

Adversarial Regularization for Explainable-by-Design Time Series Classification

Yichang Wang¹, Rémi Emonet², Elisa Fromont¹, Simon Malinowski¹, and Romain Tavenard³

¹ Univ Rennes, Inria, CNRS, IRISA, Rennes, France

² Laboratoire Hubert Curien UMR 5516, Univ Lyon, Saint-Etienne, France

³ Univ Rennes, CNRS, LETG, IRISA, Rennes, France

Abstract. Time series classification can be successfully tackled by jointly learning a shapelet-based representation of the series in the dataset and classifying the series according to this representation. This shapelet-based classification is both accurate and explainable since the shapelets are time series themselves and thus can be visualized and be provided as a classification explanation. In this paper, we claim that not all shapelets are good visual explanations and we propose a simple, yet also accurate, adversarially regularized EXplainable Convolutional Neural Network, XCNN, that can learn shapelets that are, by design, suited for explanations. We validate our method on the usual univariate time series benchmarks of the UCR repository.

Keywords: Time series · Shapelets · Adversarial networks · Explainable AI · Convolutional Neural Networks

1 Introduction

A time series (TS) Z is a series of time-ordered values, $Z = \{z^{(1)}, z^{(2)}, \dots, z^{(T)}\}$ where $z^{(t)} \in \mathbb{R}^d$, T is the length of our time series and d is the dimension of the feature vector describing each data point. If $d = 1$, Z is said univariate, otherwise it is said multivariate. In this paper, we are interested in classifying univariate time series. We are given a training set $\mathcal{T} = \{(Z_1, y_1), \dots, (Z_n, y_n)\}$, composed of n time series Z_i and their associated labels y_i . Our aim is to learn a function h such that $h(Z_i) = y_i$, in order to predict the labels of new incoming time series. The time series classification problem has been studied in countless applications (see for example [23]) ranging from stock exchange evolution, daily energy consumption, medical sensors, videos, etc.

Many methods have been developed to tackle this problem (see [2] for a review). One very successful category of methods consists in “finding” discriminative phase-independent subsequences, called *shapelets*, that can be used to classify the series. In the first papers about shapelet-based time series classification [25, 19], the shapelets were directly extracted from the training set and the selected shapelets could be used *a posteriori* to explain the classifier’s decision. However, the shapelet enumeration and selection processes were either very costly or the selection was fast but did not yield good performance (as discussed in Section 2).




Fig. 1: The three most discriminative shapelets obtained for the dataset Diatom-SizeReduction using (left column) Learning Shapelets or (right column) our XCNN architecture.

Jointly learning a shapelet-based representation of the series in the dataset and classifying the series according to this representation [15,10] allowed to obtain discriminative shapelets in a much more efficient way. An example of such learned shapelets, obtained with the method from [10], is given in blue in Figure 1 (left). However, if the learned shapelets are definitively discriminative, they are often very different (visually) from actual pieces of a real series in the dataset. As such, these shapelets might not be suited to explain a particular classifier’s decision. Note that the same interpretability issue arises with ensemble classifiers such as [3] where one decision depends on the presence of multiple shapelets. One of the main challenges nowadays is to provide Machine Learning (ML) methods that are both accurate and self-explanatory, i.e. provide mechanisms to explain their decisions to human users since, in many scenarios, it may be risky, unacceptable, or simply illegal, to let artificial intelligent systems make decisions without any human supervision [11].

In this paper, we make use of a simple convolutional network to classify time series and we show how one can leverage the principle of adversarial learning to regularize the parameters of this network such that it learns shapelets that could be more useful to interpret the classifier’s decision. Section 2 presents the most related work. We detail our XCNN method in Section 3. In Section 4, we show quantitative and qualitative results on the usual time series benchmarks [4]: XCNN performance are on par with comparable state-of-the-art methods and our explainable-by-design method provides new types of explanations for neural network’s predictions.

2 Related Work

In this section we review the literature on shapelet-based Time Series Classification (TSC) and on tools for understanding black box model predictions.

2.1 Time Series Classification

Shapelets are discriminative subseries that can either be extracted from a set of time series or learned so as to minimize an objective function. They have been introduced in [25] but in this work, the search space of all possible shapelets is explored exhaustively which makes the method intractable in practice. This high

time complexity has led to the use of heuristics in order to select the shapelets more efficiently. In Fast Shapelets (FS) [19], the authors rely on quantized time series and random projections in order to accelerate the shapelet search but they sacrifice the accuracy, as reported in [2]. The Shapelet Transform (ST) [15] consists in transforming time series into a feature vector whose coordinates represent distances between the time series and the shapelets selected beforehand. However as in [25], the shapelets selection step makes the method unfit for large scale learning.

In order to face the high complexity that comes with search-based methods, other strategies have been designed for shapelet selection. On the one hand, some attention has been paid to random sampling of shapelets from the training set [13]. On the other hand, [10] showed that shapelets could be learned using a gradient-descent-based optimization algorithm. The method, referred to as Learning Shapelets (LS) in the following, jointly learns the shapelets and the parameters of a logistic regression classifier. This makes the method very similar in spirit to a neural network with a single convolutional layer followed by a fully connected classification layer and where the convolution operation is replaced by a sliding-window local distance computation. A min-pooling aggregator should then be used for temporal aggregation.

Closely related to shapelet-based methods (as stated above), variants of Convolutional Neural Networks (CNN) have been introduced for the TSC task [24]. These are mostly mono-dimensional variants of CNN models developed in the Computer Vision field. Note however that most models are rather shallow, which is likely to be related to the moderate sizes of the benchmark datasets present in the UCR/UEA archive [4]. A review of these models can be found in [2].

Finally, ensemble-based methods, such as COTE [3] or HIVE-COTE [16], that rely on several of the above-presented standalone classifiers are now considered state-of-the-art for the TSC task. Note however that these methods tend to be computationally expensive, with high memory usage and difficult to interpret (as stated in Section 1) due to the combination of many different core classifiers.

In this paper, we propose a method that is scalable (compared to methods such as Shapelets [25] or ST [15]), yields interpretable results which can be used to explain the classifier’s decisions (compared to ensemble approaches or unconstrained approaches such as [10] or [16]), and exhibits good classification accuracy (compared to FS [19]).

2.2 Model Interpretability

Among the vast number of existing classifiers, some are considered self-explanatory (e.g. decision trees, classification rules), while others are difficult to interpret (e.g. ensemble methods, neural networks that can be considered as black-boxes). Interpretation of black box classifiers usually consists in designing an interpretation layer between the classifier and the human level. Two criteria refine the category of methods to interpret classifiers: global versus local (i.e. dedicated to one sample) explanations, and black-box dependent versus agnostic. In this category, state-of-the-art methods are Local Interpretable Model-agnostic Explanations

(LIME and Anchors) [20,21] and SHapley Additive exPlanations (SHAP) [17]. SHAP values come with the black-box local estimation advantages of LIME, but also with theoretical guarantees. A higher absolute SHAP value of an attribute compared to another means that it has a higher predictive or discriminative power. However, these methods, contrarily to XCNN, are not able to show what has been learned and is used by the classifier to explain a particular decision.

GradCAM [22] is a popular local visualization method designed to explain neural networks decisions on image classification tasks. It uses gradient-based methods to highlight (with a heat map) the discriminative pixels on a given input test image. This method was adapted in MTEX-CNN [1] as an explanation and feature selection tool for multivariate time series (MTS) classification tasks which is a closer setting to ours. In [1], the authors proposed to stack 2D and 1D convolution sequentially to capture the important feature(s) and the important time stamp(s) for the time series. The prediction results are explained by inspecting the input MTS using GradCAM on both the variable and temporal dimensions.

[6] has a similar goal as ours (to produce interpretable discriminative shapelets) and build on both the work from [15] (in this case the candidate shapelets are extracted with a piecewise aggregate approximation) and from [10] to automatically refine the “handcrafted” shapelets. Contrarily to our method, there is no explicit constraint on the learning process that ensures the interpretability of the shapelets. Besides, their experimental validation makes it hard to fully grasp the benefits and limitations of the proposed method since the algorithm is evaluated on a small subset of UCR/UEA datasets [4] and they provide visualizations for only a couple of the learned shapelets.

The work from [18] is the closest to ours. Contrarily to ours, they decouple the shapelet learning phase and the classification process resulting in a quite different adversarial architecture. Their classification process is made using the shapelet transform method [15] but, in this case, the candidate shapelets are dynamically generated for each input time series. In our case, this is learned by a simple CNN for all the dataset. In [18], an adversarial regularization is also used to constrain the generated shapelets to be similar to real pieces of the series. However, the regularization is imposed on the result of the convolutions (i.e. the feature maps) and not on the convolutions themselves as we propose to do in this paper. This is a different philosophy: we believe that the pattern detectors, i.e. the convolutions, are the shapelets. They believe that the shapelets are the series output by the convolution operation which might, in our opinion, have a very different shape than the original input signal. This difference of regularization may hinder the interpretability of the learned shapelets but this aspect is not studied in details in [18]. Besides, the proposed method does not allow global explanations (in addition to local ones) as can be done with our method. However, according to the results reported in [18], their method is more accurate than ours since it gives better results than LS [10], which gives similar results to our method, as shown in the experiments. The work proposed in [18] thus has a different trade-off explainability/accuracy than us.

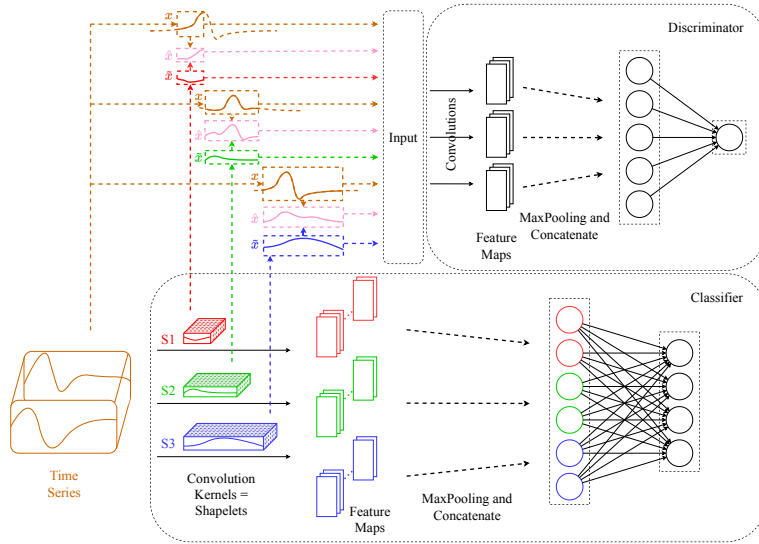


Fig. 2: Adversarial architecture of our proposed explainable CNN (XCNN).

Finally [14] also proposes a time series classification method. The authors propose to extract various symbolic representations from the time series and train a logistic regression model on top of these representations. The logistic regression weights are then inspected (using GradCAM) to extract the most discriminative features and localize the most important time series subparts. This method necessitates to discretize the original signal (and thus lose some information), it is not self-explanatory (the explanations are post-hoc) and we believe that showing the shapelets, as we can do in our method, is an important feature for explaining decisions.

3 TS Classification with Regularized Shapelets

In this section, we present our architecture, XCNN, to learn interpretable discriminative shapelets for time series classification. Our base time series *classifier* is a Convolutional Neural Network (CNN). As explained in Section 2, this model is very similar in spirit to the Learning Shapelet (LS) model presented in [10]. Both LS and CNN slide the shapelets on the series to compute local (dis)similarities. We will loosely refer to the convolution filters of our classifier as *Shapelets* in the following.

Inspired by previous work on adversarial training (e.g. [9]), in addition to our CNN classifier, we make use of an adversarial neural network (the discriminator at the top of Figure 2) to regularize the convolution parameters of our classifier. This regularization acts as a soft constraint for the classifier to learn shapelets as similar to real pieces of the training time series as possible. To obtain the best trade-off between the discriminative power of the shapelets (i.e. the final

classification performance) and their interpretability, our training procedure alternates between training the discriminator and the classifier. The training procedures are explained in the next subsection.

Contrarily to GANs, our adversarial architecture does not rely on a generator to produce fake samples from a latent space. XCNN iteratively modifies the shapelets (i.e. the convolution filters of the classifier) such that they become close to subseries from the training set. The type of data given as input to the discriminator is another major difference between a GAN and XCNN: in a GAN, the discriminator is fed with complete instances, while in XCNN, the discriminator takes subseries as input. These subseries can either be shapelets from the classifier model (denoted as \tilde{x} in Figure 2), portions of training time series (denoted as x) or interpolations between shapelets and training time series portions (\hat{x} , see the following section for more details on those). This process allows the discriminator to alter the shapelets for better interpretability.

3.1 Loss Function

As for GANs, our optimization process alternates between losses attached to the subparts of our model. Here, each training epoch consists of three main steps that are (i) optimizing the classifier parameters for correct classification, (ii) optimizing the discriminator parameters to better distinguish between real subseries and shapelets and (iii) optimizing shapelets to fool the discriminator, so that the regularized-shapelets become similar to a subsequence of time series. Each of these steps is attached to a loss function that we describe in the following.

Firstly, a multi-class cross entropy loss is used for the classifier. It is denoted by $L_c(\theta_c)$ where θ_c is the set of all classifier parameters.

Secondly, our discriminator is trained using a loss function derived from the Wasserstein GANs with Gradient Penalty (WGAN-GP) [12]:

$$L_d(\theta_d) = \mathbb{E}_{\tilde{x} \sim \mathbb{P}_S} [D(\tilde{x})] - \mathbb{E}_{x \sim \mathbb{P}_x} [D(x)] + \lambda \mathbb{E}_{\hat{x} \sim \mathbb{P}_{\hat{x}}} [(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2]$$

where \mathbb{P}_S is the empirical distribution over the shapelets, \mathbb{P}_x is the empirical distribution over the training subseries, and $\hat{x} = \epsilon x + (1 - \epsilon)\tilde{x}$, where ϵ is drawn uniformly at random from the interval $[0, 1]$.

Thirdly, shapelets are updated to fool the discriminator by optimizing on the loss $L_r(\theta_s)$ where $\theta_s \subset \theta_c$ is the set of shapelet coefficients:

$$L_r(\theta_s) = - \mathbb{E}_{\tilde{x} \sim \mathbb{P}_S} [D(\tilde{x})] \tag{1}$$

3.2 Learning Algorithm

Algorithm 1 presents the whole training procedure to update the parameters of our model. At each epoch of this algorithm, the three steps presented above are executed sequentially. Note that in the second step (lines 10–17), sampling classifier shapelets, as well as sampling subseries from the training set, is performed uniformly at random.

Algorithm 1: XCNN

```

Require : number of shapelets  $n_S$ 
Require : random initialization for the classifier/discriminator/shapelets
            $\theta_c, \theta_d, \theta_s \subset \theta_c$ 
Require : gradient penalty coefficient  $\lambda$ 
Require : number of epochs  $n_{\text{epochs}}$ , mini-batch size  $m$ 
Require : number of classifier/discriminator/regularization mini-batches per
           epoch  $n_c, n_d, n_r$ 
Require : optimizer (Adam) hyperparameters  $\alpha, \beta_1, \beta_2$ 
1 for  $i = 1, \dots, n_{\text{epochs}}$  do
2   for  $t = 1, \dots, n_c$  do
3     for  $j = 1, \dots, m$  do
4       Sample a pair  $(Z_j, y_j)$  from the training set
5        $\hat{y}_j \leftarrow h_{\theta_c}(Z_j)$ 
6        $L_c^{(j)} \leftarrow \text{CrossEntropy}(y_j, \hat{y}_j)$ 
7     end
8      $\theta_c \leftarrow \text{Adam}(\nabla_{\theta_c} \frac{1}{m} \sum_{j=1}^m L_c^{(j)}, \theta_c, \alpha, \beta_1, \beta_2)$ 
9   end
10  for  $t = 1, \dots, n_d$  do
11    for  $j = 1, \dots, m$  do
12      Sample a shapelet  $\tilde{x}_j$  from the set  $\theta_s$ , a subseries  $x_j$  from the training
13      set and a random number  $\epsilon \sim U[0, 1]$ 
14       $\hat{x}_j \leftarrow \epsilon x_j + (1 - \epsilon)\tilde{x}_j$ 
15       $L_d^{(j)} \leftarrow D(\tilde{x}_j) - D(x_j) + \lambda(\|\nabla_{\hat{x}} D(\hat{x}_j)\|_2 - 1)^2$ 
16    end
17     $\theta_d \leftarrow \text{Adam}(\nabla_{\theta_d} \frac{1}{m} \sum_{j=1}^m L_d^{(j)}, \theta_d, \alpha, \beta_1, \beta_2)$ 
18  end
19  for  $t = 1, \dots, n_r$  do
20    for  $j = 1, \dots, n_S$  do
21       $\tilde{x}_j \leftarrow \theta_s[j]$ 
22       $L_r^{(j)} \leftarrow -D(\tilde{x}_j)$ 
23    end
24     $\theta_s \leftarrow \text{Adam}(\nabla_{\theta_s} \frac{1}{n_S} \sum_{j=1}^{n_S} L_r^{(j)}, \theta_s, \alpha, \beta_1, \beta_2)$ 
25  end

```

4 Experiments

In this section, we will detail the training procedure for XCNN and present both quantitative and qualitative experimental results.

4.1 Experimental Setting

Competitors We provide experiments about the quality (for explanations) of our learned shapelets as well as their quality for classification. As explained in Section 2, our most relevant competitor is Learning Shapelets (LS) from [10] as it also describes a shapelet-based model where the shapelets are learned and where

a single model is used for classification. The quality (for explanations) of the shapelets produced by [25] and [19] is, by design, perfect since the shapelets are true subpart of the original series so we do not compare with them but only with the shapelets learned by [10]. However, we compare our classification performance to [25], Fast Shapelets [19] and the recent ELIS [6].

Datasets We use the 85 univariate time series datasets from the UCR/UEA repository for which most of our baselines results are already available [4]. Note that our CNN-based method may also be suited for multivariate time series but giving “intuitive” explanations for multivariate data is far from obvious and we decided to focus only on univariate ones in this paper. The datasets are significantly different from one to another, including seven types of data with various number of instances, lengths, and classes. The splits between training and test sets are provided in the repository.

Architecture details and parameter setting We have implemented the XCNN model using TensorFlow following the general architecture illustrated in Figure 2. The classifier is composed of one 1D convolution layer with ReLU activation, followed by a max-pooling layer along the temporal dimension and a fully connected layer with a soft-max activation. The shapelets use a Glorot uniform initializer [8] while the other weights are initialized uniformly (using a fixed range). For each dataset, three different shapelet lengths are considered, inspired by the heuristic from [10] but without resorting to hyper-parameter search: we consider 3 groups of $20 \times n_{\text{classes}}$ shapelets of length $0.2T$, $0.4T$ and $0.6T$, where n_{classes} is the number of classes in the dataset and T is the length of the time series at stake.

The convolution filters of the classifier, i.e. the shapelets, are given as input to the discriminator which has the same structure as the classifier, but with shorter convolution filters (100 filters of size $0.06T$, $0.12T$ and $0.18T$) and a single-neuron *tanh* activation instead of the soft-max in the last layer. For optimization, we use Adam optimizer with a standard parameterization ($\alpha = 10^{-3}$, $\beta_1 = 0.9$ and $\beta_2 = 0.999$) and each epoch consists in $n_c = 15$ (resp. $n_d = 20$ and $n_r = 17$) mini-batches of optimization for the classifier loss (resp. discriminator and regularizer losses).

Experimental results are reported in terms of test accuracy and aggregated over five random initializations. All experiments are run for 8,000 training epochs.

4.2 Qualitative results for explainability

We first describe how we compute the shapelet contributions to the classification of one (or multiple) example(s) and validate that our adversarial regularization actually ensures that shapelets are visually similar to the training data. We believe that the Euclidean distance is the most understandable distance for human eyes so all the figures that show shapelets and series will be displayed using this distance even though it is not the one optimized during XCNN training.

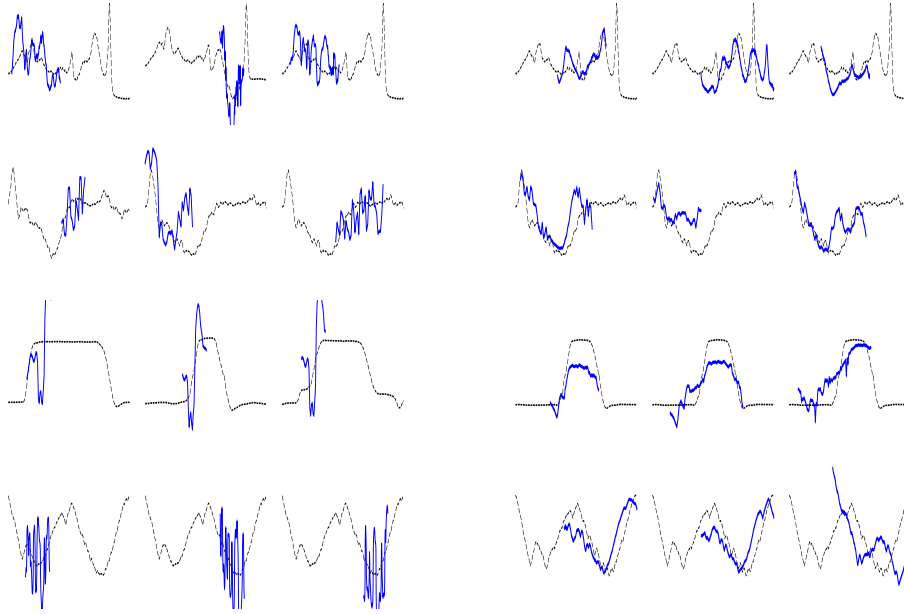


Fig. 3: Three most discriminative shapelets obtained for the datasets Beef, ECG200, GunPoint, Herring (rows 1 to 4, respectively) using (left column) Learning Shapelets or (right column) our XCNN architecture. The average discriminative power of the shapelets is evaluated using Eq. 2 and each shapelet is superimposed over its best matching time series in the test set.

Shapelet contributions The computation of the contribution of a shapelet to a decision is based on GradCAM (“Gradient-weighted Class Activation Mapping”) [22]. GradCAM is a very popular method used in computer vision to understand which parts of an original image is used by a trained neural network to make a particular classification decision. The ”interesting” parts are shown using a heat map on the original image. We recall that in a convolutional neural network, a *feature map* is the output of a particular layer of neurons. It somehow (ignoring the activation function) shows the response of a given convolution filter to the output of the previous layer. GradCAM computes the feature importance α_k^c of the feature map A^k on the classification decision c . This is computed after the final pooling layer which transforms all spatial positions (for images) A_{ij}^k of the k^{th} feature map to a single value F^k . The filter importance weight α_k^c , for a given input image (omitted for conciseness), is calculated with: $\alpha_k^c = \frac{\partial y^c}{\partial F^k}$ where y^c is the output of the network for class c .

Compared to the image classifiers used in [22], in our time series classification problem (1-dimensional) we are interested in both the **positive and negative contributions** of each learned shapelet on the classification of the (set of) series (whereas in [22] only the positive contributions matter). Those contributions are

defined for a trained network and a given time series Z_i (implicitly present in the partial derivatives) as: $p_k(Z_i) = ReLU\left(\frac{\partial y^c}{\partial F^k}\right)$ and $n_k(Z_i) = -ReLU\left(-\frac{\partial y^c}{\partial F^k}\right)$

As F^k is obtained from a global max pooling ($F^k = \max_t A_t^k$), each shapelet contribution can be associated to a timestamp $t = \arg \max_{t'} A_{t'}^k$, allowing us to localize the contribution. To produce a heat map with the positive contributions, we follow the same principle as in [22]: $L_{mask}(Z_i) = \sum_k p_k(Z_i) \tilde{A}^k(Z_i)$. where \tilde{A}^k is a vector of all zeros but at position $t = \arg \max_{t'} A_{t'}^k$ (where A_t^k is stored).

To obtain the **global positive contribution** of a shapelet k given a set of N time series examples, we compute

$$gp_k = \frac{1}{N} \sum_{i=1}^N p_k(Z_i). \quad (2)$$

The shapelets shown in Figure 3 are the 3 most contributing shapelets, according to this global criterion. In Figure 3, the shapelets learned by XCNN seem visually closer to the time series than the shapelets learned by LS. We then computed the average L_2 between a shapelet and a subpart of a time series over all the shapelets learned by XCNN and by LS for a given dataset, computed at the best matching point of the closest time series in the dataset (also in terms of L_2). The results are given in Fig. 4. This scatter-plot shows that, even if the optimized distance between the shapelets and the input series in the neural network is not the L_2 one (it is the dot product), our adversarial regularization allows XCNN to obtain closer (in terms of L_2) shapelets than LS which are deemed more suited for explanations.




Fig. 4: Average over all the shapelets learned by XCNN and by LS for a given dataset, of the L_2 distances between a shapelet and a subpart of a time series at the best matching point of the closest time series in the dataset.

Gradient-based explanations with XCNN shapelets Since we use a neural network classifier, we could directly benefit from the standard gradient-based

explanations, as also discussed in [7], to show what parts of a given time series example is important for the classifier to take its classification decision. These explanations would also be the ones produced by post-hoc methods such as LIME [20]. For lack of space, we do not show examples of such explanations but the interested reader can find many examples in [7] or in [14].

These, nowadays standard, gradient-based explanations are interesting but do not show the inner working of the classifier and, in particular, the reason why some parts of the input series were particularly useful for the classification. We believe that our ability, with XCNN, to show the shapelets that were learned and used to make the classification gives a different type of information than the gradient-based one. To illustrate this, we overlay in Figures 5 and 6 the three most positively (resp. negatively on the right) contributing shapelets on the time series at their best matching location (using L2 distance). Note that on the left side, the horizontal axis gives the length of the series (in black) while on the right, it gives the length of the shapelets which is at most 60% of the length of the series. We do not show the original series for the negative shapelets since, by definition, they are very far from the original series. In Figure 6 there is no negative shapelet used to discriminate the series of this dataset. This is due to the fact that the series for all the classes are very similar except for very small changes in the slope of the bump or in the size of the plateau at the top of the bump. These small changes can be captured by the positive shapelets but many of them are used to succeed in discriminating the classes.

We can also use our method to show the shapelets that most contribute to the classification of all examples of a *given class*. This is useful when one wants to understand the class characteristics. The global relative positive contribution of one shapelet considering all series from a given class is:

$$rp_k(c) = ReLU \left(\frac{1}{N^c} \sum_{i=1}^{N^c} \left(p_k^c(Z_i) - \frac{1}{\#c-1} \sum_{\substack{j=1 \\ j \neq c}}^{\#c-1} p_k^j(Z_i) \right) \right)$$

where N^c is the number of examples in class c , and $\#c$ is the total number of classes in the dataset. resp., we can get the $rn_k(c)$ by replacing the p_k^j with n_k^j . The time series shown in black in Figures 7 and 8 is the average over all examples of a given class. With these figures, we can draw similar conclusions as the previous ones but for an entire class.

4.3 Quantitative Results

XCNN is able to learn, by design, shapelets that are discriminative and suited for explanations. We want to quantify if this is achieved at the expense of classification accuracy and/or computation time. Our goal is to be much faster than exhaustive shapelet search methods (our baseline is Shapelets [25]), much more accurate than very fast random shapelet selection-based methods (our baseline is FS [19]) and as accurate and as fast as single model shapelet learning methods (our baselines are LS [10] and ELIS [6]).

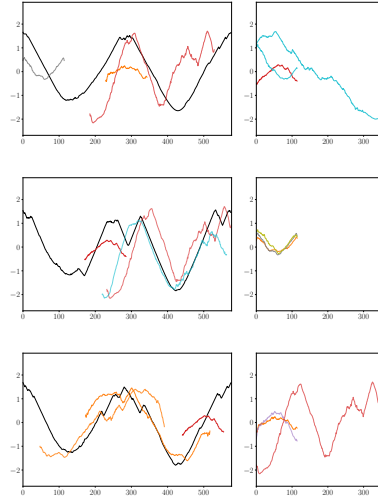


Fig. 5: Three most positively (left) and negatively (right) contributing shapelets for a random series (in black) of some of the classes (class 0, 1, 2, resp.) of the Car test set. Note that there were in total 90 positive shapelets used for this decision and about 140 negative ones.




Fig. 6: Three most positively (left) and negatively (right) contributing shapelets for a random series (in black) of some of the classes (class 0, 1, 2, resp.) of the DiatomSizeReduction test set. Note that there were in total 240 positive shapelets used for this decision and about 0 negative ones.

Accuracy We analyze the accuracies obtained by FS, LS, ELIS and our XCNN method on the 85 datasets using scatter plots. We compare FS versus XCNN in Figure 9, LS versus XCNN in Figure 10 and ELIS versus XCNN in Fig. 11. We also show how a simple CNN (without the adversarial regularization) compares against LS in Figure 12. We indicate the number of *win/tie/loss* for our method and we provide a Wilcoxon significance test [5] with the resulting p -value (> 0.01 : none of the two methods is significantly better than the other). The points on the diagonal are datasets for which the accuracy is identical for both competitors. Figure 9 shows that, as expected, our method yields significantly better performance than FS. It gives similar results (not significantly better nor worse on average) than ELIS for 52 datasets for which ELIS terminated in 48 hours. However for 33 datasets ELIS took more than 48 hours to complete. Compared to LS, for most datasets, the difference in accuracy is low, with a small edge (significant) for LS. On three datasets (namely HandOutlines, NonInvasiveFetalECGThorax1 and OliveOil), our XCNN method and its regularization seems to be strongly positive (and detrimental on one dataset), in terms of generalization. A simple CNN that would correspond to the classifier of our XCNN alone seems to give slightly better (non significant) results than LS (and thus than our XCNN). This means that

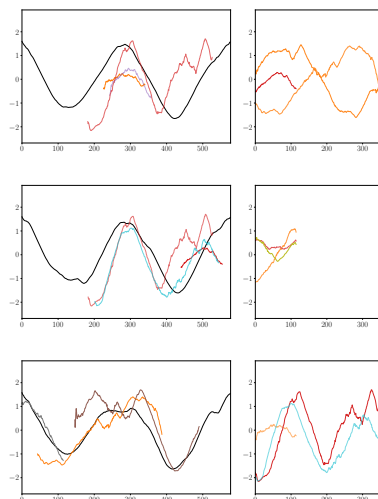


Fig. 7: Three most positively (left) and negatively (right) globally contributing shapelets for some of the classes (class 0, 1, 3, resp.) of the Car test set.

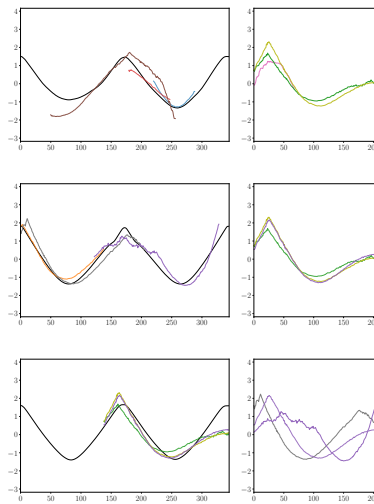


Fig. 8: Three most positively (left) and negatively (right) globally contributing shapelets for some of the classes (class 0, 1, 2, resp.) of the DiatomSizeReduction test set.

our backbone neural network architecture is a good candidate to jointly learn interpretable shapelets and classify time series with little loss on accuracy.

Table 1: Complexity of four different shapelet-based TSC algorithms (Shapelet [25], FS [19], LS [10] and XCNN). n is the number of examples in the training set, T is the average length of the time series, and c is the number of classes.

| Shapelet | FS | LS | XCNN |
|--------------------|------------------|--------------------------|--------------------------|
| $O(n^2 \cdot T^4)$ | $O(n \cdot T^2)$ | $O(n \cdot T^2 \cdot c)$ | $O(n \cdot T^2 \cdot c)$ |

Training Time We provide a theoretical complexity study (see Table 1) of all the baselines and of our XCNN method. Our method is based on a classifier and a discriminator, and both of them are simple CNNs. So the complexity of our algorithm ($O(n \cdot T^2 \cdot c)$) is related to training a CNN and should depend mainly on the number of examples (n), the average length of the time series (T), and the number of classes (c , since the latter is used to decide the number of shapelets to be learned). Note that for both LS and XCNN, the parameter n could be considered as a (quite big) constant since the number of epochs is fixed in the




Fig. 9: Accuracy comparison between Fast Shapelets (FS) and XCNN on 85 datasets (each point is a dataset) of the UCR/UEA repository [4].




Fig. 10: Accuracy comparison between Learning Shapelets (LS) and XCNN on 85 datasets (each point is a dataset) of the UCR/UEA repository [4].

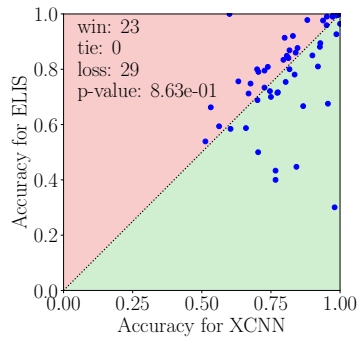


Fig. 11: Accuracy comparison between ELIS and XCNN on 85 datasets of the UCR/UEA repository [4].




Fig. 12: Accuracy comparison between Learning Shapelets (LS) and a simple CNN on 85 datasets of the UCR/UEA repository [4].

experiments. However, in LS, this number still depends on n whereas it is fixed (to 8000) in XCNN. This difference is in favor of LS for small datasets and in favor of XCNN for larger ones.

To have a better grasp on the actual training time of all methods, we ran the methods on a single dataset (ElectricDevices) and recorded the CPU time. The experiments were conducted on a Debian Cluster using Intel(R) Xeon(R) CPU E5-2650 v4 Processor (12 core 2.20 GHz CPU) with 32GB memory. The results are averaged over five runs. The implementation code of our baselines is taken from [2] (as for the accuracy results). As expected, the original Shapelet [25] method does not finish in 48 hours for this medium size dataset. FS finishes in 12.1 minutes, LS finishes in 2323 minutes, and our method takes 142 minutes. The theoretical complexity of LS and XCNN is identical so these results were

surprising. We suspected that the JAVA implementation of LS was not well optimized and we re-implemented the LS method with Keras. With this new implementation, the training phase took only 71 minutes for LS on this dataset (compared to 142 for XCNN) which shows that the time difference between the two algorithms is mainly related to the implementation (and the hyper-parameters related to the number of epochs).

5 Conclusion

We have presented a new shapelet-based time series classification method that produces shapelets that are, by design, better suited to explain decisions. The method is based on a novel adversarial architecture where one convolutional neural network is used to classify the series and another one is used to constrain the first network to learn both discriminative but also meaningful shapelets. Our results show that the expected trade-off between accuracy and interpretability is satisfactory: our classification results are comparable with similar state-of-the-art methods but with shapelets that can be used in many different way to explain the decisions.

In future work, we would like to first investigate the use of an additional regularization term to be able to determine automatically a minimal set of necessary shapelets. We also want to use our regularization on other types of data (such as multivariate time series, spatial data, graphs) and in a deep(er) CNN. Furthermore, we would like to adapt our approach to explain anomaly detections using neural network architectures such as convolutional auto-encoders or generative networks.

References

1. Assaf, R., Giurghi, I., Bagehorn, F., Schumann, A.: MTEX-CNN: Multivariate Time Series EXplanations for Predictions with Convolutional Neural Networks. In: 2019 IEEE International Conference on Data Mining (ICDM). pp. 952–957. <https://doi.org/10.1109/ICDM.2019.00106>
2. Bagnall, A., Lines, J., Bostrom, A., Large, J., Keogh, E.: 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. Bagnall, A., Lines, J., Hills, J., Bostrom, A.: Time-series classification with cote: the collective of transformation-based ensembles. *IEEE Transactions on Knowledge and Data Engineering* **27**(9), 2522–2535 (2015)
4. Chen, Y., Keogh, E., Hu, B., Begum, N., Bagnall, A., Mueen, A., Batista, G.: The ucr time series classification archive (July 2015), www.cs.ucr.edu/~eamonn/time_series_data/
5. Demšar, J.: Statistical comparisons of classifiers over multiple data sets. *J. Mach. Learn. Res.* **7**, 1–30 (Dec 2006)
6. Fang, Z., Wang, P., Wang, W.: Efficient learning interpretable shapelets for accurate time series classification. In: 2018 IEEE 34th International Conference on Data Engineering (ICDE). pp. 497–508 (2018)
7. Fawaz, H.I., Forestier, G., Weber, J., Idoumghar, L., Muller, P.A.: Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery* **33**(4), 917–963 (2019)

8. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Teh, Y.W., Titterton, M. (eds.) Proceedings of the International Conference on Artificial Intelligence and Statistics. vol. 9, pp. 249–256. PMLR (2010)
9. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in Neural Information Processing Systems. pp. 2672–2680 (2014)
10. Grabocka, J., Schilling, N., Wistuba, M., Schmidt-Thieme, L.: Learning time-series shapelets. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data mining. pp. 392–401 (2014)
11. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D.: A survey of methods for explaining black box models. *ACM Comput. Survey* **51**(5) (2018)
12. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. In: Advances in Neural Information Processing Systems (NIPS) (2017)
13. Karlsson, I., Papapetrou, P., Bostrom, H.: Generalized random shapelet forests. *Data Mining and Knowledge Discovery* **30**(5), 1053–1085 (Sep 2016)
14. Le Nguyen, T., Gsponer, S., Ilie, I., O’Reilly, M., Ifrim, G.: Interpretable time series classification using linear models and multi-resolution multi-domain symbolic representations. *Data Mining and Knowledge Discovery* pp. 1–40 (2019)
15. Lines, J., Davis, L.M., Hills, J., Bagnall, A.: A shapelet transform for time series classification. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data mining. pp. 289–297 (2012)
16. Lines, J., Taylor, S., Bagnall, A.: Time series classification with hive-cote: The hierarchical vote collective of transformation-based ensembles. *ACM Transactions on Knowledge Discovery from Data (TKDD)* **12**(5), 52 (2018)
17. Lundberg, S.M., Lee, S.: A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems (NIPS). pp. 4768–4777 (2017)
18. Ma, Q., Zhuang, W., Li, S., Huang, D., Cottrell, G.W.: Adversarial dynamic shapelet networks. In: The Thirty-Fourth Conference on Artificial Intelligence (AAAI). pp. 5069–5076 (2020)
19. Rakthanmanon, T., Keogh, E.: Fast shapelets: A scalable algorithm for discovering time series shapelets. In: proceedings of the 2013 SIAM International Conference on Data Mining. pp. 668–676. SIAM (2013)
20. Ribeiro, M.T., Singh, S., Guestrin, C.: “why should I trust you?”: Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD Int. Conference on Knowledge Discovery and Data Mining. pp. 1135–1144 (2016)
21. Ribeiro, M., Singh, S., Guestrin, C.: Anchors: High-precision model-agnostic explanations. In: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence. pp. 1527–1535 (2018)
22. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: IEEE International Conference on Computer Vision, ICCV. pp. 618–626 (2017)
23. Shumway, R.H., Stoffer, D.S.: *Time Series Analysis and Its Applications* (Springer Texts in Statistics). Springer-Verlag, Berlin, Heidelberg (2005)
24. Wang, Z., Yan, W., Oates, T.: Time series classification from scratch with deep neural networks: A strong baseline. In: Proceedings of the International Joint Conference on Neural Networks. pp. 1578–1585 (2017)
25. Ye, L., Keogh, E.: Time series shapelets: a new primitive for data mining. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data mining. pp. 947–956 (2009)