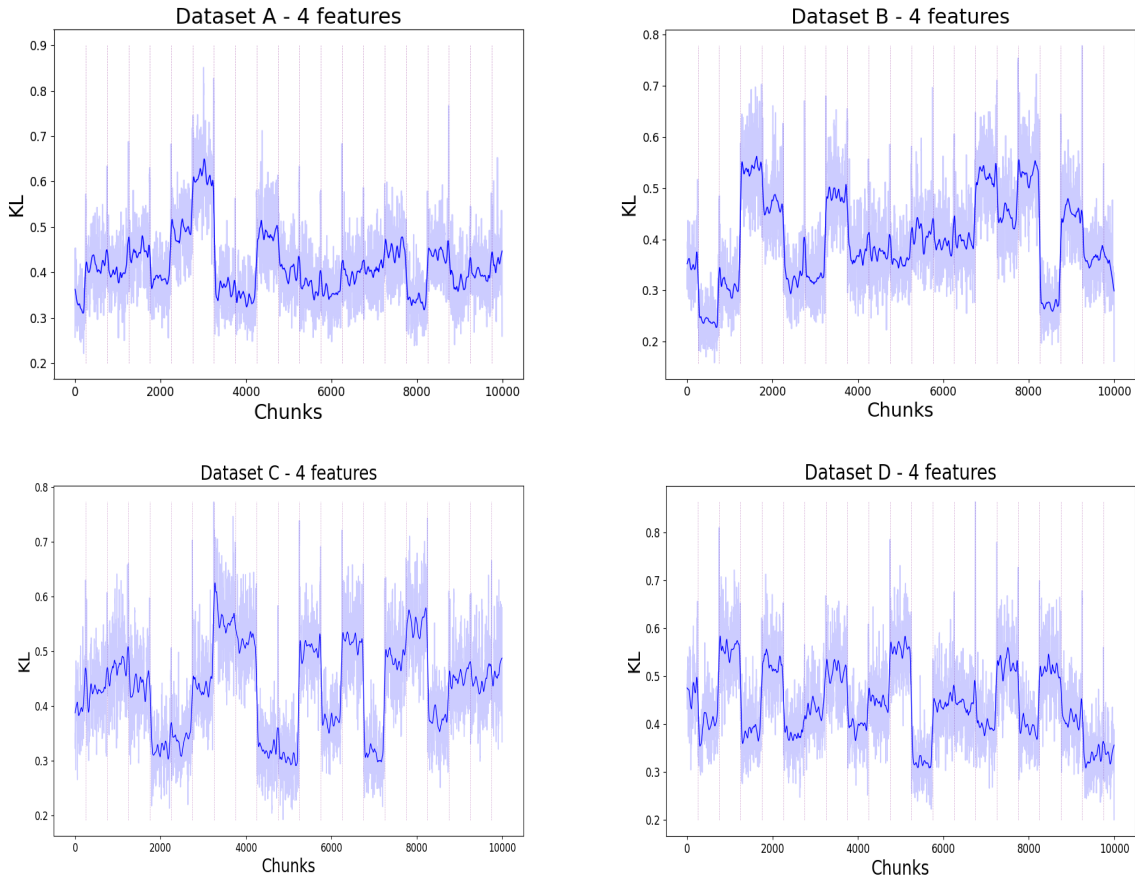

Figure 2: KL divergence values for the problems with four features. Dotted vertical lines show where was simulated the data distribution changes.
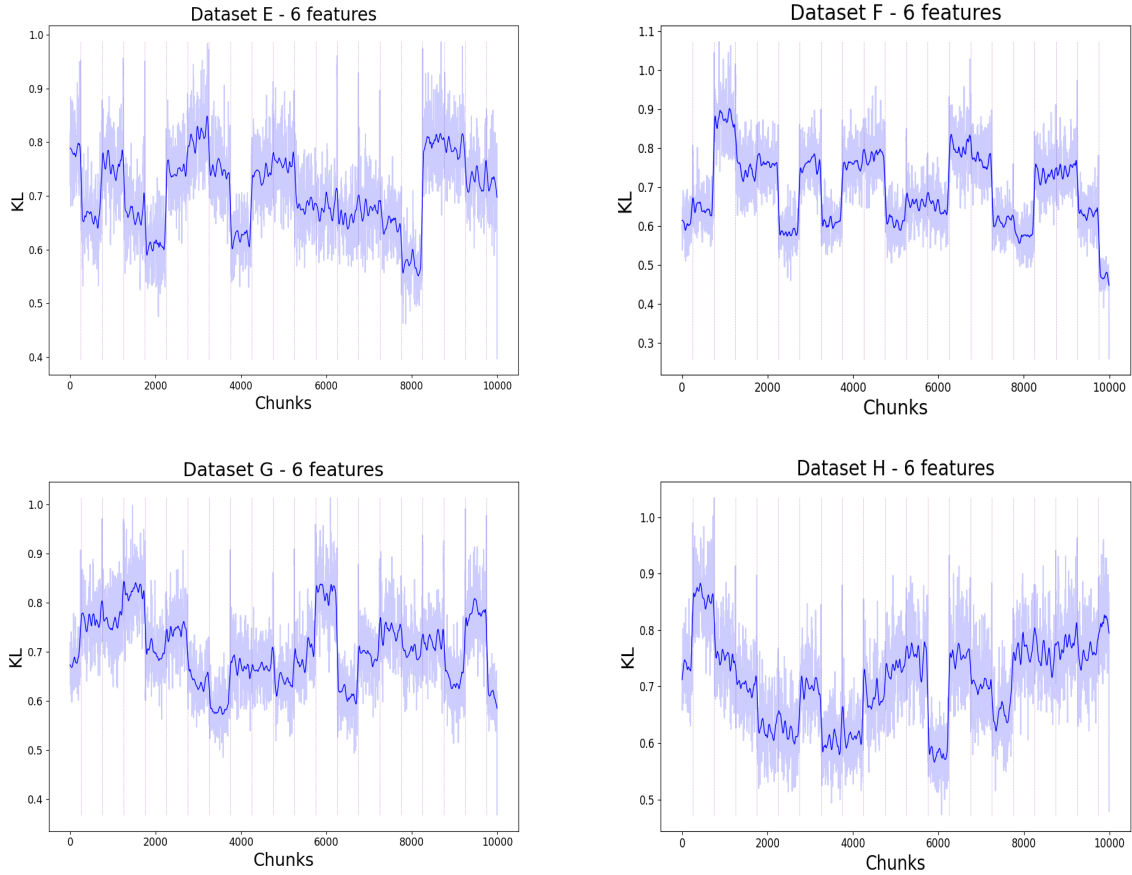
Figure 3: KL divergence values for the problems with six features. Dotted vertical lines show where was simulated the data distribution changes.

4.3 Results

Figure 2 presents four graphics associated to the datasets with four features. A similar set of graphics is presented in Figure 3, in this case the graphics were produced over datasets with six features. In each graphic, the light blue dots show the similarity metric defined in expression (5). The data points are normalized in [0,1]. Over the large cloud of light blue dots, it is shown in dark blue the smoothed curve. This curve is produced using the smoothing mapping (line 10 of Algorithm 1). Furthermore, the graphics have dotted vertical lines that show where is produced the induced drift. From the figures, it is possible to see the variability of the data, and also to see how the method *properly* reacts to the induced drifts. Figure 6 presents a visual comparison of the proposed approach with SOTA techniques. The left group of graphics shows the results for the datasets with four features, and on the right side there are graphics with the results produced over datasets with six features. It is beyond the scope of this paper to make an exhaustive comparison of the performance of our approach with change-point detection techniques based in classifiers. KLD only applies similarity analysis over the raw data. However, we also present an example of visual comparison of our method for monitoring changes and two well-known classifiers (Decision Tree (DT) and Gaussian Naive Bayes (GNB)). The two graphics in Figure 4 are curves with the first derivative of the smoothed studied metric. The curves are centered at origin. We can identify extreme values using the decision rule for detecting the critical points. It is worth mentioning that the DT works much better than GNB, it is visible in curves presented in Figure 6. In this example, it is also possible to see that KLD can identify change-point locations in places where DT classifier has low accuracy (e.g. in dataset D chunk after 4000, in dataset G chunk before 4000). Figure 5 presents a sensitive analysis of α-threshold. The figure has two graphics, one
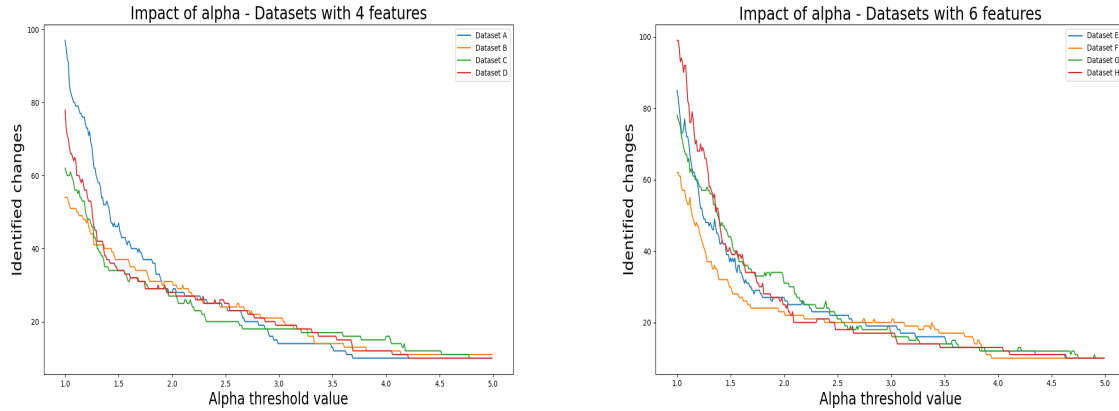
Figure 4: Two examples of the gradient of the smoothed KL divergence curve.

Figure 5: Sensitive analysis of the control parameter α. The y-axes show the number of critical points (when the gradient curve surpasses an upper-bound defined by the decision rule).

corresponds to the datasets with four features and another one to the datasets with six features. The number of changes is presented in the y-axes, we consider a change as the critical point when the gradient curve crosses an upper-bound threshold (defined by the decision rule). A larger value of α implies that the method increases the number of identified critical point locations. Note that, the simulated data has 20 induced drifts. An α value lower than 1.2 produces several false positives, and a large value of α makes that the method is not able to detect too many critical points. The behavior of the impact of α is similar through the eight analyzed problems. Figure 7 also depicts the behavior of the α-threshold. We mark the segments when the gradient curve reaches the critical interval (interval defined by the decision rule). The segments start with a blue dot and end with a red dot, i.e., blue dots represent the timestamps when the gradient curve reaches the critical area, and red dots show the timestamps when the curve escapes from the critical area. For illustrative purposes, we selected two datasets with four features (A and D) and two datasets with six features (E and H). We show the results of $\alpha \in \{1.5, 1.75, 2, 2.25, 2.5\}$. Each scatterplot has in the x-axes the evolution of chunks. As in the previous graphics, vertical lines indicate the timestamps when an incremental concept drift was induced by the simulator. According to the graphics seems that an optimal value of α is in the interval $[1.5, 1.75]$. The proposed method has a high ability to detect the change-points locations, and for α values larger than 1.5 the method almost does not predict false positives.
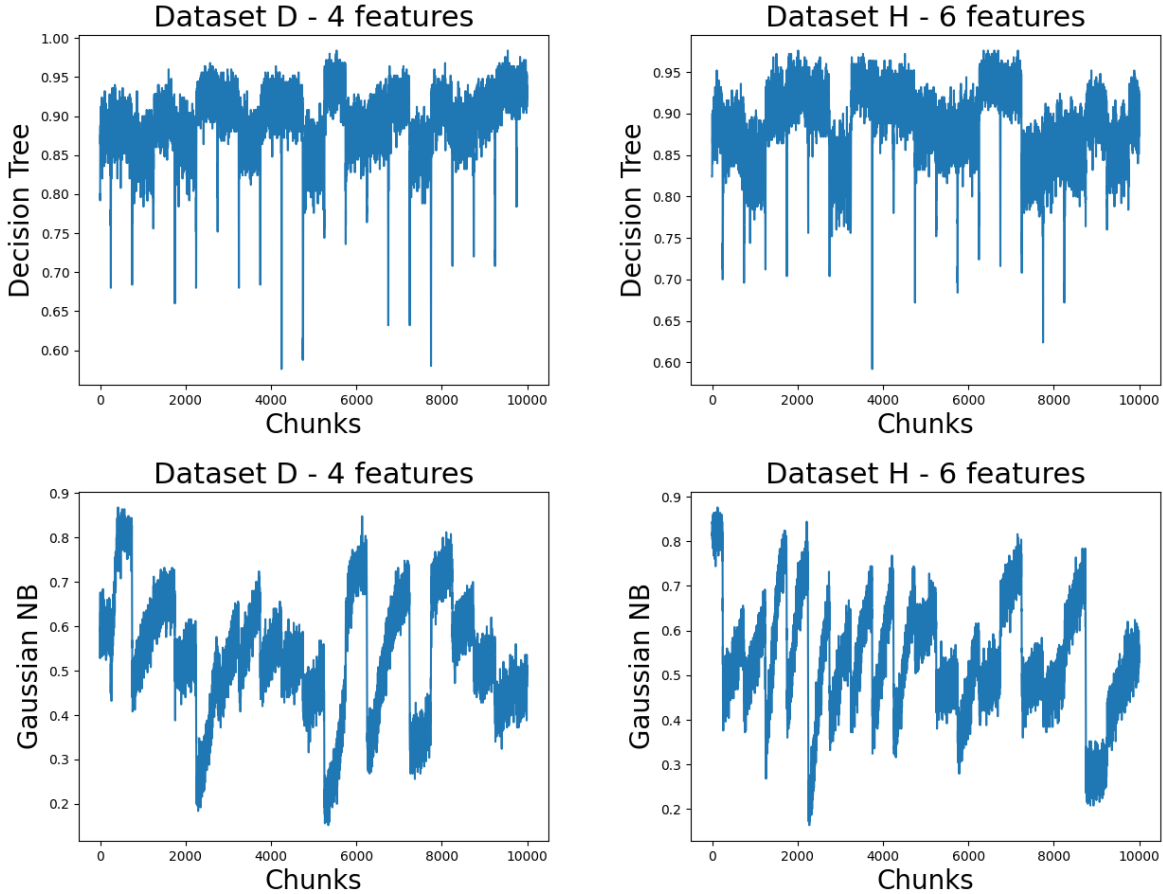
Figure 6: Example of comparison with SOTA techniques. Gradient of the smoothed KL metric curve compared with other tracking techniques (Decision Tree and Gaussian Gaussian Naive Bayes) over datasets with four and six features.

4.4 Lessons learned

Achieved results answered our initial experimental goal. We have introduced a fast approach for tracking changes in a data stream, which is based on entropy and similarity analysis over the raw data. An advantage of the proposal is that the method is universal, in the sense that, it does not assume any type of specific data distribution. The technique has a parameter (α) that is easy to adjust, and directly impacts the matching matrix. We empirically found a short interval where the α value can be tuned. Empirical results show that, independently of the input domain, it is suitable to apply the method with α value in $[1.5, 2]$.

5 Conclusions and future work

We presented a fast and effective method for tracking changes in probability distributions that may occur during continuous learning. It should be noted that the presented evaluation results of the proposed method are preliminary. However, they confirm its usefulness, which encourages us to continue the work in the future. The most important directions for further work include: (i) Evaluate the effectiveness over other data domains. Current approach has been evaluated only on binary output variables. (ii) Applying it to real-world continual learning tasks with different type of drifts. (iii) Analysis of its suitability for processing visual information, i.e., image streams with so-called syntactic drift. (iv) Consider other similarity measures such as Jensen-Shannon entropy or Wasserstein distance.
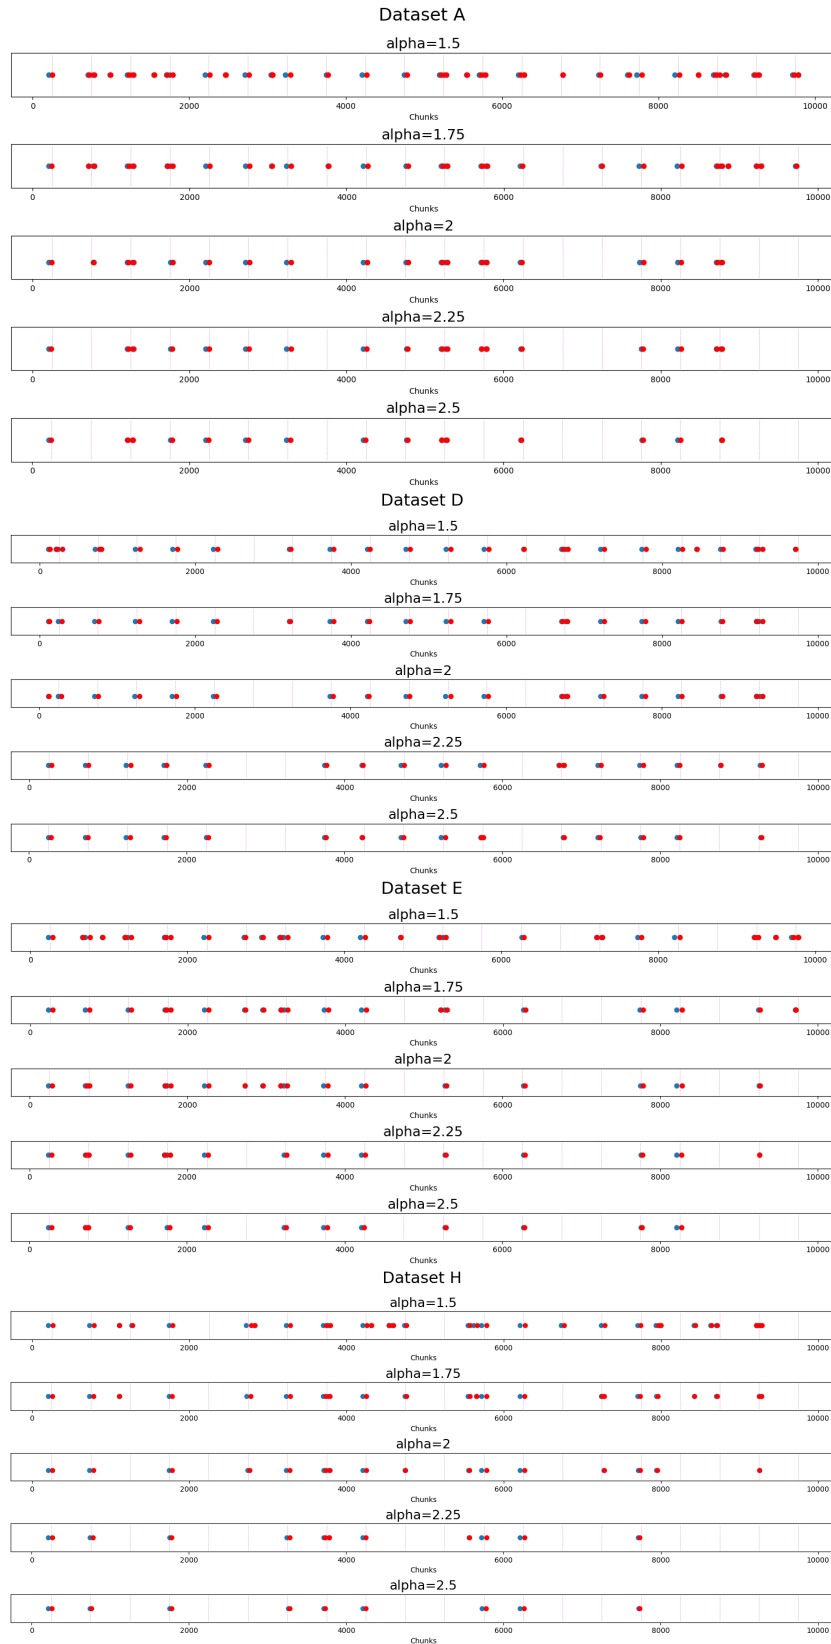
Figure 7: Sensitive analysis of the control parameter α. Horizontal axes show the evolution of chunks. Each scatter-plot presents the results obtained with a specific α value. Vertical lines indicate the timestamps when a concept drift was induced by the simulator. Blue dots represent the timestamps when the gradient curve reaches the upper-bound of the decision rule, and red dots show the moment when the curve goes outside the critical area.

12

References

[1] Manuel Baena-Garcıa, José del Campo-Ávila, Raúl Fidalgo, Albert Bifet, R Gavalda, and R Morales-Bueno. Early drift detection method. In *Fourth international workshop on knowledge discovery from data streams*, volume 6, pages 77–86, 2006.

[2] Sebastián Basterrech, Pavel Bobrov, Alexander Frolov, and Dušan Husek. Nature-inspired Algorithms for Selecting EEG Sources for Motor Imagery Based BCI. In *Artificial Intelligence and Soft Computing*, volume 9120 of *Lecture Notes in Computer Science*, pages 79–90. Springer International Publishing, 2015.

[3] Sebastián Basterrech and Pavel Krömer. A nature-inspired biomarker for mental concentration using a single-channel eeg. *Neural Computing and Applications*, 2019.

[4] Albert Bifet and Ricard Gavaldà. Kalman filters and adaptive windows for learning in data streams. In Ljupčo Todorovski, Nada Lavrač, and Klaus P. Jantke, editors, *Discovery Science*, pages 29–40, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[5] Isvani Inocencio Frias Blanco, Jose del Campo-Avila, Gonzalo Ramos-Jimenez, Rafael Morales Bueno, Agustin Alejandro Ortiz Diaz, and Yaile Caballero Mota. Online and non-parametric drift detection methods based on hoeffding's bounds. *IEEE Trans. Knowl. Data Eng.*, 27(3):810–823, 2015.

[6] Andreas Buja, Trevorh Hastie, and Robert Tibshirani. Linear smoothers and additive models. *The Annals of Statistics*, 17(2):453–555, 1989.

[7] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

[8] Imre Csiszar. The method of types [information theory]. *Information Theory, IEEE Transactions on*, 44:2505 – 2523, 11 1998.

[9] Tamraparni Dasu, Shankar Krishnan, and Suresh Venkatasubramanian. An information-theoretic approach to detecting changes in multidimensional data streams. *Interfaces*, pages 1–24, 2006.

[10] Lei Du, Qinbao Song, Lei Zhu, and Xiaoyan Zhu. A selective detector ensemble for concept drift detection. *The Computer Journal*, page bxu050, 2014.

[11] Joao Gama, Pedro Medas, Gladys Castillo, and Pedro Rodrigues. Learning with drift detection. In *In SBIA Brazilian Symposium on Artificial Intelligence*, pages 286–295. Springer Verlag, 2004.

[12] P. M Gonzalvez, Silas de Carvalho Santos, R. Barros, and D. Vieira. A comparative study on concept drift detectors. *Expert Systems with Applications*, 41(18):8144–8156, 2014.

[13] Frederik Gustafsson. *Adaptive Filtering and Change Detection*. Wiley, October 2000.

[14] T. Ryan Hoens, Robi Polikar, and Nitesh V. Chawla. Learning from streaming data with concept drift and imbalance: an overview. *Prog. Artif. Intell.*, 1(1):89–101, 2012.

[15] Bartosz Krawczyk, Leandro L. Minku, Jo ao Gama, Jerzy Stefanowski, and Michał Woźniak. Ensemble learning for data stream analysis: A survey. *Information Fusion*, 37:132 – 156, 2017.

[16] P. Ksieniewicz and P. Zyblewski. stream-learn — open-source python library for difficult data stream batch analysis. *Neurocomputing*, January 2022.

[17] Andrzej Lapinski, Bartosz Krawczyk, Pawel Ksieniewicz, and Michal Wozniak. An Empirical Insight Into Concept Drift Detectors Ensemble Strategies. pages 1–8, 07 2018.

[18] B. I. F. Maciel, S. G. T. C. Santos, and R. S. M. Barros. A lightweight concept drift detection ensemble. In *2015 IEEE 27th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 1061–1068, Nov 2015.

[19] Kyosuke Nishida and Koichiro Yamauchi. Detecting concept drift using statistical testing. In *Proceedings of the 10th International Conference on Discovery Science*, DS'07, pages 264–269, Berlin, Heidelberg, 2007. Springer-Verlag.

[20] Gustavo Oliveira, Leandro Minku, and Adriano Oliveira. Tackling virtual and real concept drifts: An adaptive gaussian mixture model, 2021.

[21] E. S. Page. Continuous Inspection Schemes. *Biometrika*, 41(1/2):100–115, 1954.

[22] Sarunas Raudys. *Statistical and Neural Classifiers: An Integrated Approach to Design*. Springer Publishing Company, Incorporated, 2014.

[23] Gordon J. Ross, Niall M. Adams, Dimitris K. Tasoulis, and David J. Hand. Exponentially weighted moving average charts for detecting concept drift. *Pattern Recognition Letters*, 33(2):191 – 198, 2012.

[24] Jeffrey C. Schlimmer and Richard H. Granger, Jr. Incremental learning from noisy data. *Mach. Learn.*, 1(3):317–354, March 1986.

[25] Ammar Shaker and Eyke Hüllermeier. Recovery analysis for adaptive learning from non-stationary data streams: Experimental design and case study. *Neurocomputing*, 150:250–264, 2015.

[26] Mingsheng Shang, Ye Yuan, Xin Luo, and Mengchu Zhou. An α-β-Divergence-Generalized Recommender for Highly Accurate Predictions of Missing User Preferences. *IEEE transactions on cybernetics*, PP, 2022.

[27] Piotr Sobolewski and Michal Wozniak. Concept drift detection and model selection with simulated recurrence and ensembles of statistical detectors. *Journal of Universal Computer Science*, 19(4):462–483, feb 2013.

[28] Gerhard Widmer and Miroslav Kubat. Effective learning in dynamic environments by explicit context tracking. In PavelB. Brazdil, editor, *Machine Learning: ECML-93*, volume 667 of *Lecture Notes in Computer Science*, pages 227–243. Springer Berlin Heidelberg, 1993.

[29] Gerhard Widmer and Miroslav Kubat. Learning in the presence of concept drift and hidden contexts. *Mach. Learn.*, 23(1):69–101, April 1996.

[30] Jie Zhang, Zhi Wei, Zhenyu Yan, MengChu Zhou, and Abhishek Pani. Online Change-Point Detection in Sparse Time Series With Application to Online Advertising. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(6):1141–1151, 2019.