

# Fast and Accurate Time Series Classification Through Supervised Interval Search

Nestor Cabello  
The University of Melbourne  
Melbourne, Australia  
ncabello@student.unimelb.edu.au

Elham Naghizade  
RMIT University  
Melbourne, Australia  
e.naghizade@rmit.edu.au

Jianzhong Qi  
The University of Melbourne  
Melbourne, Australia  
jianzhong.qi@unimelb.edu.au

Lars Kulik  
The University of Melbourne  
Melbourne, Australia  
lkulik@unimelb.edu.au

**Abstract**—*Time series classification (TSC)* aims to predict the class label of a given time series. Modern applications such as *appliance modelling* require to model an abundance of long time series, which makes it difficult to use many state-of-the-art TSC techniques due to their high computational cost and lack of interpretable outputs. To address these challenges, we propose a novel TSC method: the *Supervised Time Series Forest (STSF)*. STSF improves the classification efficiency by examining only a (set of) sub-series of the original time series, and its tree-based structure allows for interpretable outcomes. STSF adapts a top-down approach to search for relevant sub-series in three different time series representations prior to training any tree classifier, where the relevance of a sub-series is measured by feature ranking metrics (i.e., supervision signals). Experiments on extensive real datasets show that STSF achieves comparable accuracy to state-of-the-art TSC methods while being significantly more efficient, enabling TSC for long time series.

**Index Terms**—Time series classification, Interval-based classifier, Feature selection.

## I. INTRODUCTION

*Time series classification (TSC)* aims to predict the class label of a given time series (or parts of it). A time series is an ordered time-stamped sequence of data points (i.e., observations from a variable of interest). Various TSC methods have been proposed for a rich set of application areas such as medicine (e.g., classification of electrocardiograms [1]). Modern applications such as *appliance modelling* and *stress detection* pose significant challenges to the existing TSC methods because (i) their data is usually collected via sensors, resulting in large sets of time series with high update frequencies which are also prone to signal noise (e.g., sensor defects) and (ii) they require interpretable models. For example, a classifier for appliance modelling should report the time intervals that differentiate the time series from different types of appliances, enabling householders to understand their energy consumption.

Ideally, a TSC method fulfills three criteria: it is accurate, efficient and interpretable. However, accurate TSC methods suffer in either efficiency and/or interpretability for modern applications that usually involve long time series. For example, the state-of-the-art TSC methods, *hierarchical vote collective of transformation-based ensembles (HIVE-COTE)* [2] and *time series combination of heterogeneous and integrated embedding forest (TS-CHIEF)* [3] have bi-quadratic and quadratic

time complexity in the length of the series, respectively. None of these methods provide *interpretable* classification results.




Fig. 1. Discriminatory *interval 1* versus non-discriminatory *interval 2*.

To address *efficiency and interpretability*, novel techniques known as *interval-based* classifiers have focused on intervals rather than the whole signal. *Time series forest (TSF)* [4] is the most accurate interval-based method and one of the fastest TSC methods [5]. TSF randomly selects a number of intervals (i.e., sub-series). Data points within each interval are aggregated with statistical measurements (e.g., mean). The original series are thus transformed into a representation of intervals. The classification is performed using an ensemble of trees, which allows for interpretability.

Motivated by the better efficiency of interval-based methods and the intrinsic capability of providing interpretable classifications through the identification of *discriminatory intervals* (cf. Figure 1), we also build our time series classifier based on intervals. As Fig. 1 shows, *interval 1* is an example of a discriminatory interval. This interval differentiates the blue time series from the red ones. In contrast, *interval 2* is non-discriminatory because it cannot separate the red and the blue series. When classifying large datasets with long series, TSF may become computationally expensive due to the larger number of intervals to be explored. We show that a “supervised” selection of intervals is significantly more efficient through training an ensemble of trees with a pre-computed set of discriminatory intervals.

Our algorithm, *Supervised Time Series Forest (STSF)*, opportunistically creates a time series forest for classification and feature extraction based on such discriminatory intervals. For a given set of time series with length  $m$ , the number of possible interval features is quadratic to the length of the series, i.e.,  $\mathcal{O}(m^2)$ . Using a binary-inspired search strategy jointly with a feature ranking metric (i.e., Fisher score [6]), STSF reduces

the interval feature space to  $\mathcal{O}(\log m)$ . The main contributions of this paper are summarized as follows:

- We propose STSF, a supervised interval search technique with feature ranking metrics that is a highly efficient and accurate interval-based TSC method. It minimizes unnecessary computations when finding relevant intervals, making it an order of magnitude faster than TSF.
- By extracting relevant intervals from the trained trees, the generated regions of interest, i.e., groups of relevant intervals, highlight the time periods where differences in shape, distribution, or level, between time series of different classes are more pronounced, and thus provide interpretable outcomes of the classification.
- An extensive experimental study on 85 real datasets [7] shows that STSF achieves classification accuracies competitive to state-of-the-art methods but it is at least two orders of magnitude faster than them. This makes STSF suitable for large datasets with long time series.

## II. RELATED WORK

Time series classification (TSC) methods can be broadly categorised into *instance-based*, *feature-based*, and *ensemble-based*. Instance-based methods such as 1-Nearest Neighbour classifier with *Dynamic Time Warping* (1-NN DTW) [8] classify a time series according to the similarity of its ordered data points to those of time series with known class labels, where localized distortions of the data points are allowed. These methods tend to be less accurate than feature-based classifiers and computationally inefficient for long time series (i.e., quadratic to the length of the series).

Feature-based methods such as *Shapelet Transform* (ST) [9], *bag-of-SFA-symbols* (BOSS) [10], and *time series forest* (TSF) [4] represent the raw time series with a set of derived properties, i.e., features. ST uses shapelets, i.e., sub-series that are representative of class membership, and time series are classified according to their similarity to the (discriminatory) shapelets. BOSS uses the relative frequency of discriminatory sub-series for classification. It computes the number of times that a discriminatory sub-series (represented as a word) appears in the time series. ST is expensive with a bi-quadratic time complexity in the length of the series whereas BOSS has a word length complexity in the size of the alphabet, which makes it memory intensive for long series. TSF randomly selects a number of intervals to extract discriminatory features using statistical measurements. The idea is that time series from the same class tend to have intervals with similar characteristics. The *random interval spectral ensemble* (RISE) was recently introduced [2] to capture the frequency-domain features. RISE works similar to TSF, but extracts spectral features over each random interval instead of statistical measurements as in TSF. Both TSF and RISE are highly efficient but usually less accurate than other TSC methods.

Ensemble-based methods use ensembles of individual TSC methods. Two representatives are the *elastic ensemble* (EE) [11] and the *hierarchical vote collective of*

*transformation-based ensembles* (HIVE-COTE) [2]. Both approaches are highly accurate but costly. They may become too expensive on long series. EE is based on 11 instance-based TSC methods, hence it has a quadratic time complexity in the length of the time series. HIVE-COTE uses EE, ST, BOSS, TSF, and RISE for classification. HIVE-COTE is the state-of-the-art TSC method in terms of accuracy but is impractical for long series (bi-quadratic time complexity in the length of the time series). *Proximity forest* (PF) [12] and *time series combination of heterogeneous and integrated embedding forest* (TS-CHIEF) [3] focus on providing highly accurate and scalable classifications. PF builds an ensemble of proximity trees using elastic distance measures as splitting criteria. While PF is more scalable (quasi-linear time complexity in the number of time series) than HIVE-COTE, it is significantly less accurate. TS-CHIEF builds on PF and incorporates BOSS and RISE features as splitting criteria. TS-CHIEF is statistically similar to HIVE-COTE in classification accuracy, but more scalable (similar to PF). PF and TS-CHIEF are quadratic to the series length, which make them costly for long series.

Deep learning classifiers such as *fully convolutional networks* (FCN) and *residual networks* (ResNet) [13] obtain competitive classification accuracy and allow interpretations on the model decisions. However, they are also computationally expensive for long series and require GPU.

## III. PRELIMINARIES

Interval-based approaches extract sub-series for which aggregates (e.g., mean and standard deviation) are computed and used as features (i.e., interval features). To allow for interpretable classifications, we focus on *phase-dependent intervals*, i.e., discriminatory features located at the same time regions over all time series in a given dataset.

**Interval feature.** Given a set of time series  $X = \{\mathbf{x}^1, \mathbf{x}^2, \mathbf{x}^3, \dots, \mathbf{x}^n\}$ , where  $\mathbf{x}^i = \{x_1^i, x_2^i, \dots, x_m^i\}$ , an aggregation function  $f(\cdot)$ , and an interval  $(s, e)$ , an interval feature  $\mathbf{a} = f(X, s, e)$  is a vector of length  $n$ , defined as follows:

$$\mathbf{a} = \{f(\mathbf{x}^1, s, e), f(\mathbf{x}^2, s, e), f(\mathbf{x}^3, s, e), \dots, f(\mathbf{x}^n, s, e)\}$$

where  $f(\mathbf{x}^i, s, e) = f(\{x_s^i, x_{s+1}^i, \dots, x_{e-1}^i, x_e^i\})_{1 \leq s \leq e \leq m}$ .

With  $s=2$ ,  $e=4$ ,  $f=\text{mean}$ , an interval feature  $\mathbf{a} = \text{mean}(X, 2, 4)$  is represented with the dashed rectangle as follows:

$$X = \begin{pmatrix} x_1^1 & x_2^1 & x_3^1 & x_4^1 & x_5^1 & x_6^1 & \dots & x_{m-1}^1 & x_m^1 \\ x_1^2 & x_2^2 & x_3^2 & x_4^2 & x_5^2 & x_6^2 & \dots & x_{m-1}^2 & x_m^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_1^n & x_2^n & x_3^n & x_4^n & x_5^n & x_6^n & \dots & x_{m-1}^n & x_m^n \end{pmatrix}$$

**Problem statement.** We consider a set of  $n$  univariate time series  $X = \{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n\}$ , where each time series  $\mathbf{x}^i = \{x_1^i, x_2^i, \dots, x_m^i\}$  has  $m$  ordered real-valued observations, sampled at equally-spaced time intervals. Each time series  $\mathbf{x}^i$  is also associated with a class label  $y^i$ . We aim to find the set of interval features that yield the highest time series class

prediction accuracy. Finding such a set of interval features is NP-hard. For a time series of length  $m$ , there are  $\mathcal{O}(m^2)$  different intervals, and hence  $\mathcal{O}(2^{m^2})$  subsets of intervals. For a large  $m$ , it is prohibitively expensive to explore all subsets. We present an efficient heuristic to avoid the exhaustive search while retaining a high classification accuracy.

#### IV. OUR APPROACH

We take a stochastic optimization approach to select a set of interval features with a high discriminating power  $\bar{A}$  (i.e., *candidate discriminatory interval features*) from the high dimensional interval feature space  $A$  (i.e., all  $\mathcal{O}(m^2)$  possible interval features). For a time series of length  $m$ , we reduce the interval feature space size to  $\mathcal{O}(\log m)$ . We search for the best interval feature subset  $A^*$  (i.e., *discriminatory interval features*) through an ensemble of decision trees, which has a time complexity  $\mathcal{O}(r \cdot n \cdot \log n \cdot \log m)$ , where  $r$  is the total number of trees in the ensemble, and  $n$  is the number of time series instances. Fig. 2 shows an overview of STSF when training a single tree. For a given training set  $X$ , a periodogram representation of  $X$ , and a derivative representation of  $X$ , the possible number of interval features for each representation is denoted as  $A_O$ ,  $A_F$ , and  $A_D$ , respectively. For each interval feature set  $A_O$ ,  $A_F$ , and  $A_D$ , STSF selects a group of candidate discriminatory interval features  $\bar{A}_O$ ,  $\bar{A}_F$ , and  $\bar{A}_D$ . Lastly, a tree classifier, due to its intrinsic feature selection capability, enables the selection of a set of discriminatory intervals  $A^*$  with which classification is performed on the testing time series. Features from  $A^*$  may be used to add interpretability to the classification task as discussed in Section VI.




Fig. 2. Overview of STSF. Sets of candidate discriminatory interval features  $\bar{A}_O, \bar{A}_F, \bar{A}_D$  are selected from the high dimensional interval feature sets  $A_O, A_F, A_D$ , respectively. The tree classifier uses the candidate sets to select a set of discriminatory interval features  $A^*$  to further perform the TSC task.

##### A. Time series representation

We first discuss time series representations to extract intervals features. We use intervals from original (i.e., time domain), periodogram (i.e., frequency domain), and derivative representations. We focus on the latter two representations.

**Periodogram representation:** Several TSC methods [2], [14] use the *periodogram representation* when looking for time series similarities in the frequency domain. We adopt this and exploit the periodogram representation of each time series derived from the *discrete Fourier transform*. A side benefit of this representation is that it helps to *indirectly* detect *phase-independent discriminatory intervals*, i.e., discriminatory features located at different time regions of the original series as described in Fig. 3.




Fig. 3. (a) Time series  $x^i$  and  $x^k$ , with a same class label (i.e.,  $y^i = y^k$ ) present a similar sub-series but located at different time regions (i.e., phase-independent). (b) The periodogram representation  $p^i$  and  $p^k$  of series  $x^i$  and  $x^k$ , respectively. Our algorithm searches for discriminatory phase-dependent interval features and cannot find discriminatory intervals that identify both series as similar. The periodogram representation provides more flexibility as it considers the frequency of the discriminatory sub-series (ignoring its location in time) and thus helps to identify discriminatory sub-series even when they appear at different locations in time across different time series.

**Derivative representation:** Using a (first-order) difference representation of a given time series rather than the original time series improves the classification accuracy [15] as it provides trend information.

##### B. Selecting Candidate Discriminatory Intervals

Our approach to find candidate discriminatory intervals first partitions a given time series set  $X$  of size  $n \times m$  into two subsets: (i)  $X_L$  of size  $n \times u$ , and (ii)  $X_R$  of size  $n \times m - u$ . The value of  $u$  is randomly selected from the set of numbers  $\{1, 2, \dots, m\}$ . The initial random partition (according to  $u$ ) enables STSF to explore more diverse sub-series (i.e., at different locations and of different lengths). Next, a set of interval features is extracted in a supervised manner independently from sub-sets  $X_L$  and  $X_R$  using Algorithm 1. This last step runs for each aggregation function detailed later. Note that the selection of candidate discriminatory intervals runs on each of the time series representations discussed in Section IV-A.

Our supervised search algorithm for the candidate discriminatory intervals is summarized in Algorithm 1. The algorithm recursively breaks a given interval into two equally-sized intervals and computes interval features for the two resultant intervals (Line 4). The *Fisher score* [6] (detailed below) is further computed for each interval feature (Line 5). The interval feature with a higher score is added to the set of candidate discriminatory interval features (the starting and ending time indices for each interval, and the aggregation function used to represent the interval feature are also stored), and the search continues within this interval (Lines 6 to 12). The algorithm stops when the interval cannot be partitioned further (with less than two points, Lines 1 to 3).

**Feature ranking metric:** The Fisher score of an interval feature indicates how well the feature separates a class of time series from the other classes. For a given vector of class labels  $\mathbf{y} \in \{1, 2, \dots, c\}^n$ , the Fisher score of an interval feature  $\mathbf{a}$  is computed as  $\text{FisherScore}(\mathbf{a}, \mathbf{y}) = \frac{\sum_{k=1}^c n_k (\mu_k^{\mathbf{a}} - \mu^{\mathbf{a}})^2}{\sum_{k=1}^c n_k (\sigma_k^{\mathbf{a}})^2}$ . Here,  $\mu^{\mathbf{a}}$  is the overall mean of the elements in  $\mathbf{a}$ ;  $\mu_k^{\mathbf{a}}$  and  $\sigma_k^{\mathbf{a}}$  are the mean and standard deviation of the elements in  $\mathbf{a}$  labelled with the  $k$ -th class; and  $n_k$  is the number of time series labelled with the  $k$ -th class. We use the Fisher score (instead of other metrics) for its fast computation.

---

**Algorithm 1: SupervisedSearch**

---

**Input:**  $X'$ : set of  $n$  time series of length  $m'$ ;  $y$ : class label vector;  $f$ : aggregation function;  $fr$ : feature ranking metric;  $\bar{A}$ : set of candidate discriminatory intervals.

```
1 if  $m' < 2$  then
2   return  $\bar{A}$ ;
3 else
4    $a_L \leftarrow f(X', 1, m'/2)$ ;  $a_R \leftarrow f(X', m'/2, m)$ ;
5    $score_L \leftarrow fr(a_L, y)$ ;  $score_R \leftarrow fr(a_R, y)$ ;
6   if  $score_L \geq score_R$  then
7      $\bar{A} \leftarrow \bar{A} \cup \{a_L\}$ ;
8     SupervisedSearch( $X'(1 : m'/2), y, f, fr, \bar{A}$ );
9   else
10     $\bar{A} \leftarrow \bar{A} \cup \{a_R\}$ ;
11    SupervisedSearch( $X'(m'/2 : m), y, f, fr, \bar{A}$ );
12  end
13 end
14 return  $\bar{A}$ ;
```

---

### C. Classification with Discriminatory Intervals

STSF uses an ensemble of trees for classification. To train a tree, STSF uses the set of  $\mathcal{O}(\log m)$  candidate discriminatory intervals previously extracted from the series (detailed in Section IV-B). In each tree node, we split the node according to the feature with the highest *information gain* (IG). If the entropy of the node is 0, we label the node as a leaf node. The training process of STSF is less expensive than that in TSF. In TSF, as stated in the original paper: “In each time series tree node, we consider randomly sampling  $\mathcal{O}(\sqrt{m})$  interval sizes and  $\mathcal{O}(\sqrt{m})$  starting positions”. Thus, in TSF, every node (of each tree) has to compute  $\mathcal{O}(\sqrt{m})$  new features to split the node, whereas in STSF each node uses the same  $\mathcal{O}(\log m)$  features for the split. When classifying a new time series, the class is predicted by a majority vote from the tree ensemble.

## V. EXPERIMENTS

We evaluate our approach in two sets of experiments: classification accuracy (i.e., effectiveness) and training time (i.e., computational efficiency). For each set of experiments, we describe the experimental setup and present the results. STSF is implemented in Matlab. All experiments are run on a macOS 10.14.3 system with a dual core CPU (i5, 2.3GHz) using a single thread. We provide the full raw results and our code at <https://github.com/stevcabello/STSF>.

### A. Classification Accuracy

**Experimental Setup:** STSF is tested with a set of 85 benchmark datasets [7]. Each dataset provides a default train and test split. We compare with TSF [4], RISE [2], HIVE-COTE [2], TS-CHIEF [3], PF [12], BOSS [10], FCN [13], ResNet [13], and 1NN-DTW [8]. STSF and TSF are trained with 500 trees. STSF uses the mean, standard deviation (std), median, interquartile range (iqr), minimum (min), maximum

(max), and slope aggregation functions. Similar to every non-deterministic competitor (e.g., TSF, PF, and TS-CHIEF), we compute the average classification accuracy over 10 runs of STSF for each dataset.

**Results:** Table I shows the classification accuracy for STSF and the competitors. STSF shows a competitive average accuracy of 82.60, which is 4% higher than the best existing interval-based method TSF. STSF ranks higher (on average) than 1NN-DTW, TSF, RISE, PF, FCN and BOSS. Only ResNet, HIVE-COTE, and TS-CHIEF rank higher than STSF. However, as the critical difference diagram (Fig. 4) shows, STSF is *not* significantly different from ResNet. Despite trading off accuracy for interpretability (only considering *phase-dependent* discriminatory intervals), STSF achieves a highly competitive accuracy close to the state-of-the-art. Meanwhile, as our next set of experiments will show, STSF is *at least* two orders of magnitude faster than state-of-the-art TSC methods.




Fig. 4. Critical difference diagram of average ranks on 85 benchmark datasets. Bold lines show groups of statistically similar methods. STSF is competitive to exhaustive and state-of-the-art methods.

### B. Training Time

**Experimental Setup:** We measured the training times of TSF, TS-CHIEF, PF, ResNet, and 1NN-DTW over 45 datasets (underlined in Table I). These methods (except for ResNet and TSF) are implemented in Java. While comparing running times of TSC methods implemented in different programming languages may seem biased, STSF is implemented purely in Matlab which is – as an interpreted language – slower than Java (used for the baselines). Parameters of each competitor are set as suggested in its original paper.

**Results:** In Fig. 5, we present the average ranks of each TSC method and their scaled training times. STSF is two orders of magnitude faster than ResNet and TS-CHIEF, and it is also significantly faster than PF. STSF preserves its efficiency even when classifying large datasets with long time series, whereas PF, TS-CHIEF, and ResNet scale poorly. PF uses only 100 trees instead of 500 as TSF, TS-CHIEF, and STSF do. Since PF’s training time scales linearly with the number of trees, using 500 trees will make PF’s time five times larger than as shown in Fig. 5. However, PF’s average rank is unlikely to change. As the PF paper suggests: “It is unlikely that more trees would provide a very significant improvement, because the ratio of error-rates between 100 and 50 (trees) is already close to 1 (i.e., the errors are only slightly reduced)”. Due the expensive training times of HIVE-COTE, we report the training times of HIVE-COTE when classifying 21/45 small to medium datasets. The results suggest that STSF is two orders





of magnitude faster than HIVE-COTE. On large datasets with long time series, this difference is expected to be significantly larger. We also calculated the average number of extracted interval features per tree, i.e., number of computations, when training STSF and TSF. TSF is highly efficient for small datasets/shorts series. However, as the number/length of series starts increasing, TSF becomes expensive. In such cases, STSF is an order of magnitude faster than TSF and requires approximately an order of magnitude less computations.




Fig. 5. Average ranks vs training times of STSF, TSF, 1NN-DTW, PF, TS-CHIEF, and ResNet, when classifying (a) all 45 datasets, (b) the three largest datasets/longest time series NonECGT1, NonECGT2 and StarLtCurves. All training times are scaled by that of STSF, i.e., training time of STSF = 1. Interpretable TSC methods are in green. Non-interpretable ones are in red. STSF is fast, highly accurate, and allows for interpretable classifications.

## VI. INTERPRETABLE CLASSIFICATION

To provide further insights to the classification decisions, we propose the use of *regions of interest* (ROIs). An ROI is a time interval where a difference in levels (*mean, median, min, or max*), distribution (*std or iqr*), or shape (*slope*), between series of different classes is more pronounced. Discriminatory intervals are located in the tree nodes. Each discriminatory interval has information of the interval boundaries (i.e., starting and ending time indices) and the aggregation function (e.g., mean) that was used to generate the corresponding feature. To generate the ROIs, we only consider the discriminatory intervals extracted from the original time series representation. Frequency-based ROIs may be difficult to interpret, and first-order difference-based ROIs are redundant, i.e., slope aggregation on original sub-series also provides shape information. Moreover, discriminatory intervals are likely to overlap with each other, which hinders the ROIs’ interpretability. Thus, we represent ROIs using the intersected regions of such intervals. For example, we extracted *shape ROIs* when classifying the ItalyPowerDemand dataset, which comprises of series of daily household power consumption. Fig. 6 shows the extracted shape ROIs (green areas). A difference in the trend of the electrical power demand is more noticeable in the ROI between 7PM and midnight. This suggests that, in winter (red series), the power demand peaks at 7PM and starts decreasing after this time, whereas in summer (blue series), the power demand is low at 7PM and increases after this time.

## VII. CONCLUSIONS

We proposed STSF, a highly efficient algorithm for interval-based time series classification over high-dimensional datasets. STSF uses three time series representations and a supervised




Fig. 6. ROIs generated from intervals aggregated with slope statistic on ItalyPowerDemand dataset. Summer days are represented with blue series and winter days with red series. ROIs are represented with light green areas.

search strategy combined with a feature ranking metric to reduce unnecessary computations when searching for highly discriminatory interval features on each representation. Discriminatory interval features enable for interpretable outcomes of a classification. Extensive experiments on real-world datasets validate the accuracy and efficiency of STSF, as it is highly accurate, close to state-of-the-art TSC methods but orders of magnitude faster, enabling STSF to classify large datasets with long series. We plan to extend our approach to multivariate time series classification, exploring variables sampled at same and different sampling rates.

## REFERENCES

- [1] S. Karpagachelvi, M. Arthanari, and M. Sivakumar, “Classification of electrocardiogram signals with support vector machines and extreme learning machine,” *Neural Comput Appl*, vol. 21, no. 6, pp. 1331–1339, 2012.
- [2] J. Lines, S. Taylor, and A. Bagnall, “Time series classification with hive-cote: The hierarchical vote collective of transformation-based ensembles,” *TKDD*, vol. 12, no. 5, p. 52, 2018.
- [3] A. Shifaz, C. Pelletier, F. Petitjean, and G. I. Webb, “Ts-chief: A scalable and accurate forest algorithm for time series classification,” *Data Min Knowl Discov*, pp. 1–34, 2020.
- [4] H. Deng, G. Runger, E. Tuv, and M. Vladimir, “A time series forest for classification and feature extraction,” *Inf Sci*, vol. 239, pp. 142–153, 2013.
- [5] A. Bagnall, J. Lines, A. Bostrom, J. Large, and E. Keogh, “The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances,” *Data Min Knowl Discov*, vol. 31, no. 3, pp. 606–660, 2017.
- [6] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern classification*. John Wiley & Sons, 2012.
- [7] A. Bagnall, J. Lines, W. Vickers, and E. Keogh, “The uea & ucr time series classification repository,” 2019, [www.timeseriesclassification.com](http://www.timeseriesclassification.com).
- [8] C. A. Ratanamahatana and E. Keogh, “Three myths about dynamic time warping data mining,” in *SDM*, 2005, pp. 506–510.
- [9] J. Hills, J. Lines, E. Baranauskas, J. Mapp, and A. Bagnall, “Classification of time series by shapelet transformation,” *Data Min Knowl Discov*, vol. 28, no. 4, pp. 851–881, 2014.
- [10] P. Schäfer, “The boss is concerned with time series classification in the presence of noise,” *Data Min Knowl Discov*, vol. 29, no. 6, pp. 1505–1530, 2015.
- [11] J. Lines and A. Bagnall, “Time series classification with ensembles of elastic distance measures,” *Data Min Knowl Discov*, vol. 29, no. 3, pp. 565–592, 2015.
- [12] B. Lucas, A. Shifaz, C. Pelletier, L. O’Neill, N. Zaidi, B. Goethals, F. Petitjean, and G. I. Webb, “Proximity forest: an effective and scalable distance-based classifier for time series,” *Data Min Knowl Discov*, vol. 33, no. 3, pp. 607–635, 2019.
- [13] Z. Wang, W. Yan, and T. Oates, “Time series classification from scratch with deep neural networks: A strong baseline,” in *IJCNN*, 2017, pp. 1578–1585.
- [14] A. Bagnall, L. Davis, J. Hills, and J. Lines, “Transformation based ensembles for time series classification,” in *SDM*, 2012, pp. 307–318.
- [15] T. Górecki and M. Łuczak, “Using derivatives in time series classification,” *Data Min Knowl Discov*, vol. 26, no. 2, pp. 310–331, 2013.