

# On Time Series Classification with Dictionary-Based Classifiers

James Large  
James.Large@uea.ac.uk  
University of East Anglia  
Norwich Research Park  
United Kingdom

Anthony Bagnall  
Anthony.Bagnall@uea.ac.uk  
University of East Anglia  
Norwich Research Park  
United Kingdom

Simon Malinowski  
Simon.Malinowski@irisa.fr  
University of Rennes 1, University of Rennes  
IRISA, LETG-Rennes COSTEL  
France

Romain Tavenard  
Romain.Tavenard@univ-rennes2.fr  
University of Rennes 1, University of Rennes  
IRISA, LETG-Rennes COSTEL  
France

## Abstract

A family of algorithms for time series classification (TSC) involve running a sliding window across each series, discretising the window to form a word, forming a histogram of word counts over the dictionary, then constructing a classifier on the histograms. A recent evaluation of two of this type of algorithm, Bag of Patterns (BOP) and Bag of Symbolic Fourier Approximation Symbols (BOSS) found a significant difference in accuracy between these seemingly similar algorithms. We investigate this phenomenon by deconstructing the classifiers and measuring the relative importance of the four key components between BOP and BOSS. We find that whilst ensembling is a key component for both algorithms, the effect of the other components is mixed and more complex. We conclude that BOSS represents the state of the art for dictionary-based TSC. Both BOP and BOSS can be classed as bag of words approaches. These are particularly popular in Computer Vision for tasks such as image classification. We adapt three techniques used in Computer Vision for TSC: Scale Invariant Feature Transform; Spatial Pyramids; and Histogram Intersection. We find that using Spatial Pyramids in conjunction with BOSS (SP) produces a significantly more accurate classifier. SP is significantly more accurate than standard benchmarks and the original BOSS algorithm. It is not significantly worse than the best shapelet-based or deep learning approaches, and is only outperformed by an ensemble that includes BOSS as a constituent module.

**Keywords:** time series, classification, dictionary

# 1 Introduction

A family of algorithms for time series classification (TSC) involve constructing a dictionary of words from the set of time series then forming a bag of words over that dictionary for each of the time series. More specifically, they run a sliding window across each series, discretise the window to form a word, form a histogram of word counts over the dictionary, then constructing a classifier on the histograms. A recent evaluation of two of this type of algorithm, Bag of Patterns (BOP) and Bag of Symbolic Fourier Approximation Symbols (BOSS), found a significant difference in accuracy between these seemingly similar algorithms. We investigate this phenomena by deconstructing the classifiers and measuring the relative importance of the four key differences between BOP and BOSS. We find that ensembling makes both approaches significantly more accurate, but the effect of the other three components is more complex.

Both BOP and BOSS can be classed as bag of words approaches. These are particularly popular in Computer Vision for tasks such as image classification. Converting approaches for 2-D image classification to 1-D series classification from a range of domains requires careful engineering. We adapt three techniques used in Computer Vision for TSC: Scale Invariant Feature Transform; Spatial Pyramids; and Histogram Intersection. We find that using Spatial Pyramids in conjunction with BOSS (SP) produces a significantly more accurate classifier. SP is significantly more accurate than standard benchmarks and the original BOSS algorithm. It is not significantly worse than the best shapelet-based approach or a residual deep learning network, and is only outperformed by HIVE-COTE, an ensemble that includes BOSS as a constituent module.

The rest of this document is structured as follows. Section 2 provides an overview of related work, from a broad range of TSC algorithms in Section 2.1, to dictionary-based approaches in particular in Section 2.2. We provide an overview of the Computer Vision framework for bag of words classification in Section 3. Section 4 presents the results of our deconstruction of BOP and BOSS and Section 5 describes our evaluation of enhancements to BOSS. We conclude with Section 6.

## 2 Related Work

### 2.1 TSC Background

A recent experimental study [2] compared and evaluated a diverse set of eighteen TSC algorithms that have been published in leading journals and conferences in the last five years. They proposed the following taxonomy of algorithms.

#### 2.1.1 Algorithms based on raw series

Techniques based on raw series compare two series either as a vector (as with traditional classification) or by a distance measure that uses all data points. In the latter case, measures are typically combined with one-nearest-neighbour

(1-NN) classifiers and the simplest variant is to compare series using Euclidean Distance. However, this baseline is easily beaten in practice, and most research effort has been directed toward finding techniques that can compensate for small misalignments between series using specialised elastic distance measures. The almost universal benchmark for whole series measures is Dynamic Time Warping (DTW) but numerous alternatives have been proposed. The most accurate whole series approach (according to the bakeoff comparison [2]) is the Elastic Ensemble (EE) [21], an ensemble of 1-NN classifiers using various elastic measures, including DTW, combined through a proportional voting scheme.

### 2.1.2 Interval-based algorithms

Rather than use the raw series, the interval class of algorithm select one or more phase-dependent intervals of the series. At its simplest, this involves a feature selection of a contiguous subset of attributes. However, the three most effective techniques generate multiple intervals, each of which is processed and forms the basis of a member of an ensemble classifier [11, 6, 5]. There is no significant difference in accuracy between these approaches, and the simplest is the Time Series Forest (TSF) [11].

### 2.1.3 Shapelet-based algorithms

Shapelet approaches are a family of algorithms that focus on finding short patterns that define a class and can appear anywhere in the series. A class is distinguished by the presence or absence of one or more shapelets somewhere in the whole series. Shapelets were first introduced in [29]. The two leading ways of finding shapelets are through enumerating the candidate shapelets in the training set [22, 16] or searching the space of all possible shapelets with a form of gradient descent [15]. The bakeoff found that the shapelet transform algorithm used in conjunction with a heterogeneous classifier ensemble (ST-HESCA) is the most accurate approach on average.

### 2.1.4 Dictionary-based algorithms

Shapelet algorithms look for subseries patterns that identify a class through presence or absence. However, if a class is defined by the relative frequency of a pattern, shapelet approaches will be poor. Dictionary approaches address this by forming frequency counts of repeated patterns. They approximate and reduce the dimensionality of series by transforming into representative words, then compute similarity by comparing the distribution of words. Three of the approaches that have been published in the data mining literature are: Bag of Patterns (BOP) [20]; the Symbolic Aggregate Approximation Vector Space Model (SAXVSM) [28]; and the Bag of Symbolic Fourier Approximation Symbols (BOSS) [26]. We provide an overview of these algorithms in Section 2.2.

### 2.1.5 Spectral-based algorithms

The frequency domain will often contain discriminatory information that is hard to detect in the time domain. Methods include constructing an autoregressive model ([9, 1]) or combinations of autocorrelation, partial autocorrelation and autoregressive features ([3]). An interval-based spectral ensemble called Random Interval Spectral Ensemble (RISE) was proposed in [23] and shown to be more accurate on average than whole series spectral approaches.

### 2.1.6 Combinations

Two or more of the above approaches can be combined into a single classifier. For example, an approach that concatenates different feature spaces is described in [17], forward selection of features for a linear classifier is the method adopted in [13]) and transformation into a feature space that represents each group and ensembling classifiers together formed the basis of the Flat-COTE classifier [3]. A modular meta-ensemble of classifiers from each class of algorithms (EE, TSF, BOSS, ST-HESCA and RISE) called HIVE-COTE is currently the state of the art classifier for TSC when evaluated on the UCR/UEA data and simulated problems [23]. However, on individual problems, there is a wide variation between the classifiers, and the ensemble is not always the best approach. The nature of the discriminatory features will dictate the best class of algorithm.

Our basic assumption is that dictionary classifiers will be best for problems where classes are defined by the frequency of occurrence of a shape in each series rather than its binary presence or absence. For example, suppose data contains short sine waves that repeat at random intervals. In one class there are many repeating patterns, in another class there are few.

A whole series and an interval approach will fail because the positioning of the repeating patterns is random. Shapelets will not detect this phenomena because they look for the presence or absence of a pattern. Spectral approaches may do better, but not if there are large intervals between the signals. A dictionary approach should be able to detect that one pattern occurs more frequently in one class than the other. Our objective here is to develop the best dictionary-based TSC algorithm.

We describe the state of the art by summarising previously published, freely available and reproducible results<sup>1</sup>. We compare the relative performance of three base line classifiers: rotation forest with 50 trees (RotF); 1-NN with Euclidean distance (Euclid); DTW with window set through cross validation (DTW), a representative of each class of algorithm: EE, TSF, ST and RISE, the three dictionary classifiers BOP, SAXVSM and BOSS and two ensemble approaches, Flat-COTE and HIVE-COTE.

To compare multiple classifiers on multiple problems, following the recommendation of Demšar [10], we use the Friedmann test to determine if there were any statistically significant differences in the rankings of the classifiers. However, following recent recommendations in [7] and [14], we have abandoned

---

<sup>1</sup>see [www.timeseriesclassification.com](http://www.timeseriesclassification.com) for details




Figure 1: Average ranks of 12 classifiers on 100 resamples of 85 data sets. The results were first presented in [2] and [24]. A solid bar across a set of classifiers indicates there is no significant difference within that group.

the Nemenyi post-hoc test originally used by [10] to form cliques (groups of classifiers within which there is no significant difference in ranks). Instead, we compare all classifiers with pairwise Wilcoxon signed rank tests, and form cliques using the Holm correction (which adjusts family-wise error less conservatively than a Bonferroni adjustment).

HIVE-COTE is the most accurate algorithm over all, but features of these results excited our interest about dictionary classifiers. Firstly, BOP and SAXVSM performed very poorly. Neither is significantly better than 1-NN Euclidean distance and neither could beat the benchmark classifiers rotation forest and DTW. In stark contrast, BOSS is one of the best performers. It is not significantly worse than the ST-HESCA and only beaten by the two meta ensembles Flat-COTE and HIVE-COTE. HIVE-COTE contains BOSS whereas Flat-COTE does not, and the fact that HIVE-COTE is significantly better than Flat-COTE is further evidence in support of BOSS. On a head to head comparison, BOSS beats BOP on 80 of the 85 datasets. The mean difference in accuracy is over 8%. These algorithms are seemingly similar, so why is BOSS so much better than BOP? Answering this question requires a more in depth understanding of how these algorithms work.

## 2.2 Dictionary-Based Algorithms

Dictionary-based algorithms share the same basic structure. In summary, a window of length  $w$  is passed across each series. Each subseries is then represented by some string or pattern that is representative of it. In the cases considered here, each subseries from the windowing is first compressed from length  $w$  to  $l$ . The shortened subseries are then discretised, so that each of the  $l$  data is restricted

to one of  $\alpha$  values. The occurrence of the resulting ‘word’,  $\mathbf{r}$ , is recorded in a histogram, although in a stage called numerosity reduction, contiguous series of identical words are counted as a single occurrence. Each series has a separate histogram (also referred to as bag). New instances are classified based on the distance between their own histogram and those in the training set, based by default on 1-nearest neighbour classification (although other methods could be used).

There are four key stages at which major differences between dictionary based algorithms may arise:

1. the compression method to get from  $w$  real valued data to  $l$  real valued data;
2. the discretisation technique used to convert the  $l$  real valued data into  $l$  discrete data with  $\alpha$  possible values;
3. the methods of representing the collections of transformed subseries; and
4. the distance measure used to compare histograms and/or the classification algorithm used to classify new cases.

### 2.2.1 Bag of Patterns (BOP)

BOP (described in Algorithm 1) is a dictionary classifier built on the Symbolic Aggregate Approximation (SAX) [19] algorithm. SAX reduces  $w$  to  $l$  through Piecewise Aggregate Approximation (PAA) (i.e. each of the  $l$  new points is an average over an interval length  $w/l$ ) and discretises to  $\alpha$  values using quantiles of the normal distribution. If consecutive windows produce identical words, then only the first of that run is recorded. This is included to avoid the over counting of trivial matches, especially in smooth regions of the originating series. The distribution of words over a series forms a count histogram. To classify new samples, the same transform is applied to the new series and the nearest neighbour histogram within the training matrix found. BOP sets the three parameters through cross validation on the training data.

The Symbolic Aggregate Approximation - Vector Space Model (SAXVSM) [28] combines the SAX representation used in BOP with the Vector Space Model commonly used in Information Retrieval. The key differences between BOP and SAXVSM is that SAXVSM forms word distributions over classes rather than series and weights these by the term frequency/inverse document frequency ( $tf \cdot idf$ ). For SAXVSM, term frequency  $tf$  refers to the number of times a word appears in a class and document frequency  $df$  means the number of classes a word appears in.  $tf \cdot idf$  is then defined as

$$tfidf(tf, df) = \begin{cases} \log(1 + tf) \cdot \log(\frac{c}{df}) & \text{if } df > 0 \\ 0 & \text{otherwise} \end{cases}$$

where  $c$  is the number of classes. There is no significant difference in accuracy between BOP and SAXVSM, so we can without loss of generality restrict our attention to BOP.

---

**Algorithm 1** buildClassifierBOP(A list of  $n$  cases of length  $m$ ,  $\mathbf{T} = (\mathbf{X}, \mathbf{y})$ )

---

**Parameters:** the word length  $l$ , the alphabet size  $\alpha$  and the window length  $w$

```

1: Let  $\mathbf{H}$  be a list of  $n$  histograms ( $\mathbf{h}_1, \dots, \mathbf{h}_n$ )
2:  $\mathbf{p} \leftarrow \emptyset$ 
3: for  $i \leftarrow 1$  to  $n$  do
4:   for  $j \leftarrow 1$  to  $m - w + 1$  do
5:      $\mathbf{q} \leftarrow x_{i,j} \dots x_{i,j+w-1}$ 
6:      $\mathbf{r} \leftarrow \text{SAX}(q, l, \alpha)$ 
7:     if  $\mathbf{r} \neq \mathbf{p}$  then
8:        $pos \leftarrow \text{index}(\mathbf{r})$  {the function index determines the location of the
       word  $\mathbf{r}$  in the count matrix  $\mathbf{h}_i$ }
9:        $h_{i,pos} \leftarrow h_{i,pos} + 1$ 
10:     $\mathbf{p} \leftarrow \mathbf{r}$ 

```

---

### 2.2.2 Bag of Symbolic Fourier Approximation Symbols (BOSS)

BOSS also uses windows to form words over series, and represents them in a simple histogram format, but it has several major differences to BOP and SAXVSM. BOSS uses a truncated Discrete Fourier Transform (DFT) instead of a PAA on each window. Another difference is that the truncated series is discretised through a technique called Multiple Coefficient Binning (MCB), rather than using fixed intervals. MCB finds the discretising break points as a preprocessing step by estimating the distribution of the Fourier coefficients. This is performed by segmenting the series into disjoint windows, performing a DFT, then finding breakpoints for each coefficient such that each bin contains the same number of elements. The whole process of forming words is called Symbolic Fourier Approximation (SFA). BOSS then involves similar stages to BOP; it windows each series to form the term frequency through the application of DFT and discretisation by MCB, performs numerosity reduction, and forms histograms of the words in each series. A bespoke distance function is used for nearest neighbour classification. This non symmetrical function only includes distances between frequencies of words that actually occur within the first histogram passed as an argument, which refers to the test case.

Another major difference is that BOSS forms an ensemble by retaining all classifiers with training accuracy within 92% of the best during the parameter search of window sizes. New instances are classified by a majority vote of the resulting ensemble. Algorithm 2 details the construction of histograms for a given parameter set.

In a manner reminiscent of the way SAXVSM adapts BOP, BOSS-Vector Space (BOSS-VS) [27] modifies BOSS to form class histograms rather than instance histograms. Switching to class histograms massively reduces the memory requirements and speeds up classification. However, it has no significant effect on accuracy, unless it is to reduce it (see the results in [27]). In this work we are concerned with classification accuracy. The questions we address are, firstly, why is BOSS so much better than BOP (see Section 4) and secondly, can we

---

**Algorithm 2** buildClassifierBOSS(A list of  $n$  cases of length  $m$ ,  $\mathbf{T} = (\mathbf{X}, \mathbf{y})$ )

---

**Parameters:** the word length  $l$ , the alphabet size  $\alpha$ , the window length  $w$ , normalisation parameter  $p$

- 1: Let  $\mathbf{H}$  be a list of  $n$  histograms  $(\mathbf{h}_1, \dots, \mathbf{h}_n)$
- 2: Let  $\mathbf{B}$  be a matrix of  $l$  by  $\alpha$  breakpoints found by MCB
- 3:  $\mathbf{p} \leftarrow \emptyset$
- 4: **for**  $i \leftarrow 1$  to  $n$  **do**
- 5:     **for**  $j \leftarrow 1$  to  $m - w + 1$  **do**
- 6:          $\mathbf{s} \leftarrow x_{i,j} \dots x_{i,j+w-1}$
- 7:          $\mathbf{q} \leftarrow \text{DFT}(\mathbf{s}, l, \alpha, p)$  {  $\mathbf{q}$  is a vector of the complex DFT coefficients }
- 8:          $\mathbf{q}' \leftarrow (q_1 \dots q_{l/2})$
- 9:          $\mathbf{r} \leftarrow \text{MCB}(\mathbf{q}', \mathbf{B})$
- 10:         **if**  $\mathbf{r} \neq \mathbf{p}$  **then**
- 11:              $pos \leftarrow \text{index}(\mathbf{r})$
- 12:              $h_{i,pos} \leftarrow h_{i,pos} + 1$
- 13:          $\mathbf{p} \leftarrow \mathbf{r}$

---

refine BOSS to make it more accurate (see Section 5).

### 3 Computer Vision Bag of Words Framework

The histogram approach used by dictionary classifiers has similarities to many approaches used in the field of Computer Vision. A typical Computer Vision bag of words framework involves the following stages:

1. extraction of keypoints;
2. description of keypoints;
3. bag forming; and
4. classification based on bags.

BOP and BOSS extract keypoints through sliding a window over the whole series, reducing the size of the number of keypoints through numerosity reduction and a restriction of window sizes. However, approaches for dictionary-based TSC more in line with the Computer Vision approach have been proposed. [4] describes an approach for using Scale Invariant Feature Transform (SIFT) [25] features for use in TSC with dictionary classifiers. We describe this approach in detail in Section 3.1 and have implemented a version in the WEKA based TSC codebase<sup>2</sup>.

We also consider a common technique in Computer Vision called Spatial Pyramids, proposed in [18] and described in Section 3.2. We try incorporating this as a wrapper for BOSS. It could equally be applied to other dictionary approaches.

---

<sup>2</sup><https://bitbucket.org/TonyBagnall/time-series-classification>



A more complex Computer Vision approach applied to TSC is proposed in [30]. This involves a combined approach of peak finding and hybrid sampling to extract keypoints. It uses Histogram of Oriented Gradients (HOG-1D) and Dynamic Time Warping-Multidimensional Scaling (DTW-MDS) to form features describing the keypoints. It then clusters them with a  $K$  component Gaussian Mixture Model, forming bags based on Fisher Vector encoding. Finally, it constructs a linear kernel Support Vector Machine classifier. The resulting classifier, called HOG-1D+DTW-MDS, is evaluated on the standard single folds of 43 of the UEA/UCR data sets. They do not compare the results of HOG-1D+DTW-MDS to the published results for BOSS, presumably due to the lag time in publication. Using the results in Table 3 from [30] and the BOSS results presented in [2], we find no significant difference between HOG-1D+DTW-MDS and BOSS (HOG-1D+DTW-MDS wins on 22, BOSS on 19 and they tie on 2).

### 3.1 Bag of Temporal SIFT Words (BOTSW) Classifier

The Bag of Temporal SIFT Words (BOTSW) algorithm [4] adopts a version of the Computer Vision bag of words framework that is easier to reproduce than that described in [30], not least because the C++ source code is publicly available<sup>3</sup>. BOTSW first extracts keypoints from every time series through regular sampling at a rate  $r$ , which is a parameter of the method. Then, each keypoint is described by  $n_s$  feature vectors, where  $n_s$  is the number of considered scales. Each feature vector describes the keypoint at a particular scale. To obtain the feature vector of a keypoint at a scale  $s$ , the time series is filtered by a Gaussian filter of width  $s$ .  $n_b$  blocks of size  $\alpha$  are selected around the keypoint. Gradients of the filtered time series are computed for every point of every block and then weighted by a Gaussian function to give greater importance to those points nearer to the keypoint.

Each block is described by two values: the sum of positive gradients in the block and the sum of negative gradients. A feature vector that describes a keypoint at a particular scale is a  $2n_b$ -long vector. A dictionary of feature vectors is learned by a  $k$ -means clustering on the whole set of feature vectors from the time series database. Feature vectors are then quantized using the dictionary. The number of occurrences of these words in the series is computed to form a histogram, which is normalised using Signed Square Root (SSR) then  $l_2$  normalisation. This normalized histogram is the final feature vector representing this series.

In [4], a Support Vector Machine was used to classify feature vectors. However, our objective is to assess the utility of the SIFT features in relation to the BOSS features. Hence, to minimize the differences between BOSS and BOP we use 1NN classification. The parameters  $n_b$  in  $\{4, 8, 12, 16, 20\}$ ,  $\alpha$  in  $\{4, 8\}$ , and  $k$  in  $\{32, 64, 128, 256, 512, 1024\}$  are tuned through a grid search with cross validation. To further align with BOSS, we form an ensemble of BOTSW classifiers, retaining all parameter sets with training accuracy within 92% of the global maximum.

---

<sup>3</sup><https://github.com/a-bailly/dbotsw>

This homogeneous ensemble classifies new instances with a simple majority vote.

### 3.2 BOSS Ensemble with Spatial Pyramids (SP)

The essence of dictionary classifiers is to ignore temporal information through consideration of the recurrence of short subseries. Whilst this will lead to good results in problem domains with repeated discriminatory features, the disadvantage is that in some domains the location in time of a pattern is as important as the pattern itself. Spatial pyramids [18] are a method commonly used in Computer Vision, which will allow us combine temporal and phase independent features. When applied to time series, using a spatial pyramid involves recursively segmenting each series and constructing histograms on the segments.

Starting from the initial histogram across the whole series, histograms on subsections are formed by repeatedly dividing the series  $L$  times. These histograms are weighted by  $\frac{1}{2^{L-l}}$ , which is inversely proportional to the level  $l$  at which they are found. All histograms are then combined and normalised to form a single elongated histogram feature.

Because of the weighting, similarity between features found at smaller divisions on the series have a more significant effect than those found on a more global scale, as their temporal location becomes increasingly dissimilar. It is also worth noting that a pyramid with one level is equivalent to the basic bag of words, as no division has occurred.

Since BOSS ensembles over different window sizes so that discriminatory patterns of different lengths can all be accounted for, we search for  $L$  for each member of the ensemble during training. An overview of the ensemble construction is given in Algorithm 3. Feature sets formed from an optimal word length, found through CV, for a given window size are generated as usual. This feature set, implicitly produced as a pyramid with  $L = 1$ , is then augmented and further CV is performed to find the optimal  $L$  in  $\{1,2,3\}$ . This effectively defines whether the discriminatory feature type described by this parameter set is more local or global in nature. If the training accuracy of the best word length and number of levels for this window size falls within 92% of the best, it is included in the ensemble. In classification, for each member the test instance is transformed into a spatial pyramid using that member’s parameters, and the class of the train pyramid with the maximal Histogram Intersection or minimal BOSS Distance is returned. Figure 2 gives an example of the process of forming histograms for SP.

While computing the pyramids is very fast relative to the original production of the SFA words, the additional space complexity is a concern for large datasets as the final elongated histograms will be  $\sum_{l=0}^{L-1} 2^l$  times larger. This can be heavily mitigated by using sparse data representations, since histograms at higher levels will be more sparse than those at lower levels. However, a cap of 100 was also placed on the maximum size of the ensemble to keep the space requirements more reasonable. Thus if  $\lambda$  is the number of feature sets within the threshold of the max accuracy, the size of the final ensemble is  $\min(100, \lambda)$ .

---

**Algorithm 3** buildBOSSensembleSP(A list of  $n$  cases of length  $m$ ,  $\mathbf{T} = \{\mathbf{X}, \mathbf{y}\}$ )

---

```

1:  $\alpha = 4$ 
2: featureSets = [features, trainAccuracies]
3: for w in windowLengths() do
4:   bestWindowFeatureSet = null
5:   bestWindowAcc = 0
6:   for k in wordLengths() do
7:     featureSet = BOSSTransform(w,k, $\alpha$ )
8:     acc = CrossValidate(featureSet)
9:     if acc > bestWindowAcc then
10:       bestWindowAcc = acc, bestWindowFeatureSet = featureSet
11:   for L in 2,3 do
12:     featureSet = buildPyramid(bestWindowFeatureSet, L)
13:     acc = CrossValidate(featureSet)
14:     if acc > bestWindowAcc then
15:       bestWindowAcc = acc, bestWindowFeatureSet = featureSet
16:   featureSets.add(bestWindowFeatureSet, bestWindowAcc)
17: maxWindowAcc = max(featureSets.trainAccuracies)
18: for set in featureSets do
19:   if bestWindowAcc > maxWindowAcc*0.92 then
20:     addToEnsemble(set)

```

---




Figure 2: A series from the BeetleFly dataset, being divided at successive levels with Bags of SFA words being formed for each subsection. H1...7 are combined to form the final feature vector.

### 3.3 Histogram Intersection (HI) Distance

A core task in any bag of words/dictionary based technique is to compare the differences between the resulting histograms in order to define class membership. BOP uses Euclidean Distance, SAX-VSM uses Cosine Similarity, and BOSS its own measure which is a slight alteration to Euclidean. We also test the Histogram Intersection similarity measure described in [18] which is used in many different applications involving histograms. For a dictionary and resulting histogram size of  $k$ , this is defined as:

$$HI(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^k \min(a_i, b_i)$$

## 4 From BOP to BOSS

We perform all experiments using 77 of the datasets at the University of California, Riverside/University of East Anglia (UCR/UEA) time series classification repository<sup>4</sup> ([2]). There are 8 datasets that we do not use for practical reasons: their size means the classifiers take too long to train or require too much memory to complete given our time frame and the number of experiments and resampling performed. The full list of problems we used is given in Table 4. Our focus is on bridging the accuracy gap between the two classifiers; optimizing for speed and memory are of course very important, but are not the focus of this study. All of our code and data is available from a public code repository and accompanying website<sup>5</sup>. We compare classifiers by the accuracy average over 25 stratified resamples (with the same train/test size of the original data).

Both BOP and BOSS tune their parameters through a leave one out cross validation on the train data for a predefined parameter space. The results for BOSS and BOP presented in [2] were obtained using the parameter space defined in the original papers, and these parameter spaces are different. To alleviate this possible source of bias we have altered the BOP search space to match that of BOSS (see Table 1). In a pairwise comparison between the *BOP* on the old and new parameter space, the latter had higher average accuracy on 44 datasets and worse on 33. There is no significant difference between the old and new versions, and we conclude that we cannot explain the difference between BOP and BOSS on this factor.

Table 1: Parameter search spaces for BOP and BOSS.

Algorithm	Parameters
<i>BOP</i> published parameter search space	$w = 15\% \dots 36\%$ of $m$ $l =$ powers of 2 up to $w/2$ $\alpha = 2 \dots 8$
<i>BOSS</i> published parameter search space This space is used for both BOSS and BOP	$w = 10 \dots m$ $l = 8, 10, 12, 14, 16$ $\alpha = 4$

BOP and BOSS are identical except for four features.

1. **The window approximation method.** Each window of length  $w$  is reduced to a series of length  $l$  through an approximation method. BOP uses Piecewise Aggregate Approximation whereas BOSS uses the truncated Fourier terms.

<sup>4</sup>UCR/UEA TSC Repository: [www.timeseriesclassification.com](http://www.timeseriesclassification.com)

<sup>5</sup>[www.timeseriesclassification.com/DictionaryClassifiers.php](http://www.timeseriesclassification.com/DictionaryClassifiers.php)

2. **The discretisation method.** Each value in the approximate series of length  $l$  is discretised into one of  $\alpha$  values. BOP uses the fixed intervals defined in SAX whereas BOSS uses the data driven technique Multiple Coefficient Binning (MCB).
3. **The distance measure.** BOP uses 1-NN with Euclidean distance whereas BOSS uses a 1-NN with a bespoke, non-symmetric distance function that ignores zero entries in the test histogram.
4. **The classifier.** BOSS is an ensemble of multiple transforms with different parameters, whereas BOP is a single classifier.

To quantify the source of the difference in BOP and BOSS, we assess the relative importance of each of these components by adding each of the four BOSS features into BOP. We then measure their importance to BOSS by in turn replacing each BOSS feature with that used in BOP. This gives us the 10 BOP/BOSS variants listed in Table 2.

Table 2: BOP and BOSS variants, with the switching of BOSS and BOP features. For clarity, we indicate the variant by the addition or removal of the boss feature. e.g. BOP+Ens is BOP with an added BOSS-like ensemble, whereas BOSS-BD is BOSS with Euclidean distance replacing BOSS distance.

Algorithm	Label
Base BOP classifier	<i>BOP</i>
BOP with FT approximation replacing PAA	<i>BOP + FT</i>
BOP with MCB discretisation replacing Gaussian breakpoints	<i>BOP + MCB</i>
BOP with BOSS Distance measure replacing Euclidean Distance	<i>BOP + BD</i>
BOP with ensembling over best parameter sets	<i>BOP + Ens</i>
Base BOSS classifier	<i>BOSS</i>
BOSS with PAA approximation replacing FT	<i>BOSS - FT</i>
BOSS with Gaussian breakpoint discretisation replacing FT	<i>BOSS - MCB</i>
BOSS with Euclidean Distance metric replacing BOSS Distance	<i>BOSS - BD</i>
BOSS with single best parameter set	<i>BOSS - Ens</i>

Figure 3 shows the average ranks of the 10 variants of BOP and BOSS we have evaluated. Table 3 shows the mean difference in accuracy between the variants. The mean difference between the best and the worst variant is nearly 9%. To be clear, this is the absolute pairwise difference in accuracy; the worst algorithm, BOP+BD has an average accuracy over all problems of 73.8%, whereas BOSS, the best algorithm, has average accuracy of 82.65%.

We can make the following observations from these results. For BOP, using DFT with fixed boundaries, or PAA with MCB discretisation, makes no difference to using SAX. This suggests the benefit of using spectral features only comes when using bespoke bins to discretise. Using BOSS distance for BOP histograms actually makes BOP significantly worse. The only component of BOSS that significantly improves BOP is ensembling. This actually makes BOP significantly better than a single version of BOSS (the mean difference 1.32%). However,




Figure 3: Average ranks and cliques of 10 BOP/BOSS classifiers on 25 resamples of 77 data sets.

Table 3: The mean difference in accuracy between BOP and BOSS variants over 77 datasets.

	BOP+BD	BOP+MCB	BOP	BOP+FT	BOSS-FT	BOSS-Ens	BOP+Ens	BOSS-BD	BOSS-MCB	BOSS
BOP+BD	-	-0.97%	-0.97%	-0.71%	-3.10%	-3.61%	-4.94%	-5.59%	-7.60%	-8.88%
BOP+MCB	0.97%	-	-0.01%	0.26%	-2.13%	-2.65%	-3.97%	-4.62%	-6.63%	-7.91%
BOP	0.97%	0.01%	-	0.26%	-2.12%	-2.64%	-3.96%	-4.61%	-6.63%	-7.91%
BOP+FT	0.71%	-0.26%	-0.26%	-	-2.39%	-2.90%	-4.23%	-4.88%	-6.89%	-8.17%
BOSS-FT	3.10%	2.13%	2.12%	2.39%	-	-0.51%	-1.84%	-2.49%	-4.50%	-5.78%
BOSS-Ens	3.61%	2.65%	2.64%	2.90%	0.51%	-	-1.32%	-1.98%	-3.99%	-5.27%
BOP+Ens	4.94%	3.97%	3.96%	4.23%	1.84%	1.32%	-	-0.65%	-2.66%	-3.94%
BOSS-BD	5.59%	4.62%	4.61%	4.88%	2.49%	1.98%	0.65%	-	-2.01%	-3.29%
BOSS-MCB	7.60%	6.63%	6.63%	6.89%	4.50%	3.99%	2.66%	2.01%	-	-1.28%
BOSS	8.88%	7.91%	7.91%	8.17%	5.78%	5.27%	3.94%	3.29%	1.28%	-

BOP ensemble is still significantly worse than the BOSS ensemble. Hence we cannot attribute the difference to the ensembling method alone: removing any one of the four components of BOSS makes it significantly worse. The worst change to make is to switch from DFT features to PAA features. This surprised us, as we had assumed removing ensembling would cause the most harm. Clearly the four features of BOSS that differentiate it from BOP are all required and all interact to produce a better classifier. This highlights that the engineering of algorithms is often not linear, and components can interact in ways that may not be intuitively obvious. This is most clearly observable when comparing the effect of using the BOSS Distance (BD): BD makes BOP significantly worse but BOSS significantly better.

There is some difference in the structures of the resulting histograms of each transform which BOSS Distance is able to leverage over Euclidean Distance in the case of BOSS, but not BOP. Considering the actual bagging process is essentially the same between the two - both extract words from a massively sparse space,

and both use numerosity reduction - such an apparently clear, or rather ‘usable’, distinction in the final histograms is striking. BOSS Distance ignores words that do not appear in the test histogram, so for BOP these missing words are seemingly informative. However, for BOSS they are noise and removing them is beneficial. Exactly why this is so is unclear. We suspect this difference is due to the action of MCB, which creates a data driven discretisation. If the underlying distribution diverges significantly from normality, MCB will create a more accurate representation and will hence capture an underlying pattern more accurately. This could lead to the truly informative words being separated from the uninformative noise more successfully, and so the words that are ignored are more likely to be noise. This is just conjecture. Further work and analysis of the resulting histograms would need to be performed to fully understand the mechanisms at work here, however this is beyond the scope of this work.

From these experiments we conclude that BOSS, as described in [26], represents the state of the art for dictionary classifiers that were first introduced in [20]. Our next question is, can we improve on the state of the art?

## 5 BOSS Extensions

In Section 3 we described two approaches from Computer Vision that may improve dictionary classifiers: SIFT features [25], adapted for time series as described in [4] and Spatial Pyramids [18] that have not formerly been used in this context. We also described the Histogram Intersection as an alternative distance measure between histograms. We wish to assess whether adding these alternative structures to the state of the art dictionary classifier gives an overall improvement in accuracy. To try and isolate the causes of any observed differences we start with BOSS as the base classifier and make the minimum adjustments.

$SP_{BD}$  is a spatial pyramid built on top of the standard BOSS ensemble (using BD).  $SP_{HI}$  is a spatial pyramid built on top of BOSS, using histogram intersection.  $BOTSW_{BD}$  is a bag of temporal-sift classifiers that use BD, whereas  $BOTSW_{HI}$  uses histogram intersection. All four classifiers ensemble in an identical way to BOSS (retain all models within 92% of the best), and each pair of classifiers (i.e.  $BOTSW_{BD}/BOTSW_{HI}$  and  $SP_{BD}/SP_{HI}$ ) search identical parameter spaces. The average ranks and cliques are shown in Figure 4. The two SIFT-based classifiers are significantly worse than BOSS, there is no significant difference between BOSS and  $SP_{BD}$ , but  $SP_{HI}$  is significantly better than BOSS.

$SP_{HI}$  represents an advance for dictionary-based algorithms for TSC, but we do not wish to over sell the importance. Although  $SP_{HI}$  wins on 51 of the 77 problems, that actual differences are small. Figure 5 shows the scatter plot of accuracies of  $SP_{HI}$  against BOSS. The results are fairly tightly grouped on the line of equality. The overall mean difference in accuracy is just 0.64%. We would expect the improvements from  $SP_{HI}$  to be apparent on problems where discriminatory features are phase dependent. For example,  $SP_{HI}$  is over 7% more accurate than BOSS on WordSynonyms and 6% better on FiftyWords. The elastic ensemble is the most accurate classifier on these two problems, which




Figure 4: Average ranks and cliques for five variants of dictionary classifiers.

indicate the discriminatory features are time dependent. Conversely,  $SP_{HI}$  is 6% worse than BOSS on ShapeletSim and 3.5% worse on SonyAIBORobotSurface1. The phase independent classifier ST-HESCA is the most accurate classifier for these datasets, whereas EE does poorly. This indicates that the setting of the parameter  $L$ , like all parameters, is vulnerable to overfitting. Whereas it would have evidently been better to use the regular BOSS classifier (or equivalently setting  $L=1$  in the SP) on those latter datasets due to their phase-independent nature, in terms of training accuracy the parameter search process found some erroneous advantage to using more levels.




Figure 5: Scatter plot of accuracies for (a)  $SP_{HI}$  vs BOSS and (b) EE vs ST-HESCA. The latter is provided to demonstrate the type of spread observable on this data for two very different classifiers.



Table 4: The mean accuracy of five variants of dictionary classifiers. BOSS [26], spatial pyramid BOSS with histogram intersection distance (SP-HI) and BOSS distance (SP-BD) and an adapted version of bag of temporal words [4] with histogram intersection distance (BOTSW-BD) and BOSS distance (BOTSW-HI)

DataSet	BOSS	SP-HI	SP-BD	BOTSW-BD	BOTSW-HI
Adiac	<b>74.94</b>	74.38	74.77	71.62	71.59
ArrowHead	87.52	<b>88.66</b>	87.89	87.36	86.33
Beef	61.5	<b>66.13</b>	65.2	54.93	55.6
BeetleFly	<b>94.85</b>	94.2	93.6	91.6	92.8
BirdChicken	<b>98.4</b>	97.8	98	95.2	92.2
Car	85.5	85.93	85.8	<b>90.33</b>	88.2
CBF	99.81	<b>99.93</b>	99.91	99.88	99.85
ChlorineConcentration	65.96	65.96	<b>65.97</b>	56.8	58.67
CinCECGtorso	90.05	<b>94.39</b>	93.51	78.24	71.72
Coffee	<b>98.86</b>	98.57	98.57	98.29	98.71
Computers	80.23	<b>81.92</b>	81.54	67.66	65.55
CricketX	76.36	78.53	77.96	<b>81.22</b>	79.85
CricketY	74.93	77.31	76.98	<b>78.61</b>	77.04
CricketZ	77.57	78.88	78.51	<b>82.63</b>	82
DiatomSizeReduction	93.94	93.93	<b>94.21</b>	91.19	90.99
DistalPhalanxOC	<b>81.46</b>	80.52	81.25	77.61	76.09
DistalPhalanxOAG	81.41	<b>81.5</b>	81.24	77.44	77.29
DistalPhalanxTW	67.3	<b>67.34</b>	67.14	65.06	65.78
Earthquakes	74.59	74.33	74.76	74.45	<b>74.82</b>
ECG200	<b>89.05</b>	87.04	87.96	86.96	86.76
ECGFiveDays	98.33	99.3	99.21	<b>99.67</b>	99.33
FaceFour	<b>99.56</b>	98.09	98.14	93.95	91.59
FacesUCR	95.06	95.62	<b>95.67</b>	95.31	94.38
FiftyWords	70.22	76.5	<b>76.9</b>	75.94	72.32
Fish	96.87	<b>97.1</b>	96.96	94.38	95.36
GunPoint	99.41	<b>99.84</b>	99.73	98.88	98.75
Ham	83.6	83.47	<b>83.62</b>	77.26	76.15
Haptics	45.87	<b>49.36</b>	47.42	47.04	48.05
Herring	60.53	60.44	59.13	<b>60.94</b>	59.94
InlineSkate	50.26	51.08	<b>51.4</b>	42.49	40.92
InsectWingbeatSound	51.03	51.67	<b>51.82</b>	50.34	46.58
ItalyPowerDemand	86.6	88.22	87.14	<b>93.43</b>	92.96
LargeKitchenAppliances	83.66	<b>83.7</b>	82.02	79.16	78.24
Lightning2	<b>81</b>	80.59	80.2	79.93	79.21
Lightning7	66.56	67.67	68.33	<b>73.53</b>	71.07
Mallat	<b>94.85</b>	94.8	94.8	89.35	89.44
Meat	98.03	<b>98.33</b>	98.27	95.87	96.6
MedicalImages	71.46	71.59	72.19	<b>75.02</b>	72.87
MiddlePhalanxOC	<b>80.82</b>	80.52	80.78	75.86	75.6
MiddlePhalanxOAG	<b>66.6</b>	65.58	65.84	61.3	60.52
MiddlePhalanxTW	53.74	<b>53.87</b>	53.51	53.61	53.4
MoteStrain	84.6	85.49	85.33	<b>90.06</b>	89.28
OliveOil	87	<b>87.47</b>	<b>87.47</b>	86.67	86.93
OSULeaf	96.74	<b>97.79</b>	97.36	87.4	85.54
Phoneme	25.62	<b>27.84</b>	27.51	21.85	18.2
Plane	99.79	<b>99.89</b>	99.81	99.54	99.16
ProximalPhalanxOC	86.74	<b>86.89</b>	86.83	79.82	79.23
ProximalPhalanxOAG	81.9	<b>83</b>	82.85	<b>83</b>	82.79
ProximalPhalanxTW	77.28	<b>77.62</b>	77.48	75.32	75.3
RefrigerationDevices	<b>78.46</b>	77.26	77.28	57.33	67.34

ScreenType	58.6	<b>58.69</b>	58.56	51.66	45.57
ShapeletSim	<b>100</b>	94.09	<b>100</b>	99.98	98.91
SmallKitchenAppliances	75.02	<b>81.92</b>	77.74	67.07	62.28
SonyAIBORobotSurface1	<b>89.74</b>	86.22	89.36	89.26	86
SonyAIBORobotSurface2	88.77	<b>89.52</b>	88.04	88.21	86.06
SwedishLeaf	91.77	<b>92.51</b>	92.26	89.08	89.21
Symbols	96.12	96.47	96.25	<b>97.06</b>	96.23
SyntheticControl	96.79	96.15	96.68	<b>99.53</b>	98.68
ToeSegmentation1	92.88	91.88	92.44	<b>93.98</b>	91.58
ToeSegmentation2	95.97	96.03	96.15	<b>97.08</b>	95.75
Trace	99.99	<b>100</b>	<b>100</b>	<b>100</b>	<b>100</b>
TwoLeadECG	98.45	<b>98.65</b>	98.54	97.23	96.6
UWaveGestureLibraryY	66.12	<b>72.09</b>	71.79	71.75	69.67
Wine	<b>91.17</b>	90.07	90.07	83.41	84.74
WordSynonyms	65.88	<b>73.62</b>	73.47	72.03	68.56
Worms	<b>73.49</b>	71.64	72.42	72.16	72.31
WormsTwoClass	<b>80.97</b>	80.62	80.62	79.74	78.23
Yoga	90.99	<b>91.89</b>	91.52	88.86	88.71
Wins	16.5	27.25	8.25	14.75	1.25




Figure 6: A comparison of SP-HI to three alternative TSC approaches that are not dictionary-based.

Our main interest here is in understanding and improving dictionary classifiers as a particular representation within TSC. However, for context, we compare SP-HI to three state of the art algorithms from alternative domains: a Residual Network (ResNet) with results taken from [12]; a shapelet transform with a heterogenous ensemble (ST-HESCA) [8]; and an ensemble of classifiers from the different transform domains (HIVE-COTE) [24]. We compare against ST-HESCA and HIVE-COTE as they have been shown to be significantly better than competing methods [2]. The results, shown in Figure 6, demonstrate that SP-HI is not significantly worse than the competing individual representations. It is significantly worse than HIVE-COTE, but this is to be expected over so many data sets because HIVE-COTE ensembles over different representations, and contains BOSS as one of its components. On average, SP-HI is competitive with state of the art classifiers from different problem domains.

## 6 Conclusions

Dictionary classifiers are an important class of TSC algorithm that explicitly use the frequency of occurrence of repeating patterns as classification features. A previous study observed a huge difference in accuracy between two prominent approaches, BOP [20] and BOSS [26]. In order to investigate why this is so, we identify the four key differences between the two algorithms and assess their importance to both BOP and BOSS. We find that only one of these features, ensembling over different parameter values, is beneficial to both BOSS and BOP. Ensembling has proven successful in other domains for TSC [3], and it carries no train time overhead if a parameter search is being conducted. BOP with an ensemble is significantly better than a single BOSS classifier. Hence, we would recommend anyone assessing a new TSC algorithm attempt to ensemble, not least to make a better comparison to the state of the art. For example, it is quite possible that HOG-1D+DTW-MDS [30] would be significantly more accurate if ensembled.

However, there is more to BOSS than the ensemble. The three other distinguishing features: the use of the Fourier transform, data driven discretisation and bespoke distance measure, all have a significant effect overall. This demonstrates that algorithm design is not always a linear process; algorithm components interact in surprising ways. This is most clearly illustrated with distance measures. The BOSS distance makes BOSS significantly more accurate, but it makes BOP significantly worse. The importance of the distance function is further demonstrated with our experiments involving histogram interaction (HI) distance and two alternative dictionary classifiers, bag of temporal SIFT features (BOTSW) and BOSS with Spatial Pyramid (SP). Using HI made BOTSW significantly worse, but it improved SP (albeit not significantly).

Ensembling does have a memory overhead, as each base classifier must be stored. This is particularly memory intensive for histogram-based nearest neighbour classifiers such as BOP and BOSS, and it would be useful to have an algorithm that did not require storing all the histograms in the ensemble. We have experimented with using alternative less memory intensive base classifiers such as C4.5, but this significantly reduced accuracy and massively increased the time to build the classifier. A SVM approach may yield a better classifier with lower memory overhead, but our preliminary experiments showed that the extra training time made this infeasible for a large number of problems. However, it is possible to pursue this further, perhaps through using a condensed data set and/or a proxy classifier for parameter search.

The new approach to dictionary classifiers that combine temporal and dictionary features by using Spatial Pyramids in conjunction with BOSS and HI is significantly more accurate than the standard BOSS ensemble. However, the improvement is small and mostly on problems where BOSS is not the best algorithm, so it is debatable whether the extra memory overhead required by the spatial pyramid is worth the small improvement. We believe the SP approach will be best when discriminatory shape frequency features are embedded in confounding noise. In this situation, the pyramid will facilitate higher pattern

resolution in certain areas of the data. Other techniques may also improve classification for certain data, although this has yet to be conclusively shown.

The challenge for dictionary-based classifiers is to form a qualitative understanding of the type of problems that best suit this approach and to back this understanding with experimental evidence. For example, we could argue that dictionary classifiers will be a good choice of algorithm for classifying long EEG series. This seems reasonable, given BOSS is based on frequency of repetition patterns, but we have no evidence that this is actually the case. It will then be much easier to quantify whether further possible refinements based on techniques used in other fields actually improve accuracy on data for which it is sensible to use a dictionary classifier.

## Acknowledgements

This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) [grant number EP/M015807/1] and the Biotechnology and Biological Sciences Research Council (BBSRC) Norwich Research Park Biosciences Doctoral Training Partnership [grant number BB/M011216/1]. The experiments were carried out on the High Performance Computing Cluster supported by the Research and Specialist Computing Support service at the University of East Anglia and using a Titan X Pascal donated by the NVIDIA Corporation.

## References

- [1] A. Bagnall and G. Janacek. A run length transformation for discriminating between auto regressive time series. *Journal of Classification*, 31:154–178, 2014.
- [2] 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 Mining and Knowledge Discovery*, 31(3):606–660, 2017.
- [3] A. Bagnall, J. Lines, J. Hills, and A. Bostrom. Time-series classification with COTE: The collective of transformation-based ensembles. *IEEE Transactions on Knowledge and Data Engineering*, 27:2522–2535, 2015.
- [4] A. Bailly, S. Malinowski, R. Tavenard, L. Chapel, and T. Guyet. *Dense Bag-of-Temporal-SIFT-Words for Time Series Classification*, pages 17–30. 2016.
- [5] M. Baydogan and G. Runger. Time series representation and similarity based on local autopatterns. *Data Mining and Knowledge Discovery*, 30(2):476–509, 2016.

- [6] M. Baydogan, G. Runger, and E. Tuv. A bag-of-features framework to classify time series. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(11):2796–2802, 2013.
- [7] A. Benavoli, G. Corani, and F. Mangili. Should we really use post-hoc tests based on mean-ranks? *Journal of Machine Learning Research*, 17:1–10, 2016.
- [8] A. Bostrom and A. Bagnall. Binary shapelet transform for multiclass time series classification. *Transactions on Large-Scale Data and Knowledge Centered Systems*, 32:24–46, 2017.
- [9] M. Corduas and D. Piccolo. Time series clustering and classification by the autoregressive metric. *Computational Statistics and Data Analysis*, 52(4):1860–1872, 2008.
- [10] J. Demšar. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7:1–30, 2006.
- [11] H. Deng, G. Runger, E. Tuv, and M. Vladimir. A time series forest for classification and feature extraction. *Information Sciences*, 239:142–153, 2013.
- [12] H. Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P. Muller. Deep learning for time series classification: a review. *arXiv preprint arXiv:1809.04356*, 2018.
- [13] B. Fulcher and N. Jones. Highly comparative feature-based time-series classification. *IEEE Transactions on Knowledge and Data Engineering*, 26(12):3026–3037, 2014.
- [14] S. García and F. Herrera. An extension on statistical comparisons of classifiers over multiple data sets for all pairwise comparisons. *Journal of Machine Learning Research*, 9:2677–2694, 2008.
- [15] J. Grabocka and L. Schmidt-Thieme. Invariant time-series factorization. *Data Mining and Knowledge Discovery*, 28(5):1455–1479, 2014.
- [16] J. Hills, J. Lines, E. Baranauskas, J. Mapp, and A. Bagnall. Classification of time series by shapelet transformation. *Data Mining and Knowledge Discovery*, 28(4):851–881, 2014.
- [17] R. Kate. Using dynamic time warping distances as features for improved time series classification. *Data Mining and Knowledge Discovery*, 30(2):283–312, 2016.
- [18] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 2169–2178, 2006.

- [19] J. Lin, E. Keogh, L. Wei, and S. Lonardi. Experiencing SAX: a novel symbolic representation of time series. *Data Mining and Knowledge Discovery*, 15(2), 2007.
- [20] J. Lin, R. Khade, and Y. Li. Rotation-invariant similarity in time series using bag-of-patterns representation. *Journal of Intelligent Information Systems*, 39(2):287–315, 2012.
- [21] J. Lines and A. Bagnall. Time series classification with ensembles of elastic distance measures. *Data Mining and Knowledge Discovery*, 29:565–592, 2015.
- [22] J. Lines, L. Davis, J. Hills, and A. Bagnall. A shapelet transform for time series classification. In *Proc. the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2012.
- [23] J. Lines, S. Taylor, and A. Bagnall. HIVE-COTE: The hierarchical vote collective of transformation-based ensembles for time series classification. In *Proc. IEEE International Conference on Data Mining*, 2016.
- [24] J. Lines, S. Taylor, and A. Bagnall. Time series classification with HIVE-COTE: The hierarchical vote collective of transformation-based ensembles. *ACM Trans. Knowledge Discovery from Data*, 12(5), 2018.
- [25] D. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [26] P. Schäfer. The BOSS is concerned with time series classification in the presence of noise. *Data Mining and Knowledge Discovery*, 29(6):1505–1530, 2015.
- [27] P. Schäfer. Scalable time series classification. *Data Mining and Knowledge Discovery*, 30(5):1273–1298, 2016.
- [28] P. Senin and S. Malinchik. SAX-VSM: interpretable time series classification using sax and vector space model. In *Proc. 13th IEEE International Conference on Data Mining (ICDM)*, 2013.
- [29] L. Ye and E. Keogh. Time series shapelets: a novel technique that allows accurate, interpretable and fast classification. *Data Mining and Knowledge Discovery*, 22(1-2):149–182, 2011.
- [30] J. Zhao and L. Itti. Classifying time series using local descriptors with hybrid sampling. *KDE*, 28(3):623–637, 2016.