Incremental and Adaptive Feature Exploration over Time Series Stream

Jingwei Zuo, Karine Zeitouni, Yehia Taher

► To cite this version:

Jingwei Zuo, Karine Zeitouni, Yehia Taher. Incremental and Adaptive Feature Exploration over Time Series Stream. IEEE International Conference on Big Data 2019, Dec 2019, Los Angeles, United States. 10.1109/BigData47090.2019.9005660 . hal-02391802

HAL Id: hal-02391802

<https://hal.science/hal-02391802>

Submitted on 3 Mar 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

INCREMENTAL AND ADAPTIVE FEATURE EXPLORATION OVER TIME SERIES STREAM

Jingwei Zuo

DAVID Lab, University of Versailles
Université Paris-Saclay
Versailles, France
jingwei.zuo@uvsq.fr

Karine Zeitouni

DAVID Lab, University of Versailles
Université Paris-Saclay
Versailles, France
karine.zeitouni@uvsq.fr

Yehia Taher

DAVID Lab, University of Versailles
Université Paris-Saclay
Versailles, France
yehia.taher@uvsq.fr

March 3, 2020

ABSTRACT

Over past years, various attempts have been made at analysing Time Series (TS) which has been raising great interest of Data Mining community due to its special data format and broad application scenarios. An important aspect in TS analysis is Time Series Classification (TSC), which has been applied in medical diagnosis, human activity recognition, industrial troubleshooting, etc. Typically, all TSC work trains a stable model from an off-line TS dataset, without considering potential Concept Drift in streaming context. Domains like healthcare look to enrich the database gradually with more medical cases, or in astronomy, with human's growing knowledge about the universe, the theoretical basis for labelling data will change. The techniques applied in a stable TS dataset are then not adaptable in such dynamic scenarios (i.e. streaming context). Classical data stream analysis is biased towards vector or row data, where each attribute is independent to train an adaptive learning model, but rarely considers Time Series as a stream instance. Processing such type of data requires combining techniques in both communities of Time Series (TS) and Data Streams. To this end, by adopting the concepts of *Shapelet* and *Matrix Profile*, we conduct the first attempt to extract the adaptive features from Time Series Stream based on the *Test-then-Train* strategy, which is applicable in both contexts: a) under stable concept, learning model will be updated incrementally; b) for data source with Concept Drift, previous concepts that do not represent the current stream behavior will be discarded from the model.

Keywords Time Series · Data Stream · Time Series Stream · Concept Drift · Spark · Distributed Computing

1 Introduction

Time Series (TS) is a sequence of real-valued data, which can be collected from various sources, such as ECG data in medicine, IoT data in smart cities, light-curves in astronomy, GPS or accelerometer data in activity recognition, etc. Time Series Classification (TSC) is intended to predict the label of a newly input TS instance by extracting the knowledge from collected data. Various TSC approaches have been proposed by researchers in recent years which are suitable for different contexts along with dissimilar TS features. One Nearest Neighbor (1-NN) classifier for whole series similarity measure is a typical baseline of TSC research, which is usually combined with various distance measures [1–4]. Instead of considering the global feature of entire series, summary statistic features (e.g., mean, deviation, slope, etc.) can be extracted from every sub-series to build diverse ensemble classifiers [5–7]. With the emergence of TS dimensionality

reduction techniques (e.g., PAA [8], SAX [9], etc.), TS instances can be represented by high-dimensional vectors so that various techniques [10] from classic data analysis can be adapted into TS context. For instance, in case that motifs or frequent patterns are what characterize a given class, the dictionary based approaches [11–13] borrowed from Text Mining and Information Retrieval community can be adopted. As for the scenario that the occurrence of specific sub-series determines a class, TS can then be represented by such shape-based features, namely Shapelet [14]. Various Shapelet-based approaches have been proposed to optimize both the accuracy [15, 16] and the efficiency [17, 18] of the classification. Another remarkable attempt [19–21] adopting ensemble approaches on several TS representations (e.g., Shapelet-based, similarity-based, interval-based etc.) shows a superior accuracy to one single representation classifiers, where TS features are from different representation domains, and can not be presented in a single form.

The optimization of TS feature extraction and model construction process allows us to strive for a low prediction error, and stay as close as possible to Time Series’ nature Concept [22], which refers to the target variable that the learning model is trying to predict. Most TSC approaches are biased towards learning from an off-line Time Series dataset, with the assumption that data instances are independently and identically distributed (i.i.d) within a particular concept, but rarely consider the streaming context, where a gradual change of the concept happens along with the input of TS stream, that is *Concept Drift*. For instance, the most accurate ensemble classifiers [19, 20] are not good options in streaming context due to their complex architecture. Lazy classifiers on Time Series such as Nearest Neighbor (1-NN) [23] and dictionary based approaches [11] are applicable for streaming context. However, every input instance will be considered to adjust the inner concept, which requires potentially a large buffer space and will bring a huge computation cost. Recent Deep Neural Network (DNN) approaches [24–27] on TSC are capable of tuning the model incrementally, but stay always in an awkward position for the lack of explainability, which is required by domains like healthcare where questions of accountability and transparency are particularly important.

Shapelet, as a shape-based feature in TS, which is widely adopted by the community for its reliability and interpretability, provides a possibility to fulfil the aforementioned requirements. With an advanced work in [28], the explainability of Shapelet Extraction process is ensured even to non-expert, which offers us an option to further explore the streaming context conserving the current goodness.

To fill the gap between Time Series Classification and data streams processing, in this paper, we propose a TS stream learning algorithm, where TS features and models can be updated with consideration of Concept Drift. The incremental version of previous Shapelet work [28] under Spark framework allows us to further explore the *Test-then-Train* strategy, to evaluate the learning model constantly on newly input instance, then update the model regarding to the evaluation result. The cached information under old concept will be eliminated gradually by an elastic caching mechanism, which deals with the challenge of infinite streaming instances.

The main contributions of this paper are to achieve the following goals:

1. **Scalability:** Our algorithm conserves the scalability of Shapelet Extraction [28] in streaming context, which is always parallelizable in a remote Spark cluster with a minimum communication cost between distributed nodes.
2. **Shapelet Evaluation:** We propose a novel strategy to evaluate Shapelet, which shows the first attempt of transferring the techniques in Time Series community to Data Stream community.
3. **Test-then-Train:** The novel strategy, not only accelerates the incremental Shapelet extraction in stable-concept context, but also helps with detecting Concept Drift in streaming context by *Test-then-Train* technique.
4. **Explainability:** The algorithm shows not only interpretability of extracted features (i.e., Shapelets), but also a strong explainability of Shapelet Evolution in dynamic source context.
5. **Traceability:** The system allows to trace and explain the learning model at different time ticks, which gives us a possibility to supervise the system and back up the historical features.

The rest of this paper is organized as follows. In Section 2, we review the background and state-of-the-art approaches which are useful for our research problems. Then, Section 3 shows our system in both TS stream contexts of stable concept and drifting concept. Section 4 shows an empirical evaluation of our method on real-life datasets. Finally, we give our conclusions and perspectives for future work in Section 5.

2 Background

2.1 Definitions and Notations

We start with defining the notions used in the paper:

Definition 1: A Time Series T is a sequence of real-valued numbers $T=(t_1, t_2, \dots, t_i, \dots, t_n)$, where n is the length of T .

Definition 2: Time Series Stream S_{TS} is a continuous input data stream where each instance is a Time Series: $S_{TS}=(T_1, T_2, \dots, T_N, \dots)$. Notice that N increases with each new time-tick.

Definition 3: Time Series Chunk $C_{t,w}$ is a Time Series micro-batch at time-tick t with window size w in S_{TS} : $C_{t,w}=(T_{t-w+1}, T_{t-w+2}, \dots, T_t)$.

Definition 4: Cached Dataset D_t is a set of time series T_i , and class label c_i , collected after time tick t . Formally, $D_t=\langle T_t, c_{j_t} \rangle, \langle T_{t+1}, c_{j_{t+1}} \rangle, \dots, \langle T_N, c_{j_N} \rangle$, where N is the time-tick of the most recent input instance. $C = c_1, c_2, \dots, c_{|C|}$ is a collection of class labels, where $|C|$ denotes the number of labels.

Definition 5: A subsequence $T_{i,m}$ of Time Series T is a continuous subset of values from T of length m starting from index i : $T_{i,m} = (t_i, t_{i+1}, \dots, t_{i+m-1})$, where $i \in [0, n - m + 1]$.

Definition 6: Shapelet $\hat{s}$ is a time series subsequence which is particularly representative of a class. As such, it shows a shape which can distinguish one class from the others.

Definition 7: Euclidean Distance(ED) between two time series $T_{x,m}, T_{y,m}$ is expressed as follows:

$$ED_{x,y} = \sqrt{\sum_{i=1}^m (t_{x,i} - t_{y,i})^2} \quad (1)$$

Definition 8: Distance Profile DP_i is a vector which stores the Euclidean Distance between a given subsequence/query $T_{i,m}$ in source T and every subsequences $T'_{j,m}$ of target T' . Formally, $DP_i^m = (DP_{i,1}^m, \dots, DP_{i,j}^m, \dots, T_{i,n'-m+1}^m)$, where $DP_{i,j}^m = \text{dist}(T_{i,m}, T'_{j,m}), \forall j \in [0, n' - m + 1]$, n' is the length of T' .

Each element in DP_i is calculated by Euclidean Distance between z-normalized subsequences [28]. From **Fig. 1(a)**, we can visually obtain the position of Query's Nearest Neighbor (1NN) in T' from the lowest point in DP_i . Authors in [29] propose MASS which is considered as the fastest exact distance measure between two Time Series. MASS computes Distance Profile based on Fast Fourier Transform (FFT), which requires only $\mathcal{O}(n \log n)$ time and is independent of query's length, instead of $\mathcal{O}(nm^2)$ [14] by classical sliding window measure.

Figure 1: (a) Distance Profile between Query $T_{i,m}$ and target time series T' , where n' is the length of T' . Obviously, $DP_{i,j}$ can be considered as a meta TS annotating target T' ; (b) Matrix Profile between Source time series T and Target time series T' , where n is the length of T . Intuitively, MP_i shares the same offset as source T

Definition 9: Matrix Profile MP is a vector of distance between every subsequence $T_{i,m}$ in source T and its nearest neighbor $T'_{j,m}$ in target T' . Formally, $MP^m = (MP_1^m, \dots, MP_i^m, \dots, MP_{n-m+1}^m)$, where $MP_i^m = \min(DP_i^m)$, $i \in [0, n - m + 1]$, n is the length of T .

Unlike the distance profile, the matrix profile is a meta TS annotating the source TS. As shown in **Fig. 1(b)**, the lowest point in MP show the position of query which has the most similar matching in target TS.

Definition 10: *Representative Profile* RP_T^C is a vector of representative power of subsequences in T for class C : $RP_T^C = (RP_{T_1}^C, \dots, RP_{T_i}^C, \dots, RP_{T_{n-m+1}}^C)$

The representative power of subsequence T_i in class C is defined as:

$$RP_{T_i}^C = \text{avg}(MP_{T_i, T'}) \quad (2)$$

where $T' \in D_t^C$. Intuitively, $RP_{T_i}^C$ is a normalized distance between T_i and global TS instance cluster of class C , it represents the relevance between the subsequence T_i (i.e., the candidate Shapelet) and the class. As shown in **Fig. 2** (b)(d), a threshold can be set to show the starting index area in T , where the subsequences are representative for class C .

Definition 11: *Discriminative Profile* $DiscmP_T$ is a vector of discriminative power of subsequences in T :

$$DiscmP_T = RP_T^{NonC} - RP_T^C \quad (3)$$

The discriminative power of T_i in dataset shows the difference of representative power of candidate Shapelet from its own class to the others (*OVA, one-vs-all*), which follows similar heuristics in [30], where authors proved *OVA* strategy in Shapelet quality assessment performs better than traditional Decision Tree approach [14] in both accuracy and efficiency metrics. Intuitively, Discriminative Profile can give a global view of the important patterns' positions over a Time Series. As shown in **Fig. 2** (f)(g), the highest point in the profile shows the position of the sub-series in T , which has the biggest skewing of relevance between class C and other classes. Through setting a power threshold, the discriminative pattern, that is the candidate Shapelet, can be visually identified in T by their discriminative power.

Figure 2: (a) Time Series T with class C . (b,c) Matrix Profile set for T with TS instances in different classes. (d,e) Representative Profile of T in different classes. A threshold can determine the representative area in T . (f) Discriminative Profile of T in dataset. The highest point in $DiscmP_T^C$ identifies the most discriminative pattern's in T .

2.2 A Brief Review of the SMAP algorithm for TSC

Matrix profile [29] provides a meta-data which facilitates the representation of a complex correlation between two time series. Authors in [28] propose *SMAP*, a scalable algorithm for Shapelet Extraction on Matrix Profile, which ensures meanwhile the parallelism in Spark cluster, and the explainability of extraction process to a non-expert.

As shown in *Algorithm 1*, considering time series as the smallest processing unit between Spark nodes, *SMAP* firstly broadcasts the dataset to distributed nodes in order to reduce the communication cost from repetitive access of common data. Then, each cluster partition shares the computing tasks for a set of TS, and extracts the most discriminative sub-series of various length in each processing unit. Each extracted sub-series can be considered as a candidate Shapelet, which is assigned a distance threshold defined by its representative power in its own class. The threshold can determine the inclusion between the candidate Shapelet and a TS. A strategy to check if T contains a candidate Shapelet $\hat{s}$ can be defined as the following:

$$Inclusion(T, \hat{s}) = \begin{cases} true, & \text{if } dist(T, \hat{s}) \leq \hat{s}.dist_{Thresh} \\ false, & \text{otherwise} \end{cases} \quad (4)$$

Each TS unit is assigned an unique Hash ID to reduce the volume of transferred data between nodes. The TS ID, as well as the discriminative power and threshold distance of its contained candidate Shapelets, will be output as the computing results of the partition. Finally, a single aggregation process between nodes is required to obtain the Shapelet result of different classes. The whole extraction process can be visualized with a strong explainability, and generates high interpretable results.

Algorithm 1: SMAP (Shapelet extraction on MAtrix Profile on Spark)

Input: Dataset D , classSet $\hat{C}$, k

Output: $\hat{S}$

```

1  $l_{min} \leftarrow 0.1 * getMinLen(D), l_{max} \leftarrow 0.5 * getMinLen(D),$ 
2  $DiscmP \leftarrow [], dist_{Thresh} \leftarrow [], \hat{S} \leftarrow \emptyset$ 
3  $D.broadcast();$  //each TS has an unique ID
4 MapPartition (Set of  $\langle ID, T \rangle : T_{set}$ )
5   for  $\langle ID, T \rangle \in T_{set}$  do
6     for  $m \leftarrow l_{min}$  to  $l_{max}$  do
7        $DiscmP[m], dist_{Thresh}[m] \leftarrow computeDiscmP(T, D, m)$ 
8        $DiscmP[m] \leftarrow DiscmP[m] * \sqrt{1/l}$ 
9      $DiscmP, dist_{Thresh} \leftarrow pruning(DiscmP, dist_{Thresh})$ 
10    emit( $ID, DiscmP, dist_{Thresh}$ )
11 MapAggregation ( $class, (ID, DiscmP, dist_{Thresh})$ )
12   for  $c \in \hat{C}$  do
13      $\hat{S}' \leftarrow getTopk(DiscmP[c], dist_{Thresh}[c], k)$ 
14      $\hat{S} \leftarrow \hat{S} \cup \hat{S}'$ 
15 return  $\hat{S}$ 

```

To sum up this section, Discriminative Profile provides a possibility to extract the interpretable patterns in an explainable manner. The adoption of MASS essentially accelerates the extraction process compared to using pruning techniques based on brute force approach [14]. SMAP provides a parallel processing mechanism to conduct the extraction in a minimum communication cost on Spark cluster. In the next section, we will show an advanced algorithm which, when applied on Spark cluster, is capable of updating Shapelet results by adopting an incremental model in dynamic source context.

3 Algorithm and System Structure

In this section, we start by studying the incrementality of SMAP, which is a necessary condition for learning in streaming context. Then we propose the evaluation strategies to accelerate incremental learning process and adapting it to streaming context considering Concept Drift.

3.1 Incremental SMAP (ISMAP)

Typically, a non-incremental algorithm requires to re-pass the existing dataset and conduct a large amount of redundant computations. In Algorithm 2, we show Incremental Shapelet extraction on Matrix Profile on Spark (ISMAP), which avoids essentially the repetitive computations on existing dataset. As in Spark environment, the communication cost between distributed nodes is a key factor of system's efficiency. The computing task in each Spark partition should be relatively independent without frequent exchange of intermediate results with other partitions. In light of this, we need to make use of the parallel mechanism to well manage the allocation of computing tasks.

As shown in Algorithm 2, we assume that each Spark partition keeps a set of Time Series with their Discriminative Profiles and corresponding Threshold Distance sets. The newly input Time Series T_N will be broadcast to each distributed node. Information in T_N should be extracted and merged to existing knowledge base, which can be carried out into two steps:

1. **Update existing Shapelets:** With newly input instance T_N , existing candidate Shapelets should update their representative power in each class, and discriminative power in current dataset.
2. **Evaluate new candidate Shapelets:** T_N will introduce new candidate Shapelets of various length, which should be evaluated and placed into Shapelet ranking list by their discriminative power.

Algorithm 2: ISMAP(Incremental Shapelet extraction on MATrix Profile on Spark)

Input: Partition $[ID, T, DiscmP, dist_{Thresh}]$, **New input** T_N , **classSet** $\hat{C}$, k
Output: $\hat{S}$

```

1  $l_{min} \leftarrow 0.1 * getMinLen(D)$ ,  $l_{max} \leftarrow 0.5 * getMinLen(D)$ ,
2  $DiscmP \leftarrow []$ ,  $dist_{Thresh} \leftarrow []$ ,  $\hat{S} \leftarrow \emptyset$ 
3  $\langle ID_N, T_N \rangle.broadcast()$ ;
4 MapPartition ( $[ID, T, DiscmP, dist_{Thresh}]$ )
5   /* 1. compute the Matrix Profile between  $T_N$  and all TS in dataset */
6   /* 2. update the current DiscmP of all TS in dataset */
7   /* 3. prepare  $MP_{T_N}$  elements to compute  $DiscmP_{T_N}$  */
8   for  $m \leftarrow l_{min}$  to  $l_{max}$  do
9      $MP_T[m] \leftarrow computeMP(T, T_N, m)$ 
10     $MP_{T_N}[m] \leftarrow computeMP(T_N, T, m)$ 
11     $DiscmP[m], dist_{Thresh}[m] \leftarrow updateDiscmP(DiscmP[m], dist_{Thresh}[m], MP_T[m])$ 
12     $DiscmP, dist_{Thresh} \leftarrow pruning(DiscmP, dist_{Thresh})$ 
13    emit( $ID, T, DiscmP, dist_{Thresh}, MP_{T_N}$ )
14 MapAggregation ( $*, (ID, T, DiscmP, dist_{Thresh}, MP_{T_N})$ )
15    $DiscmP_{T_N}, dist_{Thresh_{T_N}} \leftarrow computeDiscmP(collect(MP_{T_N}))$ 
16    $DiscmP_{T_N} \leftarrow DiscmP * \sqrt{1/l}$ 
17    $DiscmP_{T_N}, dist_{Thresh_{T_N}} \leftarrow pruning(DiscmP_{T_N}, dist_{Thresh_{T_N}})$ 
18   cache( $ID_{T_N}, DiscmP_{T_N}, dist_{Thresh_{T_N}}$ )
19 MapAggregation ( $class, (ID, DiscmP, dist_{Thresh})$ )
20   for  $c \in \hat{C}$  do
21      $\hat{S}' \leftarrow getTopk(DiscmP[c], dist_{Thresh}[c], k)$ 
22      $\hat{S} \leftarrow \hat{S} \cup \hat{S}'$ 
23 return  $\hat{S}$ 

```

Step (1) is shown in line 9,11, from the Formula 2 and 3, we can observe that the linearity of Discriminative Profile makes the fact that each existing TS only need one single Matrix Profile computation with T_N to update the candidate Shapelets. As for Step (2), the Discriminative Profile computing of T_N is shared on different Spark partitions, where Matrix Profiles with existing TS instances are computed in line 10, an aggregation process in line 14-18 extracts the discriminative patterns in T_N , which will be aggregated with existing candidate Shapelets and update the output results in line 19-22.

Like classical incremental algorithms, ISMAP takes all input instances into account, which means every input TS instance will be imported into the system to update the Shapelet, even if the computing imports no valuable information into the system, that is, the information contained in the instance is repetitive with that in the knowledge base. Evidently, we are capable of avoiding the redundant information's computation by adopting an interleaved *Test-then-Train* strategy [31] with an extra Shapelet evaluation process over input instances.

3.2 Shapelet Evaluation

The intuition behind the evaluation procedure is that once we have a bad evaluation result, we need to import the instance batch into Shapelet Extraction process, to update the output Shapelet result. As the evaluation time $\mathcal{O}(n - m + 1)$ for a TS instance is much less than that of extraction computing ($\mathcal{O}(Nn^3 \log n)$), then an evaluation module can improve the system's efficiency by preventing the computation of certain valueless instances. However, how to define that an instance is valueless stays a problem to resolve.

The classical Shapelet-based approach [14] supposes that a Time Series T can be classified by the inclusion of a class-specified Shapelet $\hat{s}$. (i.e. if $dist(T, \hat{s}) \leq \hat{s}.dist_{thresh}$, then $T.class = \hat{s}.class$). The threshold distance of Shapelet gives a split point to decide the TS-Shapelet inclusion. As shown in [32], various approaches (e.g. Information Gain (IG), Kruskal-Wallis (KW) and Mood's Median (MM)) can be applied for both Shapelet assessment and split point decision. Representative Profile and Discriminative Profile achieve the same effect with these techniques but in a more interpretable manner. Intuitively, we are capable of deciding whether to import a TS instance into Shapelet Extraction process by evaluating its prediction results on current learning model. The Loss Measure is intended to detect the shift between the learning model and the inner concept of data source. In the context of Shapelet, the distance between the learned Shapelets and input instance is able to represent the loss to some extent. Typically, the distance is

compared with Shapelets' threshold distance, which derives the *0-1 Loss Function*:

$$L(Y, h(T)) = \begin{cases} 0, & Y = h(T) \\ 1, & Y \neq h(T) \end{cases} \quad (5)$$

where

$$h(T) = \begin{cases} C, & \text{if } \text{dist}(T, \hat{s}) \leq \hat{s}.\text{dist}_{Thresh} \\ \text{non}C, & \text{otherwise} \end{cases}$$

However, by TS-Shapelet inclusion technique, two Time Series with similar distance to a Shapelet may obtain different classes. In addition, a good prediction result of input TS instance with current Shapelets doesn't mean that the instance contains no useful information for adjusting the learning model. The *0-1 Loss Function* analyzes the surface phenomenon of the prediction but ignored the deep information behind the arbitrary split point technique. A loss measured by a crisp *0-1 Loss Function* is then ill-adapted.

When $\text{dist}(T, \hat{s}) \leq \hat{s}.\text{dist}_{Thresh}$, the prediction result is relatively acceptable. The problem then becomes how to find a balance between time efficiency and TS information checking (i.e., the exhaustive information extraction). The distance between TS and Shapelets describes the shift between real and learned concept, a small distance leads to a reliable prediction result. As the distance measure is usually data-dependant, and the absolute distance value varies with datasets, then a normalized measure describing the shift scale is required. To this end, can we just convert the TS-Shapelet inclusion problem to the possibility that a TS contains the Shapelet?

As extracted Shapelets try to separate one class to others, TSs in different classes tend to be concentrated on the split point, which causes the main error in prediction. Then we assume that $\text{dist}(T, \hat{s})$ satisfies Gaussian distribution, as shown in (6), the loss can be smoothed by *Sigmoid* function by considering distance distribution. The split point of $\hat{s}$ defines the expectation σ of the distribution.

$$L(Y, h(T)) = \frac{1}{1 + e^{-(x-\sigma)}}, \quad \sigma = \hat{s}.\text{dist}_{Thresh} \quad (6)$$

$$x = \min(\text{dist}(T^C, \hat{s})), \quad \hat{s} \in \hat{S}^C$$

As shown in Fig. 3, the smaller the loss, the greater the possibility that T will contain the Shapelet. Intuitively, a loss threshold Δ can be set by user for *ISMAP* to control the extraction from input instance, and update incrementally the Shapelet to approach the real concept of data source. When Δ is set to 0.5, it has the same effect as *0-1 Loss Function*.

Figure 3: Loss measure of Time Series by Sigmoid Function and 0-1 Loss Function, Time Series in different classes are distributed around the split point of Shapelet

However, a stable concept does not hold in several real-life scenarios. For instance, with the soundness of the knowledge in a particular domain, the labeling of newly input instances may evolve gradually, leading to a concept drift. Therefore, the most recent training instances should contribute more than the oldest ones to the target prediction. Then the problem becomes the Concept Drift detection in a Time Series Stream by monitoring the loss function. Conserving the interpretability and explainability of the algorithm, *ISMAP* can be extended to the context of TS Stream by extracting adaptive features.

3.3 Adaptive feature extraction from Time Series Stream

As shown in **Fig. 4**, the system of extracting adaptive Shapelets from Time Series Stream is composed by Shapelet Extraction, Evaluation Bloc and Caching Mechanism. We take TS Chunk $C_{t,w}$ as minimum input unit which contains a number of continuous TS instances: $C_{t,w} = (T_{t-w+1}, T_{t-w+2}, \dots, T_t)$, where t is the time-tick, w is the window size. By adopting the *Test-then-Train* strategy, the main idea here is to evaluate continuously the shift between learned concept and real concept in data source (i.e., *Test*). Once a Concept Drift is detected, the input chunk will be imported into Shapelet Extraction bloc to update the learning model (i.e., *Train*). Both Shapelet initialization and updating process are parallelizable on Spark cluster, which makes use of RAM as caching unit to lower the I/O cost.

Figure 4: System Structure in TS Stream context with Concept Drift

3.3.1 Shapelet Extraction

The computing process follows the same methodology with *ISMAP*, which allows TS instances in the input chunk to be partitioned on various Spark nodes, the discriminative patterns in each partition will be extracted individually and merged between partitions by their ranking power. The Shapelet ranking list is then composed by power-updated existing Shapelets and newly imported candidates.

3.3.2 Learned Concept Evaluation

As aforementioned, the loss of a Shapelet on input instances can describe its shift to real concept of data source. With the same methodology, once a concept drift is detected in data stream, the analysis tends to be more complicated. The challenge here is to distinguish the measured loss from two aspects:

- **Incomplete Extraction:** As main constraint in Shapelet Extraction, insufficient training instances (i.e., under-fitting) will bring a relative high loss. More data will make the learning model approach more the inner concept.
- **Concept Drift:** The measured shift can only reflect the distance to a stable concept, a big shift will be observed using out-of-date learning model.

In light of these challenges, an advanced analysis on detecting the Concept Drift from measured loss is required. That is, not only to measure the loss from each TS Chunk, but also to propose a strategy to analyse the loss. Based on the loss definition in (6), we define the average loss for a TS chunk $C_{N,w}$:

$$L_C(N) = \frac{1}{w} \sum_{k=1}^w L(Y_{N-w+k}, h(T_{N-w+k})) \quad (7)$$

Concept Enrichment: As mentioned in 3.2, an user-defined loss threshold Δ can be set to decide whether to import the chunk into the system to enrich the concept. That is: $Import_{Chunk} = True$ if $L_C(N) \leq \Delta$.

Concept Drift detection: Page-Hinkley test (PH) [31] is a typical technique used for change detection in signal processing. It allows a loss tolerance for the signal. The sequential test on the variance which considers that normal operation corresponds to a certain variance and a drift being characterized by an increase in this variance. Here we define a cumulative difference between the observed loss and their mean up until the current time:

$$m_N = \sum_{t=0}^N (L_C(t) - L_{avg}(t) - \delta) \quad (8)$$

where $L_{avg}(t)$ is the average loss until the current time tick t , δ specifies the tolerable magnitude of changes. The minimum value of m_N is defined as $M_N = \min(m_t, t = 1 \dots N)$. PH test will measure the difference between M_N and m_N :

$$PH_N = m_N - M_N \quad (9)$$

Intuitively, the difference reflects the degree of Concept Drift, when it exceeds a user-specified threshold λ , then the Concept Drift is detected.

3.3.3 Caching Mechanism

As the discriminative power of a candidate Shapelet is based on its global distribution in dataset, the fact that TS instances should be cached in memory is then a necessary condition of Shapelet Extraction. This is the main difference compared to Concept Drift detection in classical data streams, where it's possible to have one single pass on input instance. Then the main challenge here is to bridge the gap between TS and data stream analysis, that is, to consider the nature of Shapelet in Time Series Stream (i.e., part of data instances need to be cached), and propose a Shapelet-based caching mechanism in streaming context, meanwhile the caching volume should not increase indefinitely along with the never-ending TS Stream.

As Concept Drift is the fact that the prediction targets at different time tick are different, the previous learned concept is inapplicable to current input data. Conversely speaking, the fresh extracted concept doesn't match previous prediction target. This fact opens a path to optimize proactively the data caching procedure in memory. Considered as a complement to the global system shown in Fig. 4, the caching mechanism is detailed in Fig. 5. Once a Concept Drift is detected in TS stream, the TS chunk will be cached into memory. As mentioned in *ISMAP*, newly input chunk will generate its Matrix Profile set, and update those of previously cached chunks, which eventually leads to the update of Shapelet list in the learned concept. When there is a state transition in Concept Drift detection, the extraction of a fresh concept is then finished which is applicable for streaming instances coming afterwards. The detection of this transition state triggers then a cache elimination procedure.

Figure 5: Elastic Caching Mechanism for streaming instance chunk in memory

The elimination procedure is based on the fact that the prediction target of an old TS Chunk is not compatible with the fresh learned concept. By evaluating the cached chunks chronologically, we aim at finding the transition border where historical chunk starts to match the fresh updated concept, which is a reverse process to the detection of cache elimination trigger. We assume that $C_{t,w}$ is the oldest chunk cached in memory, after a trigger is detected, the evaluation will be conducted from $C_{t,w}$ to more recent chunks using the fresh learned concept. If the prediction target in $C_{t,w}$ matches the fresh concept, that means $C_{t,w}$ is in the frame of the fresh concept, the chunks in later time ticks also contributes to the concept's tuning, which can be kept in memory. Otherwise, $C_{t,w}$ should be removed from cache to eliminate the negative effect to the fresh learned concept. The process will not stop until a transition border is detected. By this proactive mechanism, the system is capable of caching a stable volume of data in TS Stream context, and generating adaptive Shapelets in the frame of drifting concept.

4 Experiments and Results

All the programs are implemented under Python 3.6. The source code, as well as a demonstration video [33] can be found in our project page¹. The Shapelet Exploration process can be either conducted at local or on a remote Spark cluster. We provide also an 1-click cluster based on Docker, to facilitate the replay of the distributed test offline by the user.

4.1 Experimental design

The experiment is conducted by two steps: A). We test the incremental feature and reliability of ISMAP after adopting Shapelet Evaluation process in Test-then-Train strategy. We evaluate the improvement of Shapelet Extraction in both efficiency and accuracy on data source with stable concept; B). We check the reliability of adaptive Shapelet Extraction from TS stream with Concept Drift.

Datasets: UCR Archive [34] is the most complete TS collection in the community, where the datasets are collected from diverse domains, such as readings from Image Outlines, Sensor Readings, Motion captures, spectrograph, and so forth. Each domain matches to certain problem type, which can be best tackled by a specific approach. Authors in [35] studied in detail the approaches and their matching problem types, and pointed out that Shapelet-based method performs relatively better in readings of *Sensor data*, *ECG data*, *Border-converted images*, etc. To this end, we conduct our first incremental experiments on 14 shapelet-characterized datasets in UCR Archive, instead of testing all datasets under various problem types.

However, when we switch to streaming context where Concept Drift happens in TS flow, to the best of our knowledge, currently, the community doesn't collect the dataset which reflects a such phenomenon, due to the fact that the problem hasn't been studied before. Therefore, we generate synthetic datasets by manually adjusting the data source to comply the test scenario. Instead of generating data from scratch to comply Shapelet features, the synthetic data is based on the datasets *Trace* and *ECG5000* (see Table 1) from two domains, which have a high reliability in Shapelet-based approaches and eligible for simulating the *Concept Drift* scenario by changing the class within some chunks on different time ticks.

Data Augmentation: Public datasets, especially those commonly used with shapelets, have relatively small size if we compare them to typical dataset in data mining². A larger number of instances is required for Concept Drift assessment in the streaming TS context. Adding noise is a typical way for data augmentation, as *ISMAP* extracts a range of patterns from each Discriminative Profile and then merge them by their discriminative power, a range-based extraction [36] rather than a single top value is noise resistant. We augment the data volume by randomly putting noise in TS instances with a random duration. The augmentation degree is set to 10 times of original volume. We put 3 Concept Drifts equally distributed on time axis. The total labelled instances are sampled into 3 equal-sized subsets with different concepts.

Reproducibility and Parameters: The initial Shapelet Extraction algorithm is based on [28], where the Shapelet length $m \in [0.1n, 0.5n]$ with a step of $0.25m$, where n is TS length. The advanced *MASS* algorithm (*mass_v2*) [37] for similarity measure is re-implemented under Python3, where flat subsequences (i.e., those where all values are equal) are ignored by the algorithm, as such a subsequence is meaningless from the definition of Shapelet, and will produce an error during *MASS* computation due to the empty value of its standard deviation. The threshold in Discriminative Profile is replaced by a top-K selection in the profile, which follows the same value as final extracted Shapelet number k ($k = 10$ for each class). As for the loss threshold for Shapelet Evaluation, $\Delta \in \{0.20, 0.25, 0.30, 0.35, 0.40, 0.45, 0.50\}$. For Concept Drift Detection, we take the loss threshold Δ which brings the highest accuracy in raw dataset. The tolerance $\delta \in \{0.10, 0.15, 0.30\}$, PH threshold $\lambda = 0.4$. The TS chunk size is fixed at 5.

4.2 Results on ISMAP

Based on the explainable approach proposed in [28], we test firstly the incremental feature of *ISMAP* by adopting interleaved Test-then-Train strategy with an evaluation procedure. The loss threshold $\Delta \in [0.2, 0.5]$, with a step of 0.05, which controls the sensitivity of system for importing TS instance into Shapelet Extraction process.

¹<https://github.com/JingweiZuo/TSSStreamMining>

²Especially for TS Classification, hundreds or thousands of TS instances can be considered as BIG DATA, due to the high time complexity of algorithms in the domain which are based on exact distance measure. Our algorithm has a time of $\mathcal{O}(N^2 n^3 \log n)$, N: number of instances, n: TS length.

Table 1: Shapelet Datasets in UCR Archive used for Incremental Test (ISMAP)

Type	Name	Train/Test	Class	Length	IG	KW	MM	ISMAP(best)	ISMAP(best)	Comp. Ratio
Simulated	SyntheticControl	300/300	6	60	0.9433	0.9000	0.8133	0.7007	0.35	46.7%
Sensor	Trace	100/100	4	275	0.9800	0.9400	0.9200	1	0.5, 0.45	26.0%
	MoteStrain	20/1252	2	84	0.8251	0.8395	0.8395	0.9169	0.45	60.0%
	SonyAIBO.I	20/601	2	70	0.8453	0.7281	0.7521	0.9151	0.4	95.0%
	SonyAIBO.II	27/953	2	65	0.8457	-	-	0.8583	0.4	63.0%
	ItalyPower.	67/1029	2	24	0.8921	0.9096	0.8678	0.9466	0.45	25.4%
ECG	ECG5000	500/4500	5	140	0.7852	-	-	0.9109	0.4	9.4%
	ECGFiveDays	23/861	2	136	0.7747	0.8721	0.8432	0.9826	0.4	51.2%
	TwoLeadECG	23/1189	2	82	0.8507	0.7538	0.7657	0.9337	0.5	47.8%
Images	Symbols	25/995	6	398	0.7799	0.5568	0.5799	0.8113	0.35	96.0%
	Coffee	28/28	2	286	0.9643	0.8571	0.8671	0.9286	0.4	78.6%
	FaceFour	24/88	4	350	0.8409	0.4432	0.4205	0.9886	except 0.45	62.5%
	DiatomSize.	16/306	4	345	0.7222	0.6111	0.4608	0.8758	0.5	50.0%
Motion	GunPoint	50/150	2	150	0.8933	0.9400	0.9000	0.9733	0.45	42.0%

Baselines: We focus on the feature itself, that is, to select the best quality Shapelets from data source. Therefore, to test the reliability of *ISMAP*, we take the Shapelet Tree methods as baselines, rather than considering classifiers learned over shapelet-transformed data [15]³. The Shapelet Tree methods utilize different quality measures to extract the Shapelet and predict target instance: a) Information Gain (IG) [14], b) Kruskal-Wallis (KW) [32], c) Mood’s Median (MM) [32]. As quality measure’s calculation is negligible compared to the total time cost, the computation time should remain at the same level when they adopt the same distance measure (e.g., *MASS*), and when *ISMAP* doesn’t adopt a Test-then-Train strategy.

Table 1 shows the accuracy performance comparison between baselines and our approach. Obviously, *ISMAP* achieved the top performance on accuracy metric on more datasets than any other classifier (12 of 14). Specifically for sensor, motion and ECG data, *ISMAP* performs no doubt better than other approaches, and achieved more than 20% accuracy improvement in *ECG5000*.

Besides the accuracy advantage compared to the baselines, the incrementality of *ISMAP* allows a flexible adjustment between accuracy and time cost. **Table 1** shows as well the parameter Δ which brings the best accuracy performance. The parameter sets a loss threshold for Shapelet evaluation during incremental extraction process, and decides whether the input TS instance contains useful information for updating existing Shapelets. The Compression Ratio is defined by the proportion of imported valuable instances over total training instances: $Comp.Ratio = \frac{nbr.instance_{imported}}{nbr.instance_{training}}$, the ratio below 1 brings a better performance in both time and memory cost. **Fig. 6** shows a global view of accuracy and time cost tested by *ISMAP* under different loss thresholds.

Figure 6: Results of Incremental Test (ISMAP) by adopting an extra Shapelet Evaluation procedure

³Nevertheless, the extracted high quality Shapelets can be concatenated with Shapelet Transform methods for a higher accuracy, though it’s not our focus here.

In general, we consider that a high loss threshold Δ leads to a high efficiency at the expense of certain accuracy. However, from the result in **Table 1**, for most of the datasets, *ISMAP* gets the highest accuracy in range $\Delta \in [0.4, 0.45]$, which is reasonable from **Fig. 3**. On the one hand, a threshold loss close to or greater than 0.5 will skip a large number of instances which are around the split point and don't make any quality improvement for current Shapelets. Instead, the instances contain valuable information for adjusting the quality of current Shapelets and could introduce new candidates. Then classifier will be overfit on small number of instances. On the other hand, the accuracy performance is affected by random input order of TS instances, the first extracted candidate Shapelets will affect the evaluation of later TS instances, which brings a small uncertainty for the number of instances to be imported into the system. For instance, in *FaceFour* dataset, time cost increases when Δ changes from 0.3 to 0.35, that means more instances are imported into system even with the increase of threshold. Which is caused by the uncertainty. A small loss threshold in early stage will be biased towards the initial randomly imported TS instances and extracted Shapelets, and reduce the acceptance space for true discriminative Shapelets coming afterwards.

Nevertheless, most of the time the accuracy keeps on a relative stable stage even with the increase of Δ , which can be explained by the fact that the instances from the same class are highly consistent, and share the common Shapelet features. Therefore, the system efficiency can be largely improved with an exchange of a negligible decrease of accuracy.

4.3 Results on adaptive features in TS Stream

Figure 7: Results of Concept Drift Detection on *ECG5000* dataset and augmented *Trace* dataset

For sake of clarity, we have selected for these experiments datasets where Shapelet-based approach has a strong reliability. Due to space limit, we only show the exploration results for two testing contexts: *ECG5000* dataset in the original space and *Trace* dataset in the augmented space. We focus on the explainable detection process of Concept Drift and the reliability of the adaptive Shapelets around drift areas.

Fig. 7 shows the Concept Drift detection process on the two datasets under different loss tolerance level. For *ECG5000*, two drifts were put at time tick 167 and 333. The Concept Drifts were detected by the system within the time periods [170, 195] and [340, 390], which are in strong accordance with the true drifts in the dataset. The extracted Shapelets before each drift period contain overall information of previous subset. During the drift periods, TS chunks are evaluated to update the current learned concept. A small adjustment time period (i.e., 25 and 50 for two drifts respectively) over the entire subset size proved the strong adaptability of the system. For augmented *Trace* dataset, two drifts were put at time tick 333 and 667. The drift detection mechanism stays reliable under different loss tolerances, which lead to different time periods for adjusting the learned concept. A high tolerance is capable of relieving the effect of outliers or excessive feedback, and allows only a continuous high loss to be considered as a Concept Drift. Therefore, less chunks are imported into the system which leads to less time cost. From the memory and time plot on right **Fig. 7**, at the end of each drift adjustment area, the cached information is largely eliminated, finally only 65 of 500 instances of *ECG5000* dataset, 100 of 1000 ($\delta = 0.15$) or 50 of 1000 ($\delta = 0.30$) instances of *Aug. Trace* dataset, are cached in the memory. The proactive caching elimination mechanism shows its elastic feature. Besides, the later imported chunks requires usually longer calculation time (the step becomes longer in the time plot), as more chunks have been cached.

Table 2: Evaluation in datasets with manually added drift

Dataset	-	i(Con. 1)	ii(Con. 2)	iii(Con. 2)	iv(Con. 3)
<i>Aug.Trace</i> ($\delta = 0.15$)	Time tick	345	380	670	790
	Test Accu.	0.9600	0.9900	0.9900	0.9800
<i>Aug.Trace</i> ($\delta = 0.30$)	Time tick	350	365	675	700
	Test Accu.	0.9600	0.9800	0.9800	0.9700
<i>ECG5000</i> ($\delta = 0.10$)	Time tick	170	195	340	390
	Test Accu.	0.9018	0.8783	0.8783	0.8927

In **Table 2**, we show the reliability of the Extracted Shapelets on 4 time ticks at the beginning/end of each drift area⁴. The extracted Shapelets perform the same accuracy at the two middle time ticks, which can be explained by the fact that no chunks were imported since the learned *Concept 2* was deemed enough reliable. Globally, the adaptive Shapelets show a high accuracy in such a streaming context with Concept Drift, although the accuracy is a little lower than that in **Table 1**, as the training on the subsets gets less information than that on the entire dataset.

5 Conclusion

In this paper, we studied the dynamic feature exploration over Time Series Stream, which is based on the interpretable Shapelet features and an explainable Shapelet extraction process. An incremental Shapelet extraction under stable concept with a novel Shapelet evaluation process is proposed, which improved largely the system’s efficiency with an exchange of a negligible decrease of accuracy. As for a non-stable concept data source, we adjust the conventional strategies of Concept Drift detection into the context of Time Series Stream, which opens the path for a proactive elimination of data cached in the memory. The system can be applied in the scenario where an existing dataset should be enriched with new knowledge but without human loop in the middle.

In the future work, we aim at exploring more challenging scenarios where TS instances are weakly labelled, and the Concept Evolution happens in the coming TS stream. More TS features beyond the Shapelet domain and their compatibility with streaming context will be considered as well.

Acknowledgements

This research was supported by DATAIA convergence institute as part of the *Programme d’Investissement d’Avenir*, (ANR-17-CONV-0003) operated by DAVID Lab, University of Versailles Saint-Quentin, and MASTER project that has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie-Slodowska Curie grant agreement N. 777695.

References

- [1] Tomasz Górecki and Maciej Łuczak. Using derivatives in time series classification. *Data Mining and Knowledge Discovery*, 26(2):310–331, 2013.

⁴We recall that the *Trace* testing dataset were augmented in the same manner, where drifts were manually added.

- [2] Tomasz Górecki and Maciej Łuczak. Non-isometric transforms in time series classification using dtw. *Know.-Based Syst.*, 61:98–108, 2014.
- [3] Jason Lines, Anthony Bagnall, J Lines, · A Bagnall, and A Bagnall. Time series classification with ensembles of elastic distance measures. *Data Min Knowl Disc*, 29:565–592, 2015.
- [4] Chotirat Ann Ratanamahatana and Eamonn Keogh. Three Myths about Dynamic Time Warping Data Mining. In *Proceedings of the 2005 SIAM International Conference on Data Mining*, pages 506–510. 2005.
- [5] Houtao Deng, George C. Runger, Eugene Tuv, and Vladimir Martynov. A time series forest for classification and feature extraction. *Inf. Sci.*, 239:142–153, 2013.
- [6] George Runger Mustafa Gokce Baydogan and Eugene Tuv. A Bag-of-Features Framework to Classify Time Series. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(10):2104–2111, 2017.
- [7] Mustafa Gokce Baydogan, George Runger, and Eamonn B Keogh Mustafa Gokce Baydogan. Time series representation and similarity based on local autopatterns. *Data Mining and Knowledge Discovery*, 30:476–509, 2016.
- [8] Eamonn J. Keogh, Kaushik Chakrabarti, Michael J. Pazzani, and Sharad Mehrotra. Dimensionality reduction for fast similarity search in large time series databases. *KAIS*, 3:263–286, 2001.
- [9] Jessica Lin, Eamonn Keogh, Stefano Lonardi, and Bill Chiu. A symbolic representation of time series, with implications for streaming algorithms. *Pro. of the 8th ACM SIGMOD workshop - DMKD '03*, page 2, 2003.
- [10] Qi Lei, Jinfeng Yi, Roman Vaculin, Lingfei Wu, and Inderjit S. Dhillon. Similarity preserving representation learning for time series clustering. In *IJCAI'19*, pages 2845–2851, 2019.
- [11] Jessica Lin, Rohan Khade, and Yuan Li. Rotation-invariant similarity in time series using bag-of-patterns representation. *Journal of Intelligent Information Systems*, 39(2):287–315, 2012.
- [12] P. Senin and S. Malinchik. Sax-vsm: Interpretable time series classification using sax and vector space model. In *ICDM'13*, pages 1175–1180, 2013.
- [13] Patrick Schäfer. The BOSS is concerned with time series classification in the presence of noise. *Data Min Knowl Disc*, 29:1505–1530, 2015.
- [14] Lexiang Ye and Eamonn Keogh. Time series shapelets: A New Primitive for Data Mining. *Proc. KDD '09*, page 947, 2009.
- [15] Jason Lines, Luke M. Davis, Jon Hills, and Anthony Bagnall. A shapelet transform for time series classification. In *Proc. KDD '12*, pages 289–297, New York, NY, USA, 2012. ACM.
- [16] Josif Grabocka, Nicolas Schilling, Martin Wistuba, and Lars Schmidt-Thieme. Learning time-series shapelets. In *Proc. KDD '14*, pages 392–401, New York, NY, USA, 2014. ACM.
- [17] Thanawin Rakthanmanon and Eamonn Keogh. Fast Shapelets: A Scalable Algorithm for Discovering Time Series Shapelets. In *Proceedings of the 2013 SIAM International Conference on Data Mining*, pages 668–676. 2013.
- [18] Zicheng Fang, Peng Wang, and Wei Wang. Efficient learning interpretable shapelets for accurate time series classification. *Proc. ICDE 2018*, pages 497–508.
- [19] Jason Lines, Jon Hills, and Anthony Bagnall. The Collective of Transformation-Based Ensembles for Time-Series Classification. *IEEE Transactions on Knowledge and Data Engineering*, 27(9):2522–2535, 2015.
- [20] J. Lines, S. Taylor, and A. Bagnall. Hive-cote: The hierarchical vote collective of transformation-based ensembles for time series classification. In *ICDM*, pages 1041–1046, Dec 2016.
- [21] Ahmed Shifaz, Charlotte Pelletier, François Petitjean, and Geoffrey I. Webb. Ts-chief: A scalable and accurate forest algorithm for time series classification. *ArXiv*, abs/1906.10329, 2019.
- [22] Albert Bifet, Geoff Holmes, Richard Kirkby, and Bernhard Pfahringer. MOA: Massive Online Analysis. Technical report, 2010.
- [23] Ken Ueno, Aopeng Xi, Eamonn Keogh, and Dah Jye Lee. Anytime classification using the nearest neighbor algorithm with applications to stream mining. *Proceedings - IEEE International Conference on Data Mining, ICDM*, (December):623–632, 2006.
- [24] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre Alain Muller. Deep learning for time series classification: a review. Technical report, 2019.
- [25] Fazle Karim, Somshubra Majumdar, Houshang Darabi, and Shun Chen. LSTM Fully Convolutional Networks for Time Series Classification. *IEEE Access*, 6:1662–1669, 2017.

- [26] Jingyuan Wang, Ze Wang, Jianfeng Li, and Junjie Wu. Multilevel Wavelet Decomposition Network for Interpretable Time Series Analysis. *KDD*, page 10, 2018.
- [27] Angus Dempster, François Petitjean, and Geoffrey I. Webb. ROCKET: Exceptionally fast and accurate time series classification using random convolutional kernels. *arXiv e-prints*, page arXiv:1910.13051, Oct 2019.
- [28] Jingwei Zuo, Karine Zeitouni, and Yehia Taher. Exploring interpretable features for large time series with se4tec. In *EDBT 2019, Lisbon, Portugal, March 26-29, 2019*, pages 606–609, 2019.
- [29] CCM Yeh, Y Zhu, L Ulanova, and N Begum. Matrix Profile I: All Pairs Similarity Joins for Time Series: A Unifying View that Includes Motifs, Discords and Shapelets. *ICDM*, 2016.
- [30] Aaron Bostrom and Anthony J. Bagnall. Binary shapelet transform for multiclass time series classification. *T. Large-Scale Data- and Knowledge-Centered Systems*, 32:24–46, 2017.
- [31] J Gama, I Zliobait E, A Bifet, M Pechenizkiy, and A Bouchachia. A Survey on Concept Drift Adaptation. *A Survey on Concept Drift Adaptation. ACM Comput. Surv.* 1, 1, Article, 1, 2013.
- [32] Jason Lines and Anthony Bagnall. Alternative quality measures for time series shapelets. In *Proc. IDEAL’12*, pages 475–483, Berlin, Heidelberg, 2012. Springer-Verlag.
- [33] Jingwei Zuo, Karine Zeitouni, and Yehia Taher. ISETS: Incremental Shapelet Extraction from Time Series Stream. *ECML PKDD 2019*, in press.
- [34] Hoang Anh Dau, Eamonn Keogh, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatana, Yanping, Bing Hu, Nurjahan Begum, Anthony Bagnall, Abdullah Mueen, and Gustavo Batista. The ucr time series classification archive, October 2018. https://www.cs.ucr.edu/~eamonn/time_series_data_2018/.
- [35] Anthony Bagnall, Jason Lines, Aaron Bostrom, James Large, and Eamonn Keogh. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery*, 31(3):606–660, 2017.
- [36] Shaghayegh Gharghabi, Shima Imani, Anthony Bagnall, Amirali Darvishzadeh, and Eamonn Keogh. An Ultra-Fast Time Series Distance Measure to allow Data Mining in more Complex Real-World Deployments. Technical report, 2018.
- [37] Abdullah Mueen, Yan Zhu, Michael Yeh, Kaveh Kamgar, Krishnamurthy Viswanathan, Chetan Gupta, and Eamonn Keogh. The fastest similarity search algorithm for time series subsequences under euclidean distance, August 2017. <http://www.cs.unm.edu/~mueen/FastestSimilaritySearch.html>.