

On the Usage and Performance of The Hierarchical Vote Collective of Transformation-based Ensembles version 1.0 (HIVE-COTE v1.0)

Anthony Bagnall, Michael Flynn, James Large, Jason Lines, and Matthew Middlehurst

University of East Anglia, Norwich, UK., ajb@uea.ac.uk

Abstract. The Hierarchical Vote Collective of Transformation-based Ensembles (HIVE-COTE) is a heterogeneous meta ensemble for time series classification. Since it was first proposed in 2016, the algorithm has undergone some minor changes and there is now a configurable, scalable and easy to use version available in two open source repositories. We present an overview of the latest stable HIVE-COTE, version 1.0, and describe how it differs to the original. We provide a walkthrough guide of how to use the classifier, and conduct extensive experimental evaluation of its predictive performance and resource usage. We compare the performance of HIVE-COTE to three recently proposed algorithms.

Keywords: Time series · Classification · Heterogeneous ensembles · HIVE-COTE

1 Introduction

The Hierarchical Vote Collective of Transformation-based Ensembles (HIVE-COTE) is a heterogeneous meta ensemble for time series classification. The key principle behind HIVE-COTE is that time series classification (TSC) problems are best approached by careful consideration of the data representation, and that with no expert knowledge to the contrary, the most accurate algorithm design is to ensemble classifiers built on different representations.

HIVE-COTE was first described in 2016 [14,15]. At the time, HIVE-COTE was significantly more accurate on average than other known approaches [3] on the 85 datasets that were then the complete UCR archive [6]. HIVE-COTE was an improvement over the 2015 version, called just COTE on publication [4] but later renamed Flat-COTE to differentiate it from its successor. Flat-COTE is a standard ensemble of a range of classifiers built on different representations. It was itself a natural extension of the Elastic Ensemble [13] which only contains nearest neighbour classifiers using different distance measures. HIVE-COTE takes a more structured approach than Flat-COTE. The original HIVE-COTE, which we will henceforth refer to as HIVE-COTE alpha, contained the following classification modules: Elastic Ensemble (EE) [13]; Shapelet Transform

Classifier (STC) [11]; Time Series Forest (TSF) [8]; and Bag of Symbolic-Fourier Approximation Symbols (BOSS) [19]. Each module is encapsulated and built on the train data independently of the others. For new data, each module passes an estimate of class probabilities to the control unit, which combines them to form a single prediction. It does this by weighting the probabilities of each module by an estimate of its testing accuracy formed from the training data.

Our goal with HIVE-COTE alpha was to achieve the highest level of accuracy without concern for the computational resources. This has led to the perception that HIVE-COTE is very slow and does not scale well. Whilst this is true if used in its basic form, it is in fact very simple to restructure HIVE-COTE so it achieves the same level of accuracy in orders of magnitude less time. We have made many small changes to HIVE-COTE with the goal of making it scalable and more useful. We describe these improvements and encapsulate them as HIVE-COTE version 1.0. The changes in HIVE-COTE are both algorithmic and engineering in nature.

The two slowest components of HIVE-COTE alpha are STC (in its old format) and EE. STC used to conduct a full enumeration of all possible shapelets. We have found that this enormous computational effort is not only unnecessary, but often results in over fitting. EE requires cross validation of numerous nearest neighbour classifiers and is very slow on training and testing. EE resulted from a comparative study of nearest neighbour distance measures. Our hypothesis was there was no significant difference between the numerous distance measures and dynamic time warping when used in nearest neighbour classifiers, which is true. We only ensemble as an afterthought. We were surprised to see significant improvement. Its design was necessarily ad hoc to avoid over fitting. We have found that dropping EE all together does not make HIVE-COTE much worse. The main changes are:

1. STC no longer fully enumerates the shapelet space.
2. EE is dropped altogether from HIVE-COTE.
3. The STC, BOSS and RISE components include revisions to improve efficiency.
4. All components and HIVE-COTE 1.0 itself are now contractable (you can set a run time limit), checkpointable (you can save a version to continue building later) and tuneable (select parameters based on train set performance).
5. HIVE-COTE 1.0 can be threaded, built from existing results and easily configured.

The aim of this report is to describe the changes, showcase the usage of HIVE-COTE and to present some new benchmark results that should be used in all future experiments, and to demonstrate the scalability of HIVE-COTE.

2 HIVE-COTE 1.0 Design

Figure 1 provides an overview of the HIVE-COTE structure. The top level ensemble structure and the implementation of each component are described below.

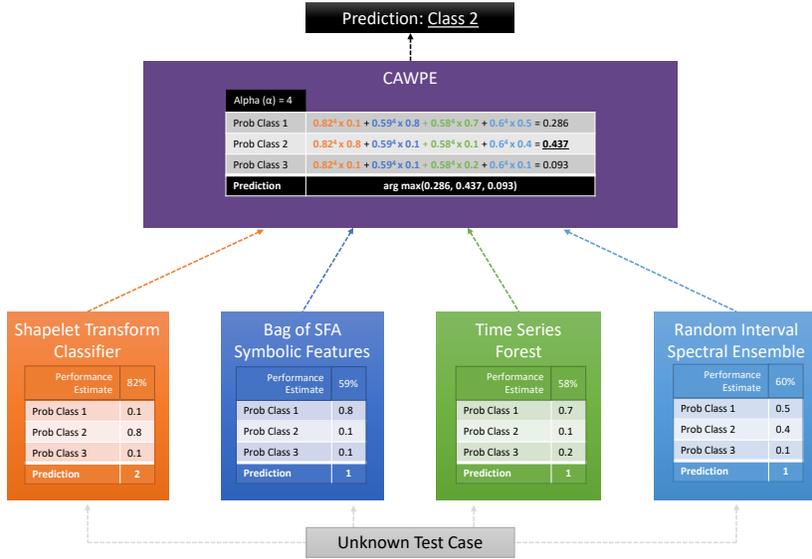


Fig. 1. An overview of the ensemble structure of HIVE-COTE 1.0. Each module produces an estimate of the probability of membership of each class. The control unit (CAWPE) combines these probabilities, weighted by an estimate of the quality of the module found on the train data.

2.1 Ensemble Structure

HIVE-COTE adopts the Cross-validation Accuracy Weighted Probabilistic Ensemble (CAWPE) ensemble structure [12], summarised in Algorithm 1. CAWPE

Algorithm 1 HIVE-COTE classify(A test case \mathbf{x})

Return: prediction for case \mathbf{x}

Parameters: A set of classifiers $\langle M_1, \dots, M_k \rangle$, an exponent α , a set of weights w_i , and the number of classes c

- 1: $\{\hat{p}_1, \dots, \hat{p}_c\} = \{0, \dots, 0\}$
 - 2: **for** $i \leftarrow 1$ to k **do**
 - 3: **for** $j \leftarrow 1$ to c **do**
 - 4: $\hat{q}_j \leftarrow \hat{p}((y = j) | M_i, \mathbf{x})$
 - 5: $\hat{p}_j \leftarrow \hat{p}_j + w_i^\alpha \cdot \hat{q}_j$
 - 6: **return** $\arg \max_{j=1 \dots c} \hat{p}_j$
-

uses an estimate of the accuracy of each classifier to weight the probability estimates of each component. It constructs a tilted distribution through exponentiation (using α) to extenuate differences in classifiers. The weight for each component is found either through ten fold cross validation, or, if the classifier has the capability to estimate its own performance, internally.

2.2 Time Series Forest (TSF)

TSF [8] aims to capture basic summary features from intervals of a time series. For any given time series of length m there are $m(m-1)/2$ possible intervals that can be extracted. TSF takes a random forest-like approach to sampling these intervals. A formal description of TSF is provided in Algorithm 2.

For each tree, r intervals are randomly selected (lines 5-7), each with a random start position and length. Each interval is summarised by the mean, standard deviation and slope (lines 8-11), and the summaries of each interval are concatenated into a single feature vector of length $3r$ for each time series. A decision tree is built on this concatenated feature vector (line 12). New cases are classified using a majority vote of all trees in the forest. The version of TSF used in the bake off [3] employed the random tree used by random forest. However, the decision tree described in [8] has some minor differences to the random tree. It makes no difference in terms of accuracy, but the tree from [8], the time series tree, has advantages in terms of interpretability. Hence, the current version of TSF uses the original version, which we call TimeSeriesTree. Contracting is enforced by the method timeRemaining. TSF simply builds until it has the required number of trees or the time runs out.

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

Parameters: the number of trees, k ; the minimum interval length, p ; the number of intervals per tree, r . (default $k \leftarrow 500$, $p \leftarrow 3$, and $r \leftarrow \sqrt{m}$)

```

1: Let  $\mathbf{F} = (\mathbf{F}_1 \dots \mathbf{F}_k)$  be the trees in the forest
2:  $i \leftarrow 1$ 
3: while  $i < k$  and timeRemaining() do
4:   Let  $\mathbf{S}$  be a list of  $n$  cases ( $\mathbf{s}_1 \dots \mathbf{s}_n$ ) with  $3r$  attributes
5:   for  $j \leftarrow 1$  to  $r$  do
6:      $b \leftarrow \text{randBetween}(1, m - p)$ 
7:      $e \leftarrow \text{randBetween}(b + p, m)$ 
8:     for  $t \leftarrow 1$  to  $n$  do
9:        $s_{t,3(j-1)+1} \leftarrow \text{mean}(\mathbf{x}_t, b, e)$ 
10:       $s_{t,3(j-1)+2} \leftarrow \text{standardDeviation}(\mathbf{x}_t, \mathbf{b}, \mathbf{e})$ 
11:       $s_{t,3(j-1)+3} \leftarrow \text{slope}(\mathbf{x}_t, b, e)$ 
12:    $\mathbf{F}_i.\text{buildTimeSeriesTree}(\mathbf{S}, \mathbf{y})$ 

```

2.3 Random Interval Spectral Ensemble (RISE)

Like TSF, RISE [15] is a tree based interval ensemble. Unlike TSF, it uses a single interval for each tree, and it uses spectral features rather than summary statistics. RISE was recently updated to be faster and contractable [10]. During the build process, summarised in Algorithm 3, a single random interval is selected for each tree. The first tree is a special case in which the whole series is used (lines 5 and 6). Otherwise, an interval with length that is a power of 2 (line

9) is chosen. The same interval for each series is then transformed using the Fast Fourier Transform (FFT) and Auto Correlation Function (ACF). This is a change on the original RISE which also used the partial autocorrelation function (PACF) and autoregressive (AR) model features. Restriction to the Fourier and ACF coefficients does not decrease accuracy, but makes the algorithm much faster. The power spectrum coefficients are concatenated with the first 100 ACF coefficients to form a new training set. In the `tsml` implementation of RISE the base classifier used is the RandomTree classifier used by random forest (line 13). In the test process class probabilities are assigned as a proportion of base classifier votes.

RISE controls the contract run time by creating an adaptive model of the time to build a single tree (lines 4 and 14). This is important for long series (such as audio), where very large intervals can mean very few trees. Details are in [10].

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

Parameters: the number of trees, k ; the minimum interval length, p . (default $k \leftarrow 500$, $p \leftarrow \min(16, m/2)$)

```

1: Let  $\mathbf{F} \leftarrow \langle \mathbf{F}_1 \dots \mathbf{F}_k \rangle$  be the trees in the forest.
2:  $i \leftarrow 1$ 
3: while  $i < k$  and timeRemaining() do
4:   buildAdaptiveTimingModel()
5:   if  $i = 1$  then
6:      $r \leftarrow m$ 
7:   else
8:      $max \leftarrow \text{findMaxIntervalLength}()$ 
9:      $r \leftarrow \text{findPowerOf2Interval}(p, max)$ 
10:   $b \leftarrow \text{randBetween}(1, m - r)$ 
11:   $\mathbf{T}' \leftarrow \text{removeAttributesOutsideOfRange}(\mathbf{T}, \mathbf{b}, \mathbf{r})$ 
12:   $\mathbf{S} \leftarrow \text{getSpectralFeatures}(\mathbf{T}')$ 
13:   $\mathbf{F}_i.\text{buildRandomTreeClassifier}(\mathbf{S}, \mathbf{y})$ 
14:  updateAdaptiveModel( $r$ )
15:   $i \leftarrow i + 1$ 

```

2.4 Bag of SFA Symbolic Features (BOSS)

Dictionary based classifiers convert real valued time series into a sequence of discrete symbol words, then base classification on these words. Commonly, a sliding window of length w is run across a series. For each window, the real valued series of length w is converted through approximation and discretisation processes into a symbolic string of length l (referred to as a word), which consists of α possible letters. The occurrence in a series of each word from the dictionary defined by l and α is counted and, once the sliding window has completed, the series is transformed into a histogram. Classification is based on the histograms

of the words extracted from the series, rather than the raw data. The Bag of Symbolic-Fourier-Approximation Symbols (BOSS) [19] ensemble was found to be the most accurate dictionary based classifier in the bake off [3]. Hence, it forms our benchmark for new dictionary based approaches.

Algorithm 4 gives a formal description of the bag forming process of an individual BOSS classifier. Windows may or may not be normalised (lines 6 and 7). Words are created using Symbolic Fourier Approximation (SFA) (lines 8-13). SFA first finds the Fourier transform of the window (line 8), ignoring the first term if normalisation occurs (lines 9-12). It then discretises the first l Fourier terms into α symbols to form a word in the method SFALookup, using a bespoke supervised discretisation algorithm called Multiple Coefficient Binning (MCB) (line 13). Lines 14-16 implement the process of not counting self similar words: if two consecutive windows produce the same word, the second occurrence is ignored. This is to avoid a slow-changing pattern relative to the window size being over-represented in the resulting histogram.

BOSS uses a non-symmetric distance function in conjunction with a nearest neighbour classifier. Only the words contained in the test instance’s histogram (i.e. the word count is above zero) are used in the distance calculation, but it is otherwise the Euclidean distance.

Algorithm 4 baseBOSS(A list of n time series of length m , $\mathbf{T} = (\mathbf{X}, \mathbf{y})$)

Parameters: the word length l , the alphabet size α , the window length w , normalisation parameter z

```

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: for  $i \leftarrow 1$  to  $n$  do
4:   for  $j \leftarrow 1$  to  $m - w + 1$  do
5:      $\mathbf{s} \leftarrow x_{i,j} \dots x_{i,j+w-1}$ 
6:     if  $z$  then
7:        $s \leftarrow \text{normalise}(s)$ 
8:      $\mathbf{q} \leftarrow \text{DFT}(\mathbf{s}, l, \alpha, p)$  {  $\mathbf{q}$  is a vector of the complex DFT coefficients}
9:     if  $z$  then
10:       $\mathbf{q}' \leftarrow (q_2 \dots q_{l/2+1})$ 
11:     else
12:       $\mathbf{q}' \leftarrow (q_1 \dots q_{l/2})$ 
13:      $\mathbf{r} \leftarrow \text{SFALookup}(\mathbf{q}', \mathbf{B})$ 
14:     if  $\mathbf{r} \neq \mathbf{p}$  then
15:        $pos \leftarrow \text{index}(\mathbf{r})$ 
16:        $h_{i,pos} \leftarrow h_{i,pos} + 1$ 
17:      $\mathbf{p} \leftarrow \mathbf{r}$ 

```

The final classifier is an ensemble of individual BOSS classifiers. The original BOSS ensemble built all models over a pre-defined parameter space for w , l , z and α and retained all base classifiers with accuracy of 92% or higher of the best. This introduces instability in memory usage and carries a time overhead. HIVE-

COTE 1.0 uses contractable BOSS (cBOSS) [17] as its dictionary based classifier. cBOSS changes the method used by BOSS to form its ensemble to improve efficiency and allow for a number of usability improvements. cBOSS was shown to be an order of magnitude faster than BOSS on both small and large datasets from the UCR archive while showing no significant difference in accuracy [17]. It randomly samples the parameter space without replacement (line 8), subsamples the data for each base classifier (line 10), and retains a fixed number of base classifiers. An exponential weighting scheme based on train accuracy, such as the one used in HIVE-COTE, is introduced for ensemble members.

Algorithm 5 cBOSS(A list of n cases length m , $\mathbf{T} = (\mathbf{X}, \mathbf{y})$)

Parameters: the maximum number of base classifiers, k ; the number of parameter samples, s ; the proportion of cases to sample, p . (default $k \leftarrow 50$; $s \leftarrow 250$; $p \leftarrow 0.7$)

```

1: Let  $w$  be window length,  $l$  be word length,  $z$  be normalise/not normalise and  $\alpha$  be
   alphabet size.
2: Let  $\mathbf{B} \leftarrow \langle \mathbf{B}_1 \dots \mathbf{B}_k \rangle$  be a list of  $k$  BOSS classifiers
3: Let  $\mathbf{W} \leftarrow \langle w_1, \dots, w_k \rangle$  be a list of classifier weights
4: Let  $\mathbf{R}$  be a set of possible BOSS parameter combinations
5:  $i \leftarrow 0$ 
6:  $minAcc \leftarrow \infty$ 
7: while  $i < s$  and  $|\mathbf{R}| > 0$  and timeRemaining() do
8:    $[l, a, w, p] \leftarrow \text{randomSampleParameters}(\mathbf{R})$ 
9:    $\mathbf{R} = \mathbf{R} \setminus \{[l, a, w, p]\}$ 
10:   $\mathbf{T}' \leftarrow \text{subsampleData}(\mathbf{T})$ 
11:   $\mathbf{cls} \leftarrow \text{baseBOSS}(\mathbf{T}', l, a, w, z)$ 
12:   $acc \leftarrow \text{estimateAccuracy}(\mathbf{T}', \mathbf{cls})$  { estimate accuracy on train data}
13:  if  $i < k$  then
14:     $\mathbf{B}_i \leftarrow \mathbf{cls}$ ,  $w_i \leftarrow acc^4$ 
15:    if  $acc < minAcc$  then
16:       $minAcc \leftarrow acc$ ,  $idx \leftarrow i$ 
17:    else if  $acc > min\_acc$  then
18:       $\mathbf{B}_{idx} \leftarrow \mathbf{cls}$ ,  $w_{idx} \leftarrow acc^4$ 
19:     $[minAcc, idx] \leftarrow \text{findLowestAcc}(\mathbf{B})$ 
20:   $i \leftarrow i + 1$ 

```

2.5 Shapelet Transform Classifier (STC)

There are two significant changes to the STC used in HC 1.0. Firstly, it only fully enumerates the shapelet space when there is sufficient time to do so. Secondly, it uses a Rotation Forest classifier [18] rather than a heterogeneous ensemble. The shapelet transform is highly configurable: it can use a range of sampling/search techniques in addition to alternative quality measures. We present the default settings and direct the interested reader to the code. The amount of time for the

shapelet search is now a parameter. The algorithm calculates how many possible shapelets there are in a data set, then estimates how many shapelets it can sample from each series. After searching, it updates its timing model using simple reinforcement learning. These operations are encapsulated in operations `estimateNumberOfShapelets` (line 7) and `updateTimingModel` (line 11). Shapelets are randomly sampled in the method `sampleShapelets` (line 8). If the algorithm is allowed more shapelets than the series contains, it evaluates them all. We have experimented with a range of alternative neighbourhood search algorithms, but nothing is much better than random search. Once the shapelets are generated, they are evaluated using information gain (line 9). We use a one vs all evaluation for multi-class problems [5]. Overlapping shapelets are removed in line 10, before the candidates are merged into the overall pool, with the weakest members of the population being deleted. Once the search is complete, the transform is performed (line 13) and the classifier constructed (line 14).

Algorithm 6 STC(A list of n cases length m , $\mathbf{T} = (\mathbf{X}, \mathbf{y})$)

Parameters: the maximum number of shapelets to keep, k ; the shapelet search time, t . (default $k \leftarrow 1000$, $t \leftarrow 1$ hour).

- 1: Let \mathbf{S} be a list of up to k shapelets
- 2: Let \mathbf{R} be a rotation forest classifier.
- 3: $i \leftarrow 0$
- 4: $\text{minIG} \leftarrow 0$
- 5:
- 6: **while** `shapeletTimeRemaining(t)` **do**
- 7: $p \leftarrow \text{estimateNumberOfShapelets}(t, m, n)$
- 8: $\mathbf{S}' \leftarrow \text{sampleShapelets}(\mathbf{x}_i, p)$
- 9: $\mathbf{s} \leftarrow \text{evaluateShapelets}(\mathbf{S}', \mathbf{T})$
- 10: $\mathbf{S}' \leftarrow \text{removeSelfSimilar}(\mathbf{S}', \mathbf{s})$
- 11: `updateTimingModel()`
- 12: $\mathbf{S} \leftarrow \text{merge}(\mathbf{S}, \mathbf{S}')$
- 13: $\mathbf{X}' \leftarrow \text{shapeletTransform}(X, \mathbf{S})$
- 14: $\mathbf{R}.\text{buildRotationForest}(\mathbf{X}', \mathbf{y})$

3 HIVE-COTE 1.0 Usability

We help maintain two toolkits that include time series classification functionality. `sktime`¹ is an open source, Python based, sklearn compatible toolkit for time series analysis. `sktime` is designed to provide a unifying API for a range of time series tasks such as annotation, prediction and forecasting (see [16] for a description of the overarching design of `sktime` and [1] for an experimental comparison of some of the classification algorithms available). The Java toolkit

¹ <https://github.com/alan-turing-institute/sktime>

for time series machine learning, `tsml`², is Weka compatible and is the descendent of the codebase used to perform the bake off. The two toolkits will eventually converge to include all the features described here. Experiments reported in this paper are conducted with `tsml`, as it has more functionality.

3.1 Java Implementation of HIVE-COTE 1.0 in `tsml`

The `HIVE_COTE` class is in the package `tsml.classifiers.hybrids` and can be used as any other Weka classifier. The default configuration is that described in this paper. The code described here is all available in the class `EX07_HIVE_COTE_Examples` with more detail and comments. A basic build is described in Listing 1.1. It cannot handle missing values, unequal length series or multivariate problems yet.

```

1  HIVE_COTE hc = new HIVE_COTE();
2  //this setup called in default constructor in April 2020
3  hc.setupHIVE_COTE_1_0();
4  Instances[] trainTest =
5      DatasetLoading.sampleItalyPowerDemand(0);
6  hc.buildClassifier(trainTest[0]);

```

Listing 1.1. A most basic way using HIVE-COTE 1.0 in `tsml`

We rarely build the classifier in this way. Instead, we build the components and post process the meta ensemble. This is most easily done using our `Experiments` class, which formats the output in a standard way. An example code snippet is in Listing 1.2. Details on optional input flags not given below can be found in the code.

```

1  String[] settings=new String[6];
2  //Where to get data
3  settings[0]="-dp=src/main/java/experiments/data/tsc/";
4  //Where to write results
5  settings[1]="-rp=Temp/";
6  //Whether to generate train files or not
7  settings[2]="-gtf=true";
8  //Classifier name: See ClassifierLists for valid options
9  settings[3]="-cn=TSF";
10 //Problem file
11 settings[4]="-dn=Chinatown";
12 //Resample number: 1 gives the default train/test split
13 settings[5]="-f=1";
14 Experiments.ExperimentalArguments expSettings =
15     new Experiments.ExperimentalArguments(settings);
16 Experiments.setupAndRunExperiment(expSettings);

```

Listing 1.2. Using `Experiments.java` to build a single component.

² <https://github.com/uea-machine-learning/tsml>

HIVE_COTE can read component results directly from file using syntax of the form given in Listing 1.3. It will look in the directory structure created by Experiments. Currently, this file loading method requires all the classifier results to be present in order to build.

```

1  HIVE_COTE hc=new HIVE_COTE();
2  hc.setBuildIndividualsFromResultsFiles(true);
3  hc.setResultsFileLocationParameters("C:/Temp/", "
    Chinatown", 0);
4  String[] components={"TSF","RISE","cBOSS","STC"};
5  hc.setClassifiersNamesForFileRead(components);

```

Listing 1.3. Buildind HIVE-COTE from existing results files

HIVE_COTE is configurable for different components, threadable (see Listing 1.4) and contractable (see Listing 1.5). In sequential mode, it simply divides the time equally between components. When threaded, it gives the full contract time to each component. It does not yet thread individual components; it is on our development list. You can set the maximum build time for HIVE_COTE if the components all implement the `TrainTimeContractable` interface.

```

1  HIVE_COTE hc = new HIVE_COTE();
2  EnhancedAbstractClassifier[] c=new
    EnhancedAbstractClassifier[2];
3  c[0]=new RISE();
4  c[1]=new TSF();
5  String[] names={"RISE","TSF"};
6  hc.setClassifiers(c,names,null);
7  hc.enableMultiThreading(2);

```

Listing 1.4. Threaded build of HIVE-COTE with bespoke classifiers

```

1  //Ways of setting the contract time
2  HIVE_COTE hc = new HIVE_COTE();
3  //Minute, hour or day limit
4  hc.setMinuteLimit(10);
5  //Specify units
6  hc.setTrainTimeLimit(30, TimeUnit.HOURS);
7  //Or just give it in nanoseconds
8  hc.setTrainTimeLimit(10000000000L);

```

Listing 1.5. Contracting HIVE-COTE for a rom existing results files

Finally, HIVE_COTE is tuneable. Our method of implementing tuned classifiers is to wrap the base classifier in a `TunedClassifier` object which interacts through the method `setOptions`. An example of tuning the α parameter is given in Figure 1.6. We advise tuning using results loaded from file. We have found tuning α makes no significant difference. We have not finished evaluating tuning which components to use.

```

1  HIVE_COTE hc=new HIVE_COTE();
2  hc.setBuildIndividualsFromResultsFiles(true);

```

```

3 hc.setResultsFileLocationParameters(resultsPath, dataset,
4 fold);
5 hc.setClassifiersNamesForFileRead(cls);
6 TunedClassifier tuner=new TunedClassifier();
7 tuner.setClassifier(hc);
8 ParameterSpace pc=new ParameterSpace();
9 double[] alphaVals={1,2,3,4,5,6,7,8,9,10};
10 pc.addParameter("A",alphaVals);
tuner.setParameterSpace(pc);

```

Listing 1.6. Tuning HIVE-COTE α parameter from existing results files

3.2 Python Implementation of HIVE-COTE 1.0 in `sktime`

A Python implementation of HIVE-COTE is under development and available in `sktime`. As previously discussed, this version of the algorithm is less mature in terms of its development and the results reported in this paper are from the Java version of the algorithm. The `sktime` implementation is in an alpha state, and will eventually converge on the same functionally as the more developed Java implementation, but currently has a number of limitations (such as building from file and running constituents in parallel). Further, the Python implementations of the constituent classifiers are less efficient than the Java implementations, and as such, HIVE-COTE 1.0 in `sktime` is slower than the `tsml` implementation on equivalent inputs.

The interface and basic usage of HIVE-COTE in `sktime` is very similar to that of `tsml`. The terminology is slightly different however as the `sktime` version uses `fit` and `predict` derived from `scikit-learn` while the Java version uses `build` and `classify` from `Weka/tsml`. Notionally the process of constructing and making predictions with HIVE-COTE are equivalent and a simple example of fitting and predicting with HIVE-COTE in `sktime` is given in Listing 1.7. The most up-to-date implementation of HIVE-COTE can be found in the `sktime` toolkit on GitHub under the `hive-cote` branch ³.

```

1 def basic_hive_cote(data_dir, dataset_name):
2     # using the default constructor for the HIVE-COTE class
3     hc = HIVE-COTE()
4
5     # loading training data
6     train_x, train_y = load_data(
7         data_dir + dataset_name + "_TRAIN.ts")
8
9     # building HIVE-COTE 1.0 sequentially
10    hc.fit(train_x, train_y)
11
12    # loading testing data

```

³ https://github.com/alan-turing-institute/sktime/blob/hive_cote/sktime/contrib/meta/ensembles.py

```

13 test_x, test_y = load_data(
14     data_dir + dataset_name + "_TEST.ts")
15
16 # predict class values of the test data
17 preds = hc.predict(train_x)
18
19 # calculate the test accuracy
20 acc = accuracy_score(self.train_y, preds)

```

Listing 1.7. A simple example of using HIVE-COTE 1.0 in `sktime`

4 Performance

To measure performance of the new HIVE-COTE, we evaluate each component and the algorithm itself on 112 of the 128 UCR archive datasets. These 112 datasets are all equal length and have no missing values. Figure 2 shows the critical difference diagram for HIVE-COTE 1.0 (HC 1.0) and its four components. This broadly mirrors the performance presented in [15]. We have compared the

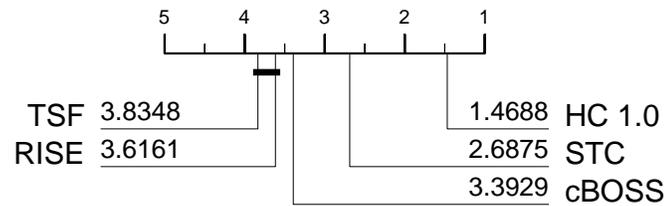


Fig. 2. Critical difference diagram for HIVE-COTE 1.0 and its four components on 112 UCR TSC problems. Full results are available from www.timeseriesclassification/results.php

results of the four components to the original results and found there is no significant difference. Comparison of HIVE-COTE alpha version and 1.0 identify a small, but significant, difference. Removing EE makes HIVE-COTE worse on 48 and better on 33 of the 85 datasets used in the original experiments. The mean reduction in accuracy is 0.6%. The differences in accuracy on specific problems identify those where a distance based approach may be the best. MedicalImages, SonyAIBORobotSurface1, WordSynonyms and Lightning7 were all more than 5% less accurate with EE removed. We are trading this small loss in average accuracy for significant gains in run time and reduction in memory usage. We have explored ways of making EE more efficient, but as yet none of these approaches have met the criteria for maintaining accuracy and providing sufficient speed up. Since HIVE-COTE alpha was proposed, three new algorithms have achieved equivalent accuracy. TS-CHIEF [20] is a tree ensemble that embeds dictionary, spectral and distance based representations. InceptionTime [9] is a

deep learning ensemble, combining five homogeneous networks each with random weight initialisations for stability. ROCKET [7] uses a large number (10,000 by default) of randomly parameterised convolution kernels in conjunction with a linear ridge regression classifier. We use the configurations of each classifier described in their respective publications. Figure 3 shows the critical difference diagram for these three classifiers and HC 1.0. There is no significant difference between any of them. The differences between HIVE-COTE and the other three

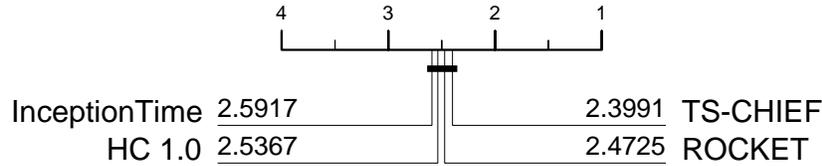


Fig. 3. Critical difference diagram for current state of the art on 109 UCR TSC problems.

are summarised in Table 1. TS-CHIEF is the most similar to HIVE-COTE, with an average accuracy just 0.25% lower and a high correlation between accuracies. ROCKET has a high variation in performance, as reflected in the high standard deviation of differences and the difference between the mean and median of the differences. It also has the lowest correlation to HC 1.0.

Table 1. Summary of differences in accuracy between HIVE-COTE and the other three algorithms. A negative difference indicates HIVE-COTE is more accurate.

	Mean	Median	Std Dev of differences	Correlation
TS-CHIEF	-0.25%	0.00%	3.801	95.54%
InceptionTime	-0.65%	0.00%	5.82	89.46%
ROCKET	-1.41%	0.05%	8.64	80.28%

Table 2 shows the results for problems with the 10 biggest differences between HIVE-COTE and ROCKET. ROCKET does very poorly on some problems (hence the large average difference in Table 1). InceptionTime also shares this characteristic of occasionally simply failing on a problem. It is not obvious why this happens to both ROCKET and InceptionTime. It may be the result of over fitting. It is worthwhile considering the run time of these algorithms. However, comparisons are made more difficult because of the different software. HIVE-COTE and TS-CHIEF are both built using `tsml`, so are directly comparable. Table 3 summarises the time taken to train the classifiers on 109 UCR problems. Three problems (`HandOutlines`, `NonInvasiveFetalECGThorax1` and

Table 2. Accuracy for ten problems with the biggest difference between HIVE-COTE and ROCKET (five negative and five positive differences).

	TS-CHIEF	ROCKET	HC 1.0	InceptionTime
PigAirwayPressure	96.01%	19.55%	95.77%	92.21%
SemgHandMovementCh2	88.50%	65.26%	88.90%	55.10%
EthanolLevel	60.56%	62.53%	84.90%	87.55%
CinCECGTorso	95.34%	86.41%	99.37%	83.28%
ScreenType	59.42%	60.90%	72.42%	70.56%
FiftyWords	84.27%	82.51%	77.16%	82.68%
ChlorineConcentration	66.08%	79.61%	73.39%	86.36%
MedicalImages	79.91%	80.51%	74.04%	79.63%
WordSynonyms	79.37%	76.44%	69.32%	75.18%
SonyAIBORobotSurface1	88.97%	95.81%	82.63%	95.42%

NonInvasiveFetalECGThorax2) are omitted because TS-CHIEF did not complete within 7 days (the job limit on our cluster). Of the HC-1.0 components, STC is by far the slowest. This is caused by the classifier, Rotation Forest, not the transform, which is contracted to take at most one hour. The STC design principle is to choose a large number of shapelets (up to 1000) and let the classifier sort out their importance. Rotation forest is the best approach for problems with all real valued attribute [2], and we have developed a contracted version that can, if necessary, be used to control the build time (see Listing 1.6). Most of the extra computation required by HIVE-COTE is in forming the estimates of the accuracy on the train data. TS-CHIEF is the slowest algorithm, but it is approximately the same as HIVE-COTE.

Table 3. Time in hours to train a classifier for 112 of the UCR problems on a single core.

	TSF	cBOSS	RISE	STC	HC 1.0	TS-CHIEF
Mean	0.13	0.11	0.15	2.30	4.26	5.75
Total	14.15	12.31	16.05	251.12	464.49	626.33

Test time can be a factor for deploying classifiers in time critical situations. Table 4 summarises the time (in minutes) taken to predict the test cases. STC is the slowest component when testing, again caused by the classifier not the transform. TS-CHIEF is the slowest in testing. We do not have reliable times for ROCKET, which is run in `sktime`. It is undoubtedly faster than HIVE-COTE and TS-CHIEF though. We can also measure maximum memory usage of the classifiers, as this is also often a serious bottleneck for scalability. Table 5 summarises the memory usage of the six `tsml` classifiers. The pattern is similar to that observed with run time. TSF, RISE and cBOSS have a light memory footprint. STC, and hence HIVE-COTE, have a larger memory usage. TS-CHIEF is the most memory hungry, and it seemingly does not scale well in terms of

Table 4. Time in minutes to make predictions on the test data for 109 of the UCR problems

	TSF	cBOSS	RISE	STC	HC 1.0	TS-CHIEF
Mean	0.01	0.08	0.07	0.62	0.78	4.09
Total	1.42	8.63	7.09	67.50	84.64	445.55

memory. For, example, on the 11 problems where HIVE-COTE required more than 1 GB, TS-CHIEF required approximately four times the memory. The highest memory usage on the 109 problem it could complete was 18 GB on FordA and FordB. It seems highly likely that the memory usage of TS-CHIEF could be improved without loss of accuracy.

Table 5. Memory usage in MB for 109 UCR problems

	TSF	cBOSS	RISE	STC	HC 1.0	TS-CHIEF
Mean	162	255	263	1,464	1,618	2,004
Max	1061	3432	740	3,533	4,426	18,532

5 Conclusions

The purpose of this report is to present a more practical version of HIVE-COTE and compare it to recent advances in the field of TSC. On average, there is no real difference between InceptionTime, TS-CHIEF, ROCKET and HIVE-COTE in terms of accuracy. ROCKET is undoubtedly the fastest, but it is prone to fail badly on the occasional data set. InceptionTime is slow and requires a GPU. It also fails on some data sets. HIVE-COTE and TS-CHIEF are broadly comparable. HIVE-COTE is presently more configurable and controllable.

We hope the results presented here and on the accompanying website serve to act as a baseline comparison for any new research in the field in the future. We have presented HIVE-COTE as is, without improvements that we have been working on. Our next goal is to find a more accurate version that is comparable in run time and memory usage, and to demonstrate scalability by introducing some much larger problems into the UCR archive. We are also developing versions that can handle unequal length series and multivariate TSC problems.

References

1. A. Bagnall, F. Király, M. Löning, M. Middlehurst, and G. Oastler. A tale of two toolkits, report the first: benchmarking time series classification algorithms for correctness and efficiency. *ArXiv e-prints*, arXiv:1909.05738, 2019.
2. A. Bagnall, A. Bostrom, G. Cawley, M. Flynn, J. Large, and J. Lines. Is rotation forest the best classifier for problems with continuous features? *ArXiv e-prints*, arXiv:1809.06705, 2018.
3. 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.

4. 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.
5. 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.
6. H. Dau, A. Bagnall, K. Kamgar, M. Yeh, Y. Zhu, S. Gharghabi, C. Ratanamahatana, A. Chotirat, and E. Keogh. The ucr time series archive. *IEEE/CAA Journal of Automatica Sinica*, 6(6):1293–1305, 2019.
7. A. Dempster, F. Petitjean, and G. I. Webb. Rocket: Exceptionally fast and accurate time series classification using random convolutional kernels. *arXiv preprint arXiv:1910.13051*, 2019.
8. H. Deng, G. Runger, E. Tuv, and M. Vladimir. A time series forest for classification and feature extraction. *Information Sciences*, 239:142–153, 2013.
9. H. Fawaz, B. Lucas, G. Forestier, C. Pelletier, D. Schmidt, J. Weber, G. Webb, L. Idoumghar, P. Muller, and F. Petitjean. Inceptiontime: Finding alexnet for time series classification. *ArXiv*, 2019.
10. M. Flynn, J. Large, and A. Bagnall. The contract random interval spectral ensemble (c-RISE): The effect of contracting a classifier on accuracy. In *International Conference on Hybrid Artificial Intelligence Systems*, volume 11734 of *Lecture Notes in Computer Science*, pages 381–392. 2019.
11. 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.
12. J. Large, J. Lines, and A. Bagnall. A probabilistic classifier ensemble weighting scheme based on cross validated accuracy estimates. *Data Mining and Knowledge Discovery*, 33(6):1674–1709, 2019.
13. J. Lines and A. Bagnall. Time series classification with ensembles of elastic distance measures. *Data Mining and Knowledge Discovery*, 29:565–592, 2015.
14. J. Lines, S. Taylor, and A. Bagnall. HIVE-COTE: The hierarchical vote collective of transformation-based ensembles for time series classification. In *Proc. 16th IEEE International Conference on Data Mining*, 2016.
15. 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.
16. M. Löning, A. Bagnall, S. Ganesh, V. Kazakov, J. Lines, and F. J. Király. A unified interface for machine learning with time series. *ArXiv e-prints*, arXiv:1909.07872, 2019.
17. M. Middlehurst, W. Vickers, and A. Bagnall. Scalable dictionary classifiers for time series classification. In *International Conference on Intelligent Data Engineering and Automated Learning*, Lecture Notes in Computer Science, pages 11–19. 2019.
18. J. Rodriguez, L. Kuncheva, and C. Alonso. Rotation forest: A new classifier ensemble method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(10):1619–1630, 2006.
19. 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.
20. A. Shifaz, C. Pelletier, F. Petitjean, and G. Webb. TS-CHIEF: A scalable and accurate forest algorithm for time series classification. *ArXiv e-prints*, arXiv:1906.10329, 2019.