

dCAM: Dimension-wise Class Activation Map for Explaining Multivariate Data Series Classification

Paul Boniol
Université Paris Cité
boniol.paul@gmail.com

Mohammed Meftah
EDF R&D
mohammed.meftah@edf.fr

Emmanuel Remy
EDF R&D
emmanuel.remy@edf.fr

Themis Palpanas
Université Paris Cité & IUF
themis@mi.parisdescartes.fr

ABSTRACT

Data series classification is an important and challenging problem in data science. Explaining the classification decisions by finding the discriminant parts of the input that led the algorithm to some decision is a real need in many applications. Convolutional neural networks perform well for the data series classification task; though, the explanations provided by this type of algorithms are poor for the specific case of multivariate data series. Addressing this important limitation is a significant challenge. In this paper, we propose a novel method that solves this problem by highlighting both the temporal and dimensional discriminant information. Our contribution is two-fold: we first describe a convolutional architecture that enables the comparison of dimensions; then, we propose a method that returns dCAM, a Dimension-wise Class Activation Map specifically designed for multivariate time series (and CNN-based models). Experiments with several synthetic and real datasets demonstrate that dCAM is not only more accurate than previous approaches, but the only viable solution for discriminant feature discovery and classification explanation in multivariate time series.

CCS CONCEPTS

• **Information systems** → *Data management systems; Information systems applications*; • **Computing methodologies** → *Machine learning*.

KEYWORDS

database; time series; classification; deep learning

ACM Reference Format:

Paul Boniol, Mohammed Meftah, Emmanuel Remy, and Themis Palpanas. 2022. dCAM: Dimension-wise Class Activation Map for Explaining Multivariate Data Series Classification. In *Proceedings of the 2022 International Conference on Management of Data (SIGMOD '22)*, June 12–17, 2022, Philadelphia, PA, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3514221.3526183>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMOD '22, June 12–17, 2022, Philadelphia, PA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9249-5/22/06...\$15.00
<https://doi.org/10.1145/3514221.3526183>

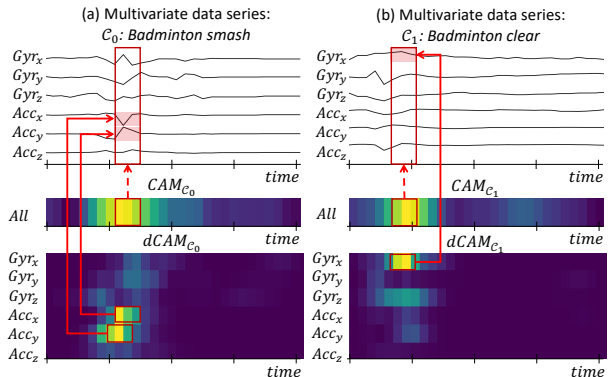


Figure 1: CAM and dCAM computed for two instances (of (a) "badminton smash" and (b) "badminton clear") of Racket-Sport UCR/UEA dataset

1 INTRODUCTION

Several applications across many domains produce big collections of data series¹, which need to be processed and analyzed [5, 40, 41, 43]. Typical analysis tasks include pattern matching (or similarity search) [15–17, 22, 34, 42, 49–51, 64], classification [14, 28, 38, 54, 55, 59, 61, 69], clustering [35, 44–46, 62], anomaly detection [6–8, 20, 47, 56, 68], motif discovery [19, 36, 75], and others [29].

Data series classification is a crucial and challenging problem in data science [18, 67]. To solve this task, various data series classification algorithms have been proposed in the past few years [3], applied on a large number of use cases. Standard data series classification methods are based on distances to the instances' nearest neighbors, with k-NN classification (using the Euclidean or Dynamic Time Warping (DTW) distances) being a popular baseline method [12]. Nevertheless, recent works have shown that ensemble methods using more advanced classifiers achieve better performance [4, 37]. Following recent breakthroughs in the computer vision community, new studies successfully propose deep learning methods for data series classification [9, 11, 27, 32, 63, 72, 73], such as Convolutional Neural Network (CNN), Residual Neural Network (ResNet) [65], and InceptionTime [28].

[Classification Explanation] While having a trained and accurate classification model, finding explanations of the classification result (i.e., finding the discriminative features that made the model decide which class to attribute to each instance) is a challenging but essential problem, e.g., in manufacturing for anomaly-based predictive maintenance [70], or in medicine for robot-assisted surgeon training [26]. Such discriminant features can be based on patterns of interest that occur in a subset of dimensions at different timestamps

¹A data series, or data sequence, is an ordered sequence of points. If the dimension that imposes the ordering of the sequence is time, then we talk about *time series*. In this paper, we use the terms *sequence*, *data series*, and *time series* interchangeably.

or the same timestamp. For some CNN-based models, the Class Activation Map (CAM) [74] can be used as an explanation for the classification result. CAM has been used for highlighting the parts of an image that contribute the most to a given class prediction and has also been adapted to data series [27, 65].

[Challenges] Nevertheless, CAM for data series suffers from one important limitation. Since CAM is a univariate time series (of the same length as the input instances) with high values aligned with the subsequences of the input that contribute the most for a given class identification, in the specific case of multivariate data series as input, no information can be retrieved from CAM on the level of contribution of specific dimensions. As an example, Figure 1 illustrates CAM (top heatmaps) applied on two instances (belonging to two different classes) of the RacketSport UCR/UEA dataset. We observe that CAM explains why the data series correspond to a badminton "smash" or "clear" gesture by highlighting the same temporal window across all dimensions (variables). It is thus not clear what aspect of the gesture distinguishes it from the other. Addressing this significant limitation is a sought-after challenge.

[Contributions] In this paper, we present a novel approach that fills-in the gap by addressing this limitation for the popular CNN-based models. We propose a novel data organization and a new CAM technique, dCAM (Dimension-wise Class Activation Map), that is able to highlight both the *temporal and dimensional* information at the same time. For instance, in Figure 1, dCAM (bottom heatmaps) is pointing to specific subsequences of particular dimensions that explain why the two gestures are different. Our method requires only a single training phase, is not constrained by the architecture type, and can efficiently and effectively retrieve discriminant features thanks to a technique that exploits information from different permutations of the input data dimensions. Thus, any kind of architecture in which we can apply CAM can benefit from our approach. Our contributions are as follows.

- We develop a new method that transforms convolutional-based neural network architectures: whereas previous network architectures can only provide an explanation for all the dimensions together, our transformation represents the only deep learning solution that enables explanation in individual dimensions. Our approach can be used with any deep network architecture with a Global Average Pooling layer.
- We demonstrate how we can apply our method to three modern deep learning classification architectures. We first describe dCNN, inspired by the traditional CNN architecture. We then describe how more advanced architectures, such as ResNet and InceptionTime, can be transformed as well. We name these transformed architectures dResNet and dInceptionTime.
- We propose dCAM, a novel method (based on dCNN/ dResNet/dInceptionTime) that returns a multivariate CAM, identifying the important parts of the input series for *each* dimension.
- We experimentally demonstrate with several synthetic and real datasets that (among Class Activation Map-based methods) dCAM is not only more accurate in classification than previous approaches, but the only viable solution for discriminant feature discovery and classification explanation in multivariate time series. Finally, we make our code available online [1].

2 BACKGROUND AND RELATED WORK

We first present useful notations and definitions, and discuss related work. Table 1 summarizes the symbols we use in this paper.

[Data Series] A multivariate, or D -dimensional data series $T \in \mathbb{R}^{(D,n)}$ is a set of D univariate data series of length n . We note $T = [T^{(0)}, \dots, T^{(D-1)}]$ and for $j \in [0, D-1]$, we note the univariate data series $T^{(j)} = [T_0^{(j)}, T_1^{(j)}, \dots, T_{n-1}^{(j)}]$. A subsequence $T_{i,\ell}^{(j)} \in \mathbb{R}^\ell$ of the dimension $T^{(j)}$ of the multivariate data series T is a subset of contiguous values from $T^{(j)}$ of length ℓ (usually $\ell \ll n$) starting at position i ; formally, $T_{i,\ell}^{(j)} = [T_i^{(j)}, T_{i+1}^{(j)}, \dots, T_{i+\ell-1}^{(j)}]$.

[Neural Network Notations] We are interested in classifying data series using a neural network architecture model. We define the neural network input as $X \in \mathbb{R}^n$ for univariate data series (with x_i the i^{th} value and $X_{i,\ell}$ the sequence of ℓ values following the i^{th} value), and $\mathbf{X} \in \mathbb{R}^{(D,n)}$ for multivariate data series (with $x_{j,i}$ the i^{th} value on the j^{th} dimension and $\mathbf{X}_{j,i,\ell}$ the sequence of ℓ values following the i^{th} value on the j^{th} dimension).

Dense Layer: The basic layer of neural network is a fully connected layer (also called *Dense layer*) in which every input neuron is weighted and summed before passing through an activation function. For univariate data series, given an input data series $X \in \mathbb{R}^n$, given a vector of weights $W \in \mathbb{R}^n$ and a vector $B \in \mathbb{R}^n$, we have:

$$h = f_a \left(\sum_{x_i, w_i, b_i \in (X, W, B)} w_i * x_i + b_i \right) \quad (1)$$

f_a is called the activation function and is a non-linear function. The commonly used activation function f_a is the rectified linear unit (ReLU) [39] that prevents the saturation of the gradient (other functions that have been proposed are *Tanh* and *Leaky ReLU* [66]). For the specific case of multivariate data series, all dimensions are concatenated to give input $X, W \in \mathbb{R}^{D*n}$. Finally, one can decide to have several output neurons. In this case, each neuron is associated with a different W and B , and Equation 1 is executed independently. **Convolutional Layer:** This layer has played a significant role for image classification [31, 33, 65], and recently for data series classification [27]. Formally, for multivariate data series, given an input vector $\mathbf{X} \in \mathbb{R}^{(D,n)}$, and given matrices weights $\mathbf{W}, \mathbf{B} \in \mathbb{R}^{(D,\ell)}$, the output $h \in \mathbb{R}^n$ of a convolutional layer can be seen as a univariate data series. The tuple (W, B) is also called kernel, with (D, ℓ) the size of the kernel. Formally, for $h = [h_0, \dots, h_n]$, we have:

$$h_i = f_a \left(\sum_{\substack{X^{(j)}, W^{(j)}, B^{(j)} \in \\ (X, W, B)}} \sum_{\substack{x_k, w_k, b_k \in \\ (X^{(j)}_{i-\lfloor \frac{\ell}{2} \rfloor, i+\lfloor \frac{\ell}{2} \rfloor), W^{(j)}, B^{(j)}}} w_k * x_k + b_k \right) \quad (2)$$

In practice, we have several kernels of size (D, ℓ) . The result is a multivariate series with dimensions equal to the number of kernels, n_f . For a given input $\mathbf{X} \in \mathbb{R}^{(D,n)}$, we define $A \in \mathbb{R}^{(n_f, n)}$ to be the output of a convolutional layer $conv(n_f, \ell)$. A_m is thus a univariate series corresponding to the output of the m^{th} kernel. We denote with $A_m(T)$ the univariate series corresponding to the output of the m^{th} kernel, when T is used as input.

Global Average Pooling Layer: Another type of layer frequently used is pooling. Pooling layers compute average/max/min operations, aggregating values of previous layers into a smaller number of values for the next layer. A specific type of pooling layer is Global

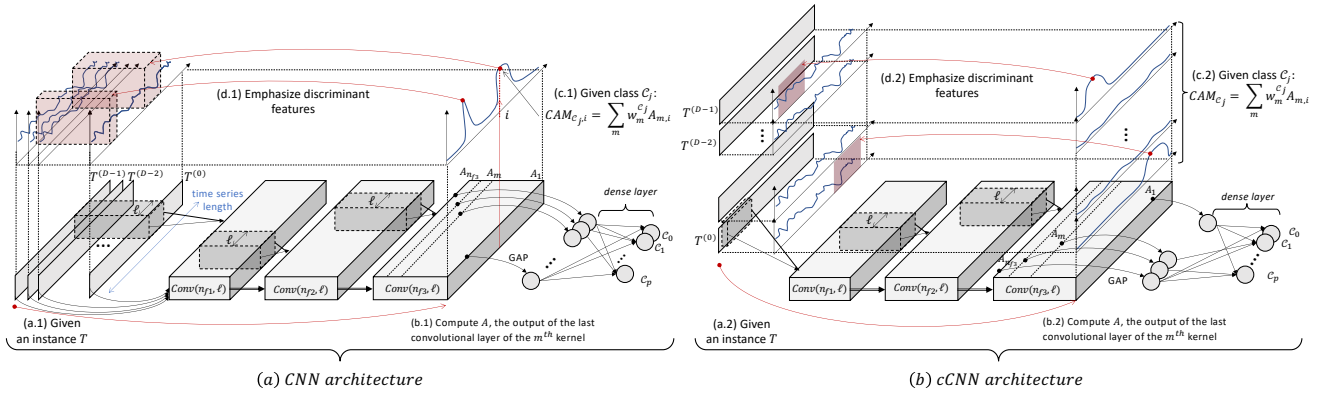


Figure 2: Illustration of Class Activation Map for (a) CNN architecture and (b) cCNN architecture with three convolutional layers (n_{f1} , n_{f2} , and n_{f3} different kernels respectively of size all equal to ℓ).

Average Pooling (GAP). This operation is averaging an entire output of a convolutional layer $A_m(T)$ into one value, thus providing invariance to the position of the discriminative features.

Learning Phase: The learning phase uses a loss function \mathcal{L} that measures the accuracy of the model and optimizes the various weights. For the sake of simplicity, we note Ω the set containing all weights (e.g., matrices \mathbf{W} and \mathbf{B} defined in the previous sections). Given a set of instances \mathcal{X} , we define the average loss as: $J(\Omega) = \frac{1}{|\mathcal{X}|} \sum_{X \in \mathcal{X}} \mathcal{L}(X)$. Then for a given learning rate α , the average loss is back-propagated to all weights in the different layers. Formally, back-propagation is defined as follows: $\forall \omega \in \Omega, \omega \leftarrow \omega - \alpha \frac{\partial J}{\partial \omega}$. In this paper, we use the stochastic gradient descent using the ADAM optimizer [30] and cross-entropy loss function.

2.1 Convolutional-based Neural Network

We now describe the standard architectures used in the literature. The first is Convolutional Neural Networks (CNNs) [27, 65]. CNN is a concatenation of convolutional layers (joined with *ReLU* activation functions and batch normalization). The last convolutional layer is connected to a Global Average Pooling layer and a dense layer. In theory, instances of multiple lengths can be used with the same network. A second architecture is the Residual Neural Network (ResNet) [27, 65]. This architecture is based on the classical CNN, to which we add residual connections between successive blocks of convolutional layers to avoid that the gradients explode or vanish. Other methods have been proposed in the literature [11, 27, 28, 58], though, CNN and ResNet have been shown to perform the best for multivariate time series classification [27]. InceptionTime [28] has not been evaluated on multivariate data series, but demonstrated state-of-the-art performance on univariate data series. Finally, other kinds of architectures than convolutional ones have been proposed in the literature. Attention-based models have been introduced, such as TapNet [71]. For the specific case of temporal data, recurrent-based models, such as Recurrent Neural Networks [53] (RNN), Long-Short Term Memory [23] (LSTM), and Gated Recurrent Unit [10] (GRU) have received a lot of attention. These three models are relevant to the data series classification task, and we include them in our experimental study.

2.2 Class Activation Map (CAM)

Once the model is trained, we need to find the discriminative features that led the model to decide which class to attribute to each instance. Several methods have been proposed to extract meaningful information from CNNs, such as grad-CAM [57] that uses the gradients of the weights to compute the discriminant features, and CAM [74]. The latter has been proposed to highlight the parts of an image that contributes the most to a given class identification. The latter has been experimented on data series [27, 65] (univariate and multivariate). This method explains the classification of a certain deep learning model by emphasizing the subsequences that contribute the most to a certain classification. Note that the CAM method can only be used if (i) a Global Average Pooling layer has been used before the *softmax* classifier, (ii) the model accuracy is high enough. Thus, only the standard architectures CNN and ResNet proposed in the literature can benefit from CAM. We now define the CAM method [27, 65]. For an input data series T , let $A(T)$ be the result of the last convolutional layer $conv(n_f, \ell)$, which is a multivariate data series with n_f dimensions and of length n . $A_m(T)$ is the univariate time series for the dimension $m \in [1, n_f]$ corresponding to the m^{th} kernel. Let $w_m^{C_j}$ be the weight between the m^{th} kernel and the output neuron of class $C_j \in \mathcal{C}$. Since a Global Average Pooling layer is used, the input to the neuron of class C_j can be expressed by the following equation:

$$z_{C_j}(T) = \sum_m w_m^{C_j} \sum_{A_{m,i}(T) \in A_m(T)} A_{m,i}(T).$$

The second sum represents the averaged time series over the whole time dimension. Note that weight $w_m^{C_j}$ is independent of index i . Thus, z_{C_j} can also be written by the following equation:

$$z_{C_j}(T) = \sum_{A_{m,i}(T) \in A_m(T)} \sum_m w_m^{C_j} A_{m,i}(T).$$

Finally, $CAM_{C_j}(T) = [CAM_{C_j,0}(T), \dots, CAM_{C_j,n-1}(T)]$ that underlines the discriminative features of class C_j is defined as follows:

$$\forall i \in [0, n-1], CAM_{C_j,i}(T) = \sum_m w_m^{C_j} A_{m,i}(T).$$

As a consequence, $CAM_{C_j}(T)$ is a univariate data series where each element at index i indicates the significance of the index i (regardless of the dimensions) for the classification as class C_j . Figure 2(a) depicts the process of computing CAM and finding the discriminant subsequences in the initial series.

2.3 CAM Limitations for Multivariate Series

As mentioned earlier, a CAM that highlights the discriminative subsequences of class C_j , $CAM_{C_j}(T)$, is a univariate data series. The information provided by $CAM_{C_j}(T)$ is sufficient for the case of univariate series classification, but not for multivariate series classification. Even though the significant temporal index may be correctly highlighted, no information can be retrieved on which dimension is significant or not. Solving this serious limitation is a significant challenge in several domains. For that purpose, one can propose rearranging the input structure to the network so that the CAM becomes a multivariate data series. A new solution would be to decide to use a 2D convolutional neural network with kernel size $(\ell, 1)$, such that each kernel slides on each dimension separately. Thus, for an input data series T , $A_m(T)$ would become a multivariate data series for the variable $m \in [1, n_f]$, and $A_m^{(d)}(T) \in A_m(T)$ would be a univariate time series that would correspond to the dimension d of the initial data series. We call this solution *cCNN*, and we use *cCAM* to refer to the corresponding Class Activation Map. Figure 2(b) illustrates *cCNN* architecture and *cCAM*. Note that if a GAP layer is used, then architectures other than CNN can be used, as well, such as ResNet and InceptionTime. We denote these baselines as *cResNet* and *cInceptionTime*.

Nevertheless, new limitations arise from this solution. The dimensions are not compared together: each kernel of the input layer will take as input only one of the dimensions at a time. Thus, features depending on more than one dimension will not be detected.

Recent studies study the specific case of multivariate data series classification explanation. A benchmark study analyzing the saliency/explanation methods for multivariate time series concluded that the explainable methods work better when the multivariate data series is handled as an image [25], such as in the *cCNN* architecture. This confirms the need to propose a method specifically designed for multivariate data series. Finally, some recently proposed approaches [2, 24] address the problems of identifying the discriminant features and discriminant temporal windows independently from one another. For instance, MTEX-CNN [2] is an architecture composed of two blocks. The 1st block is similar to *cCNN*. The 2nd block consists of merging the results of the 1st block into a 1D convolutional layer, which enables comparing dimensions. A variant of CAM [57] is applied to the last convolutional layer of the 1st block in order to find discriminant features for each dimension. The discriminant temporal windows are detected with the CAM applied to the last convolutional layer of the 2nd block. In practice however, this architecture does not manage to overcome the limitations of *cCNN*: discriminant features that depend on several dimensions are not correctly identified by MTEX-CNN, which has similar accuracy to *cCNN* (we elaborate on this in Section 5).

In our experimental evaluation, we compare our approach to the MTEX-CNN, *cCNN*, *cResNet* and *cInceptionTime*, and further demonstrate their limitations when addressing the problem at hand.

3 PROBLEM FORMULATION

Given a set \mathcal{T} of multivariate data series $T = \{T^{(0)}, T^{(1)}, \dots, T^{(D-1)}\}$ of D dimensions belonging to classes $C_j \in \mathcal{C}$, and a model $f : \mathcal{T} \rightarrow \mathcal{C}$, we aim to find a function $g(T, f)$ that returns a multivariate series $g(T, f, C_j) = \{T^{(0)'}, T^{(1)'}, \dots, T^{(D-1)'}\}$, in which $T^{(i)'}$

Symbol	Description
T	a data series
$ T $	length of T
$T^{(i)}$	i^{th} dimensions of T
D	number of dimension
\mathcal{C}	set of all classes
C_j	one class of \mathcal{C}
$w_m^{C_j}$	weight of connecting the m^{th} convolutional layer and class C_j neuron
$A_m(T)$	output of the m^{th} convolutional layer for input T
$z_{C_j}(T)$	output of C_j neuron for input T
$CAM_{C_j}(T)$	Class Activation Map for class C_j and input T
Σ_T	set of all possible permutations of T dimensions
S_T^i	T with one possible permutation of its dimensions ($S_T^i \in \Sigma_T$)
k	number of permutations
n_g	number of permutations that have been correctly classified by the model

Table 1: Table of symbols.

is a series that has high values if the corresponding subsequences in T_i discriminate T of belonging to another class than C_j .

4 PROPOSED APPROACH

Based on a new architecture that we call *dCNN* (and variant architectures, e.g., *dResNet*, *dInceptionTime*), *dCAM* aims to provide a multivariate CAM pointing to the discriminant features within each dimension. Contrary to the previously described baseline (*cCNN*, *cResNet* and *cInceptionTime*), one kernel on the first convolutional layer will take as input all the dimensions together with different permutations. Thus, similarly to the standard CNN architecture, features depending on more than one dimension will be detectable while still having a multivariate CAM. Nevertheless, the latter has to be processed such that the significant subsequences are detected.

We first describe the proposed architecture *dCNN* that we need in order to provide a *dCAM*, while still being able to extract multivariate features. We then demonstrate that the transformation needed to change CNN to *dCNN* can also be applied to other, more sophisticated architectures, such as ResNet and InceptionTime, which we denote as *dResNet* and *dInceptionTime*. We demonstrate that using permutations of the input dimensions makes the classification more robust when important features are localized into small subsequences within some specific dimensions.

We then present in detail how we compute *dCAM* (based on a *dCNN*/*dResNet*/*dInceptionTime* architecture). Our solution benefits from the permutations injected into the *dCNN* to identify the most discriminant subsequences used for the classification decision.

4.1 Dimension-wise Architecture

As mentioned earlier, the classical CNN architecture mixes all dimensions in the first convolutional layer. Thus, the CAM is a univariate data series and does not provide any information on which dimension is the discriminant one for the classification. To address this issue, we can use a two-dimensional CNN architecture by reorganizing the input (i.e., the *cCNN* solution we described earlier). In this architecture, one kernel (of size $(1, \ell, 1)$) slides on each dimension independently. Thus, for a given data series $(T^{(0)}, \dots, T^{(D-1)})$ of length n , the convolutional layers returns an array of three dimensions (n_f, D, n) , each row $m \in [0, D - 1]$ corresponding to the

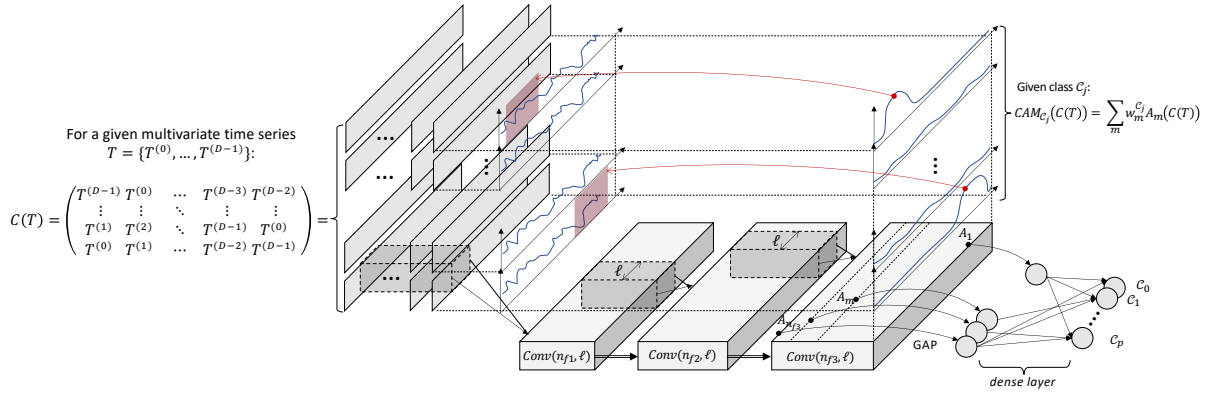


Figure 3: dCNN architecture and application of the CAM.

extracted features on dimension m . Nevertheless, the kernels $(1, \ell, 1)$ get as input each dimension independently: such an architecture cannot learn features that depend on multiple dimensions.

4.2 A first Architecture: dCNN

In order to have the best of both cases, we propose the dCNN architecture, where we transform the input into a cube, in which each row contains a given combination of all dimensions. One kernel (of size $(D, \ell, 1)$) slides on all dimensions D times. This allows the architecture to learn features on multiple dimensions simultaneously. Moreover, the resulting CAM is a multivariate data series. In this case, one row of the CAM corresponds to a given combination of the dimensions. However, we still need to be able to retrieve information for each dimension separately, as well. To do that, we make sure that each row contains a different permutation of the dimensions. As the weights of the kernels are at fixed positions (for specific dimensions), a permutation of the dimensions will result in a different CAM. Formally, for a given data series T , we note $C(T) \in \mathbb{R}^{(D,D,n)}$ the input data structure of dCNN:

$$C(T) = \begin{pmatrix} T^{(D-1)} & T^{(0)} & \dots & T^{(D-3)} & T^{(D-2)} \\ \vdots & \vdots & \dots & \vdots & \vdots \\ T^{(1)} & T^{(2)} & \dots & T^{(D-1)} & T^{(0)} \\ T^{(0)} & T^{(1)} & \dots & T^{(D-2)} & T^{(D-1)} \end{pmatrix}$$

Note that each row and column of $C(T)$ contains all dimensions. Thus, a given dimension $T^{(i)}$ is never at the same position in $C(T)$ rows. The latter is a crucial property for the computation of dCAM. In practice, we guarantee the latter property by shifting by one position the order of the dimensions. Thus $T^{(0)}$ in the first row is aligned with $T^{(1)}$ in the second row. A different order of T dimensions will thus generate a different matrix $C(T)$.

Figure 3 depicts the dCNN architecture. The input $C(T)$ is forwarded into a classical two-dimensional CNN. The rest of the architecture is independent of the input data structure. The latter means that any other two-dimensional architecture (containing a Global Average Pooling) can be used (such as ResNet), by only adapting the input data structure. Similarly, the training procedure can be freely chosen by the user. For the rest of the paper, we will use the cross-entropy loss function and the ADAM optimizer.

Observe that multiple permutations of the original multivariate series will be processed by several convolutional layers, enabling the kernel to examine multiple different combinations of dimensions

and subsequences. Note that the kernels of the dCNN will be sparse, which has a significant impact on overfitting.

4.3 The dResNet/dInceptionTime Architectures

As mentioned earlier, any architecture using a GAP layer after the last convolutional layer can benefit from dCAM. Thus, different (and more sophisticated) architectures can be used with our approach. To that effect, we propose two new architectures dResNet and dInceptionTime, based on the state-of-the-art architectures ResNet [65] and InceptionTime [28]. The transformations that lead to dResNet and dInceptionTime are very similar to that from CNN to dCNN, using $C(T)$ as input to the transformed networks. The convolutional layers are transformed from 1D (as originally proposed [28, 65]) to 2D. Similarly to dCNN, the kernel sizes are $(D, \ell, 1)$ and convolute over each row of $C(T)$ independently.

We demonstrate in the experimental section that these architectures do not affect the usage of our proposed approach dCAM, and we evaluate the choice of architecture on both classification and discriminant features identification.

4.4 Dimension-wise Class Activation Map

At this point, we have our network trained to classify instances among classes C_0, C_1, \dots, C_p . We now explain how to compute dCAM that will identify discriminant features within dimensions. We assume that the network has to be accurate enough in order to provide a meaningful dCAM. We evaluate in the experimental section the relation between the classification accuracy of the network and the discriminant features identification accuracy of dCAM.

At first glance, we can compute the regular CAM $CAM_{C_j}(C(T)) = \sum_m w_m^j A_m(C(T))$. However, a high value on the i^{th} row at position t on $CAM_{C_j}(C(T))$ does not mean that the subsequence at position t on the i^{th} dimension is important for the classification. It instead means that the combination of dimensions at the i^{th} row of $C(T)$ is important.

4.4.1 Random Permutation Computations. Given those different combinations of dimensions (i.e., one row of $C(T)$) produce different outputs (i.e., the same row in $CAM_{C_j}(C(T))$), the positions of the dimensions within the $C(T)$ rows have an impact on the CAM. Consequently, for a given combination of dimensions, we

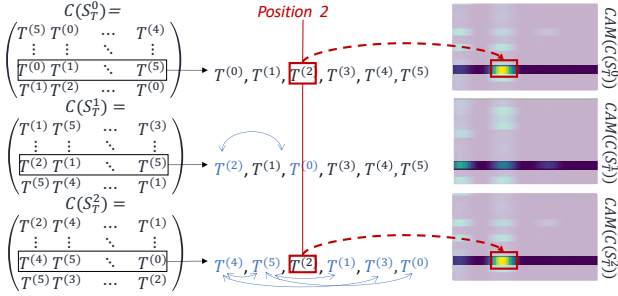


Figure 4: Example of Class Activation Map results for different permutations.

can assume that at least one dimension at a given position is responsible for the high value in the CAM row. For the rest of this paper, we use Σ_T as the set of all possible permutations of T dimensions, and $S_T^i \in \Sigma_T$ for a single permutation of T . E.g., for a given data series $T = \{T^{(0)}, T^{(1)}, T^{(2)}\}$, one possible permutation is $S_T^i = \{T^{(1)}, T^{(0)}, T^{(2)}\}$.

Figure 4 depicts an example of CAMs for different permutations. In this Figure, for three given permutations of T (i.e., S_T^0 , S_T^1 and S_T^2), we notice that when $T^{(2)}$ is in position two of the second row of $C(S_T^i)$, the Class Activation Map $CAM(C(S_T^i))$ is greater than when $T^{(2)}$ is not in position two. We infer that the second dimension of T in position two is responsible for the high value. Thus, we may examine different dimension combinations by keeping track of which dimension at which position is activating the CAM the most. We now describe the steps necessary to retrieve this information.

DEFINITION 1. For a given data series $T = \{T^{(0)}, T^{(1)}, \dots, T^{(D-1)}\}$ of length n and its input data structure $C(T)$, we define function idx , such that $idx(T^{(i)}, p_j)$ returns the row indices in $C(T)$ that contain the dimension $T^{(i)}$ at position p_j .

We can now define the following transformation \mathcal{M} .

DEFINITION 2. For a given data series $T = \{T^{(0)}, T^{(1)}, \dots, T^{(D-1)}\}$ of length n , a given class C_j and Class Activation Map, we define $\mathcal{M}(CAM_{C_j}(C(T))) \in \mathbb{R}^{(D, D, n)}$ (with $CAM_{C_j}(C(T)) \in \mathbb{R}^{(D, n)}$ and $CAM_{C_j}(C(T))_i$ its i^{th} row) as follows:

$$\mathcal{M}(CAM_{C_j}(C(T))) = \begin{pmatrix} CAM_{C_j}(C(T))_{idx(T^{(0)}, 0)} & \dots & CAM_{C_j}(C(T))_{idx(T^{(0)}, D-1)} \\ CAM_{C_j}(C(T))_{idx(T^{(1)}, 0)} & \dots & CAM_{C_j}(C(T))_{idx(T^{(1)}, D-1)} \\ \vdots & \dots & \vdots \\ CAM_{C_j}(C(T))_{idx(T^{(D-1)}, 0)} & \dots & CAM_{C_j}(C(T))_{idx(T^{(D-1)}, D-1)} \end{pmatrix} \quad (3)$$

Figure 5 depicts the \mathcal{M} transformation. As explained in Definition 2, the \mathcal{M} transformation enriches the Class Activation Map by adding the dimension position information. Note that if we change the dimension order of T , their $\mathcal{M}(CAM_{C_j}(C(T)))$ changes as well. Indeed, for a given dimension $T^{(i)}$ and position p_j , $idx(T^{(i)}, p_j)$ will not have the same value for two different dimension orders of T . Thus, computing $\mathcal{M}(CAM_{C_j}(C(T)))$ for different dimension orders of T will provide distinct information regarding the importance of a given position (subsequence) in a given dimension. We expect that subsequences (of a specific dimension) that discriminate one class from another will also be associated (most of the time) with a high value in the Class Activation Map.

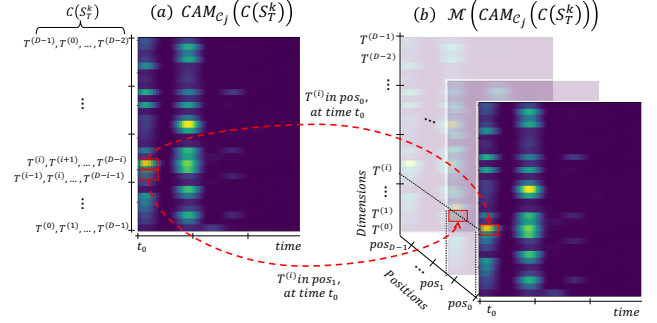


Figure 5: Transformation \mathcal{M} for a given data series T .

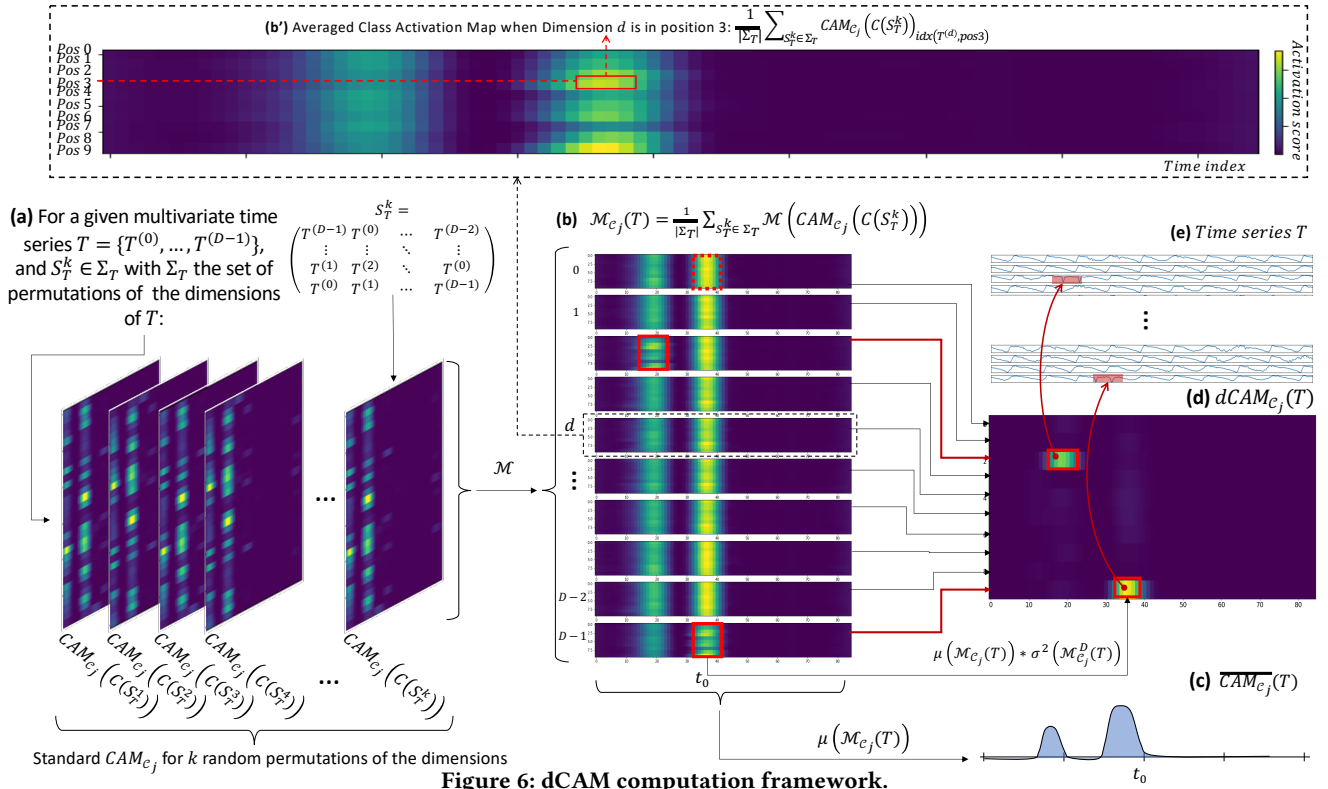
4.4.2 Merging Permutations. We compute $\mathcal{M}(CAM_{C_j}(C(S_T)))$, for different $S_T \in \Sigma_T$. Note that the total number of permutations for high-dimensional data series is enormous: $|\Sigma_T| = D!$. In practice, we only compute \mathcal{M} for a randomly selected subset of Σ_T . We thus merge $k = |\Sigma_T|$ permutations S_T^k , by computing the averaged matrix $\bar{\mathcal{M}}_{C_j}(T)$ of all the \mathcal{M} transformations of the permutations:

$$\bar{\mathcal{M}}_{C_j}(T) = \frac{1}{|\Sigma_T|} \sum_{S_T^k \in \Sigma_T} \mathcal{M}(CAM_{C_j}(C(S_T^k)))$$

Figure 6 illustrates the process of computing $\bar{\mathcal{M}}_{C_j}(T)$ from the set of permutations of T , Σ_T . $\bar{\mathcal{M}}_{C_j}(T)$ can be seen as a summarization of the importance of each dimension at each position in $C(T)$, for all the computed permutations. Figure 6(b') (at the top of the figure) depicts $\bar{\mathcal{M}}_{C_j}(T)_d$, which corresponds to the d^{th} row (i.e., the dotted box in Figure 6(b)) of $\bar{\mathcal{M}}_{C_j}(T)$. Each row of $\bar{\mathcal{M}}_{C_j}(T)_d$ corresponds to the average activation of dimension d (for each timestamp) when dimension d is in a given position in $C(T)$.

Note that all permutations of T are forwarded into the dCNN network without training it again. Thus, even though the permutations of T generate radically different inputs to the network, the network can still classify most of the instances correctly. For k permutations, we use n_g to denote the number of permutations that the model has correctly classified. We provide an analysis (see Section 5) of n_g/k w.r.t the classification accuracy of the model and the impact that n_g/k has on the discriminant features identification accuracy.

4.4.3 dCAM Extraction. We can now use the previously computed $\bar{\mathcal{M}}_{C_j}$ to extract explanatory information on which subsequences are considered important by the network. First, we note that each row of $C(T)$ corresponds to the input format of the standard CNN architecture. Thus, we expect that the result of a row of $\bar{\mathcal{M}}_{C_j}$ (one of the ten lines in Figure 6(b)) is similar to the standard CAM. Hence, we can assume that $\mu(\bar{\mathcal{M}}_{C_j}(T)) = \sum_{d \in [0, D-1]} \sum_{p \in [0, D-1]} \bar{\mathcal{M}}_{C_j}^{d,p}(T) / (2 * D)$ is equivalent to standard Class Activation Map $CAM_{C_j}(T)$ (this approximation is depicted in Figure 6(d)). Moreover, in addition to the temporal information, we can extract temporal information per dimension. We know that for a given position p and a given dimension d , $\bar{\mathcal{M}}_{C_j}^{d,p}(T)$ represents the averaged activation for a given set of permutations. If the activation $\bar{\mathcal{M}}_{C_j}^{d,p}(T)$ for a given dimension is constant (regardless of its value, or the position p), then the position of dimension d is not important, and no subsequence in that dimension d is discriminant. On the other hand, a high or low value at a



specific position p means that the subsequence at this specific position is discriminant. While it is intuitive to interpret a high value, interpreting a low value is counterintuitive. Usually, a subsequence at position p with a low value should be regarded as non-discriminant. Nevertheless, if the activation is low for p and high for other positions, then the subsequence at position p is the consequence of the low value and is thus discriminant. We experimentally observe this situation, where a non-discriminant dimension has a constant activation per position (e.g., see dotted red rectangle in Figure 6(b): this pattern corresponds to a non-discriminant subsequence of the dataset). On the contrary, for discriminant dimensions, we observe a strong variance for the activation per position: either high values or low values (e.g., see solid red rectangles in Figure 6(b): these patterns correspond to the (injected) discriminant subsequences highlighted in red in Figure 6(e)). We thus can extract the significant subsequences per dimension by computing the variance of all positions of a given dimension. We can filter out the irrelevant temporal windows using the averaged $\mu(\mathcal{M}_{C_j}(T))$ for all dimensions, and use the variance to identify the important dimensions in the relevant temporal windows. Formally, we define $\text{dCAM}_{C_j}(T)$ as follows.

DEFINITION 3. For a data series T and class C_i , $\text{dCAM}_{C_j}(T)$ is:

$$\text{dCAM}_{C_j}(T) = \begin{pmatrix} \sigma^2(\tilde{\mathcal{M}}_{C_j}^0(T)_{t_0}) * \mu(\tilde{\mathcal{M}}_{C_j}(T)_{t_0}) & \dots & \sigma^2(\tilde{\mathcal{M}}_{C_j}^0(T)_{t_n}) * \mu(\tilde{\mathcal{M}}_{C_j}(T)_{t_n}) \\ \vdots & \dots & \vdots \\ \sigma^2(\tilde{\mathcal{M}}_{C_j}^{D-2}(T)_{t_0}) * \mu(\tilde{\mathcal{M}}_{C_j}(T)_{t_0}) & \dots & \sigma^2(\tilde{\mathcal{M}}_{C_j}^{D-2}(T)_{t_n}) * \mu(\tilde{\mathcal{M}}_{C_j}(T)_{t_n}) \\ \sigma^2(\tilde{\mathcal{M}}_{C_j}^{D-1}(T)_{t_0}) * \mu(\tilde{\mathcal{M}}_{C_j}(T)_{t_0}) & \dots & \sigma^2(\tilde{\mathcal{M}}_{C_j}^{D-1}(T)_{t_n}) * \mu(\tilde{\mathcal{M}}_{C_j}(T)_{t_n}) \end{pmatrix} \quad (4)$$

4.5 Time Complexity Analysis

[Training step]: CNN/ResNet/InceptionTime require $O(\ell * |T| * D)$ computations per kernel, while dCNN/dResNet/dInceptionTime require $O(\ell * |T| * D^2)$ computations per kernel. Thus, the training time per epoch is higher for dCNN than CNN. However, given that the size of the input of dCNN is larger (containing D permutations of a single series) than CNN, the number of epochs to reach convergence is lower for dCNN when compared to CNN. Intuitively, dCNN trains on more data during a single epoch. This leads to similar overall training times (see Section 5.7).

[dCAM step]: The CAM computation complexity is $O(|T| * D * n_f)$, where n_f is the number of filters in the last convolutional layer. Let $N_f = [n_{f_1}, \dots, n_{f_n}]$ be the number of filters of the n convolutional layers. Then, a forward pass has time complexity $O(\ell * |T| * D^2 * \sum_{n_{f_i} \in N_f} n_{f_i})$. In dCAM, we evaluate k different permutations. Thus, the overall dCAM complexity is $O(k * \ell * |T| * D^2 * \sum_{n_{f_i} \in N_f} n_{f_i})$. Observe that since the k permutations can be computed in parallel, the most important parameter for the execution time is D .

4.6 Further Observations

We note that since in real use cases, labels are not available, the number of correctly classified permutations (called n_g) could be used as a proxy to assess the quality of the explanation (see Section 5.6).

Moreover, when analyzing sets of series, we can use dCAM on each one independently, and then aggregate the dCAM results to identify global discriminant features (see Section 5.8).

5 EXPERIMENTAL EVALUATION

5.1 Experimental Setup

We implemented our algorithms in Python 3.5 using the PyTorch library [48]. The evaluation was conducted on a server with Intel Core i7-8750H CPU 2.20GHz x 12, with 31.3GB RAM, and Quadro P1000/PCIe/SSE2 GPU with 4.2GB RAM, and on Jean Zay cluster with Nvidia Tesla V100 SXM2 GPU with 32 GB RAM.

Our code and datasets are available online [1].

5.1.1 Datasets. We conduct our experimental evaluation using real datasets from the UCR/UEA archive [12] to evaluate the classification performance of the competing methods. The real datasets are injected with known discriminant patterns and a real use case from the medical domain to evaluate the discriminant features identification. We use the StarLightCurves (classes 2 and 3 only), ShapesAll (classes 1 and 2 only), and Fish (class 1 and 2 only) datasets from the UCR archive [12], in which we inject subsequences that will generate discriminant features. We build two types of datasets to study the ability of the algorithms to identify the discriminant patterns guiding the classification decision, (1) when these patterns occur in a subset of the dimensions at different timestamps, and (2) when these patterns occur in a subset of the dimensions at the same timestamp.

(1) For the *Type 1* datasets, we build each dimension of Class 1 by concatenating random instances from one class of one of our two UCR seed datasets. We build Class 2 by injecting in the series of the other class of our two UCR datasets a pattern in 2 random dimensions at a random position in the series.

(2) For the *Type 2* datasets, we build each dimension of Class 1 by concatenating random instances from one of the classes of our two UCR datasets and injecting patterns from the other class in x random dimensions and at different positions. We build Class 2 by injecting patterns at the same positions of 2 random dimensions.

Examples of *Type 1* and *Type 2* 5-dimensional datasets based on StarLightCurves are depicted in Figures 7(a), and 7(b), respectively; we use 1000 such datasets. In addition, we consider a use case from medicine related to robot-assisted surgeon training (Section 5.8).

5.1.2 Evaluation Measures. We first evaluate the classification accuracy, $C\text{-acc}$. This measure corresponds to the ratio of correctly classified instances among all instances in the test dataset.

We then evaluate the discriminant features accuracy, $Dr\text{-acc}$, for Class 1 (see Figure 7). We define $Dr\text{-acc}$ as the PR-AUC for CAM/cCAM/dCAM obtained from the models and the ground-truth. The ground-truth is a series that has 1 at the positions of discriminant features (see Figure 7(a.2): ground-truth contains 1 at the positions of the injected patterns, marked with the red rectangles, and 0 otherwise). We motivate the choice of PR-AUC (instead of ROC-AUC) because we are more interested in measuring the accuracy of identifying the injected patterns (representing at max 0.02 percent of the dataset) than measuring the accuracy of not detecting the non-injected patterns. In this very unbalanced case, PR-AUC is more appropriate than ROC AUC [13].

Note that even though we annotate each point of the injected subsequences as discriminant, only some subparts of these sequences may be discriminant, thus, leading to $Dr\text{-acc}$ less than 1. Finally, for CNN/ResNet/InceptionTime, we compute the $Dr\text{-acc}$ scores by

assuming that their (univariate) CAM values are the same for all dimensions. We mark their $Dr\text{-acc}$ scores with a star in Table 2.

5.2 Baselines and Training Setup

We compare our model, dCNN/dResNet/dInceptionTime, to the classical CNN/ResNet/InceptionTime model [26–28, 65], and the cCNN/cResNet/cInceptionTime baseline we introduced in Section 2. We are using the same architecture setup for all models. We then use CAM for CNN, ResNet, InceptionTime, cCAM for cCNN, cResNet and cInceptionTime and dCAM for dCNN, dResNet and dInceptionTime to identify discriminant features. For CNN, cCNN and dCNN, we are using 5 convolutional layers with (64, 128, 256, 256, 256) filters respectively. We are using a kernel size of 3 and a padding of 2. For ResNet, cResNet, and dResNet, we are using three blocks with three convolutional layers of 64 filters (for the first two blocks) and 128 layers (for the last block). We are using kernel sizes equal to 8, 5, and 3 for each block for the three layers of the block. For InceptionTime, cInceptionTime and dInceptionTime, we are using the same architecture as originally defined [28].

We also include MTEX-CNN [2](MTEX) as a baseline, representative of other kinds of architectures that can provide a multivariate CAM. The explanation is computed separately for discriminant features and timestamps using grad-CAM [57] (MTEX-grad). The latter is a variant of the usual CAM using the gradients of the weights instead of the GAP layer to compute the activation.

We finally include three recurrent neural networks: the usual Recurrent Neural Network [53] (RNN), Long-Short Term Memory [23] (LSTM), and Gated Recurrent Unit [10] (GRU) to our benchmark. As following previous evaluation work conducted in the UCR/UEA archive [60], we use for all networks one recurrent hidden layer (RNN, LSTM, and GRU respectively) of 128 neurons. We then add one dense layer connecting the 128 neurons to the classes neurons.

We split our dataset into training and validation sets with 80 and 20 percent of the total dataset, respectively (equally balanced between the two classes). The training dataset is used to train the model, and the validation dataset is used as a validation dataset during the training phase. We generate a fully new test dataset for synthetic datasets and evaluate $C\text{-acc}$ and $Dr\text{-acc}$. We train all models with a learning rate $\alpha = 0.00001$, a maximum batch size of 16 instances (less if GPU memory cannot fit 16 instances), and a maximal number of epochs equal to 1000 (we use early stopping and stop before 1000 epochs if the model starts overfitting the test dataset). For dCAM, we use $k = 100$ (number of random permutations), a value that we empirically verified (due to lack of space, a detailed analysis of the effect of k is in the full version of the paper).

5.3 Classification Accuracy evaluation

We first evaluate the classification performance of our proposed approaches (denoted as c -Baselines and d -Baselines in Table 2) and the different baselines (denoted as Baselines in Table 2) over the UCR/UEA multivariate data series. We run each method ten times and report the average $C\text{-acc}$.

We first observe that the recurrent models (RNN, GRU, LSTM) are less accurate by approximately 0.10 than CNN-based models (CNN, ResNet and InceptionTime). These results confirm the observations of previous works [26, 27, 52, 65]. We then observe that

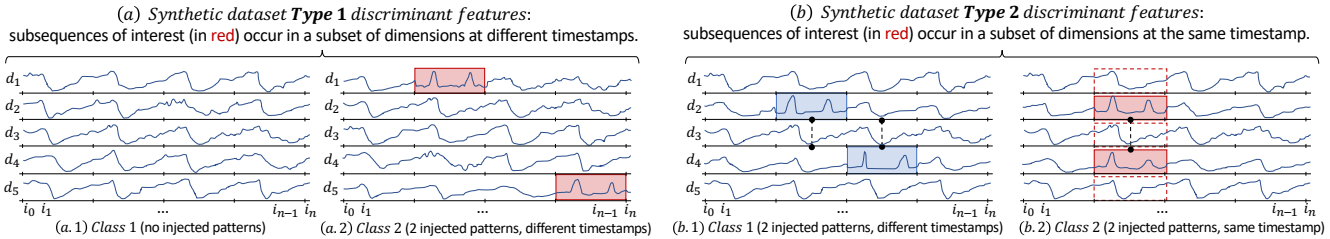


Figure 7: Synthetic datasets: (a) *Type 1*, in which the discriminant subsequence is two injected patterns from class 2 StarLightCurves dataset in random dimensions at random positions, (b) *Type 2*, in which the discriminant factor is the fact that the two injected patterns are injected at the same position.

Metadata				C-acc (averaged on 10 runs)												
Datasets name	C	T	D	Baselines							c-Baselines				d-Baselines	
				RNN	GRU	LSTM	MTEX	CNN	ResNet	InceptionT.	cCNN	cResNet	cInceptionT.	dCNN	dResNet	dInceptionT.
AtrialFibrillation	3	640	2	0.66	0.70	0.66	0.72	0.41	0.40	0.64	0.56	0.53	0.68	0.49	0.45	0.61
Libras	15	45	2	0.86	0.84	0.75	0.93	0.96	0.96	0.82	0.80	0.82	0.65	0.91	0.94	0.66
BasicMotions	4	100	2	0.68	0.87	0.83	0.91	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
RacketSports	4	30	6	0.70	0.78	0.75	0.81	0.94	0.99	0.90	0.95	0.98	0.85	0.94	0.98	0.92
Epilepsy	4	206	3	0.63	0.83	0.83	0.97	1.00	1.00	0.97	1.00	1.00	1.00	1.00	1.00	0.99
StandWalkJump	3	2500	4	0.60	0.80	0.90	0.66	0.70	0.66	0.65	0.83	1.00	0.81	0.95	1.00	0.75
UWaveGest.Lib.	8	315	3	0.88	0.93	0.88	0.93	0.88	0.89	0.89	0.76	0.74	0.64	0.84	0.89	0.83
Handwriting	26	152	3	0.45	0.43	0.40	0.34	0.83	0.90	0.55	0.42	0.70	0.38	0.76	0.89	0.52
NATOPS	6	51	24	0.87	0.91	0.82	0.91	0.99	1.00	0.95	0.86	0.89	0.83	0.97	0.99	0.91
PenDigits	10	8	2	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99
FingerMovements	2	50	28	0.56	0.58	0.58	0.62	0.70	0.68	0.71	0.57	0.63	0.55	0.72	0.71	0.66
Artic.WordRec.	25	144	9	0.98	0.97	0.90	0.98	0.99	0.99	0.93	0.82	0.94	0.74	0.98	0.99	0.88
HandMov.Dir.	4	400	10	0.46	0.40	0.37	0.44	0.44	0.42	0.51	0.34	0.35	0.40	0.45	0.44	0.33
Cricket	12	1197	6	0.98	0.90	0.80	0.95	1.00	1.00	0.98	0.94	0.97	0.87	1.00	1.00	0.98
LSST	14	36	6	0.54	0.53	0.52	0.53	0.62	0.66	0.40	0.56	0.59	0.49	0.62	0.66	0.51
Eth.Concentration	4	1751	3	0.33	0.37	0.32	0.58	0.35	0.36	0.34	0.36	0.36	0.36	0.35	0.39	0.34
SelfReg.SCP1	2	896	6	0.91	0.88	0.89	0.88	0.86	0.83	0.87	0.88	0.84	0.88	0.86	0.86	0.88
SelfReg.SCP2	2	1152	7	0.58	0.57	0.60	0.58	0.59	0.58	0.62	0.60	0.59	0.59	0.57	0.60	0.63
Heartbeat	2	405	61	0.73	0.73	0.73	0.75	0.83	0.86	0.83	0.76	0.76	0.76	0.84	0.86	0.83
PhonemeSpectra	39	217	39	0.11	0.11	0.11	0.15	0.31	0.37	0.27	0.31	0.33	0.28	0.33	0.40	0.32
EigenWorms	5	17984	6	0.57	0.50	0.66	0.57	0.90	0.92	0.82	0.71	0.92	0.73	0.92	0.92	0.81
MotorImagery	2	3000	64	0.53	0.59	0.58	0.59	0.58	0.57	0.56	0.56	0.57	0.56	0.65	0.68	0.66
FaceDetection	2	62	144	0.64	0.58	0.60	0.72	0.57	0.59	0.71	0.55	0.70	0.70	0.57	0.61	0.63
Mean				0.662	0.686	0.672	0.717	0.758	0.766	0.735	0.701	0.747	0.684	0.770	0.793	0.723
Rank				8.26	7.65	8.73	6.39	4.73	4.13	6.08	7.30	5.73	7.73	4.56	2.65	6.56

Table 2: C-acc averaged accuracy for 10 runs over UCR/UEA datasets.

ResNet-based architecture performs better than CNN-based and InceptionTime-based architectures. Moreover, we note that, overall, dCNN and dResNet have a better *C-acc* than CNN and ResNet, respectively. This observation confirms that our proposed architectures (dResNet, dCNN) do not result in any loss in accuracy; on the contrary, they are slightly more accurate than usual architectures (ResNet, CNN). We notice that dResNet is, on average, one rank higher than ResNet. Similar observations can be made when comparing dCNN and CNN.

Moreover, Table 2 confirms that using cCNN baselines (or cResNet and cInceptionTime) implies a drop in classification accuracy. For instance, CNN architecture is 0.05 more accurate than cCNN architecture. Thus, *c-Baselines* cannot guarantee at least equivalent accuracy. Figure 8(a) depicts the comparison between dCNN *C-acc* (on the y-axis) and CNN/cCNN *C-acc* (on the x-axis; CNN: blue circles; cCNN: red crosses). The dotted line corresponds to cases when both classifiers have the same accuracy. We observe that almost all cCNN *C-acc* (red crosses) are above the dotted line, which shows that dCNN is more accurate for most datasets. Similarly, we observe that most of the CNN *C-acc* (blue circles) are above the dotted lines, which means that dCNN is more accurate than CNN. The same observation can be made when examining Figure 8(b), in which dResNet is compared with ResNet and cResNet.

However, the same observation is not true when comparing dInceptionTime with InceptionTime and cInceptionTime. Even though in Figure 8(c) most of the red crosses are above the dotted line, indicating that dInceptionTime is most of the time more accurate than cInceptionTime, the blue circles are equally distributed above and under the dotted line. Thus, dInceptionTime is not more accurate than InceptionTime. The results in Table 2 also show that the averaged *C-acc* across all datasets (as well as the averaged rank) is lower for dInceptionTime than for InceptionTime. Nevertheless, the performance of dInceptionTime is very close to that of InceptionTime. Thus, transforming the original architecture into one that supports dCAM does not penalize classification performance.

Finally, we observe that the accuracy of MTEX-CNN is lower than that of *Baselines* and the *d-Baselines*. We note that MTEX-CNN and cCNN have very similar performance (average accuracy of 0.71 and 0.70, and average rankings of 6.39 and 7.30). As we explained earlier (see Section 2.3), the MTEX-CNN architecture is divided into two blocks. The experiments demonstrate that the 2nd block cannot capture all discriminant features, and thus, cannot reach the accuracy of a traditional CNN. We conclude that MTEX-CNN is not as accurate as traditional architectures (such as CNN) or our proposed architectures (such as dCNN).

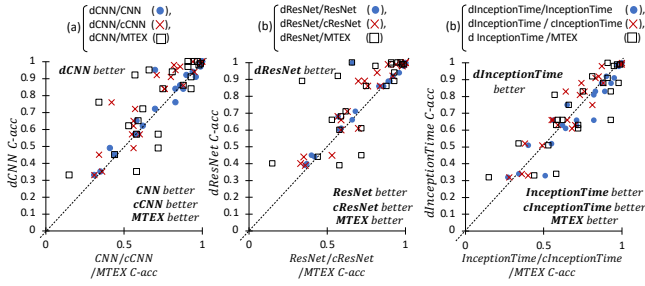


Figure 8: C -acc comparison of (a) dCNN with cCNN, CNN and MTEX, (b) dResNet with cResNet, ResNet and MTEX, and (c) dInceptionTime with cInceptionTime, InceptionTime and MTEX on UCR/UEA datasets.

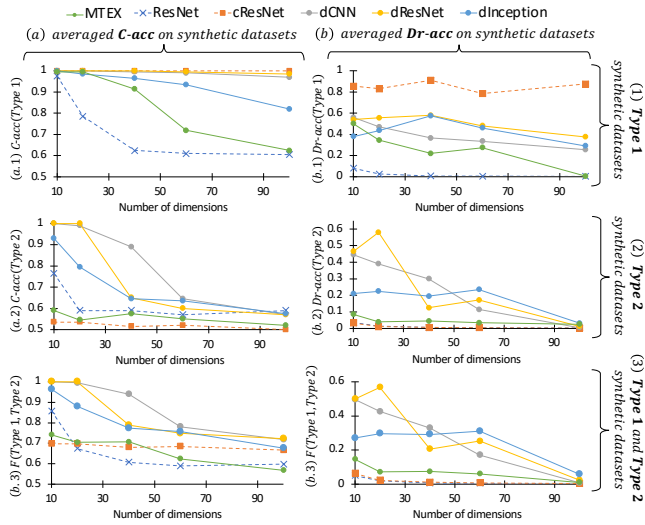


Figure 9: Evaluation of influence of number of dimensions on our approaches and the baselines C -acc and Dr -acc.

5.4 Discriminant Features Identification

We now evaluate the classification accuracy (C -acc) and the discriminant features identification accuracy (Dr -acc) on synthetically built datasets. Table 3 depicts both C -acc and Dr -acc on *Type 1* and *Type 2* datasets, when varying the number of dimensions from 10 to 100. In this experiment, we keep as baselines only ResNet and cResNet, which are the most accurate methods among all other baselines.

Overall, all methods have better performance (both C -acc and Dr -acc) on *Type 1* datasets than on *Type 2*. This was expected: discriminant features located in single dimensions are easier to find than discriminant features that depend on several dimensions.

We then notice that for low dimensional ($D=10$) datasets, ResNet, dResNet, dCNN, and dInceptionTime are performing nearly perfect C -acc. Moreover, ResNet and MTEX-CNN are performing well for low-dimensional data series but start to fail for a more significant number of dimensions. While the drop is already significant for the *Type 1* dataset built from the StarLightCurve dataset, it is even stronger for *Type 2* datasets, for which ResNet fails to classify instances with a number of dimensions $D \geq 20$. On the contrary, dCNN, dResNet, and dInceptionTime, which use the random permutations in the input, are not sensitive to the number of dimensions and have an almost perfect C -acc for most of *Type 1* datasets. We

observe a C -acc drop for dCNN, dResNet and dInceptionTime as dimensions increase for *Type 2* datasets. However, this drop is significantly less pronounced than that of ResNet. Overall, dCNN, dResNet, and dInceptionTime, which have on average the three highest ranks, are the most accurate methods.

Regarding cResNet, although it achieves a nearly perfect C -acc for *Type 1* datasets, we observe that it fails to classify correctly instances of *Type 2* datasets. As explained in Section 2, the input data structure is not rich enough to allow comparisons among dimensions, which is the main way to find discriminant features between the two classes of *Type 2* datasets. We also observe that MTEX-CNN fails to classify instances of *Type 2* datasets. Thus, this architecture does not correctly detect the discriminant features across different dimensions. Overall, Figure 9(a) shows that dCNN, dResNet and dInceptionTime are equivalent to cResNet for *Type 1* (Figure 9(a.1)), outperforming all the baselines for *Type 2* (Figure 9(a.2)), and in general are better than the baselines (ResNet and cResNet) for both types (Figure 9(a.3) with $F(\textit{Type 1}, \textit{Type 2}) = \frac{2 * C\text{-acc}(\textit{Type 1}) * C\text{-acc}(\textit{Type 2})}{C\text{-acc}(\textit{Type 1}) + C\text{-acc}(\textit{Type 2})}$).

We now compare the different methods using the Dr -acc measure. We observe that the baseline cCAM (computed with cCNN) is outperforming CAM (computed with ResNet) and dCAM (with all of dCNN, dResNet and dInceptionTime) for *Type 1* datasets. This is explained by the fact that these classes can be discriminated by treating dimensions independently. Thus, cCAM (with no comparisons between dimensions) is naturally the best solution. Nevertheless, as *Type 2* datasets require comparisons among dimensions to discriminate the classes, cCAM fails on them, with a Dr -acc very similar to the one of a random classifier. This confirms that such a baseline cannot be considered as a general solution for multivariate data series classification. We also observe that Dr -acc of the explanation method of MTEX-CNN (MTEX-grad) is lower than dCAM for *Type 1* and close to Dr -acc of cCAM for *Type 2*, meaning that it cannot identify discriminant features of *Type 2* datasets.

We then compare CAM and dCAM (used with dCNN/dResNet/dInceptionTime). Figure 9(b) shows that dCAM significantly outperforms CAM, and that Dr -acc reduces for all models as the number of dimensions increases. Nevertheless, Dr -acc of dCAM remains relatively high for both *Type 1* (Figure 9(b.1)) and *Type 2* (Figure 9(b.2)) datasets (for less than 60 dimensions).

This result demonstrates the superiority of dCAM over state-of-the-art methods. Besides, the average ranks in Table 3 indicate that dCAM computed from ResNet has the highest rank of 2.15.

5.5 Influence of k

This section analyzes the influence of the number of permutations k on the discriminative features identification accuracy (Dr -acc). We compute the Dr -acc for 20 different instances for which dCAM is computed using a value of k between 1 and 400. We randomly select the 20 instances from the 9 ShapesAll datasets, *Type 1* and *Type 2*. (We excluded the *Type 2* ShapesAll dataset with 100 dimensions, because no model trained on this dataset leads to reasonably accurate results: see Table 3.) Figure 10(a.1) for *Type 1* datasets and (a.2) for *Type 2* datasets depicts the evolution of Dr -acc (on average for the 20 instances) as k increases, for dCNN, dResNet and dInceptionTime. The results show that the model architecture

Datasets			C-acc (averaged on 10 runs)						Dr-acc (averaged on 50 instances)						
Dataset name	Type	Dimensions	MTEX	ResNet	cResNet	dCNN	dResNet	dInception	MTEX-grad	CAM	cCAM	dCNN	dResNet	dInception	Random
StarLightCurve	Type 1	10	0.99	0.95	1.00	1.00	1.00	1.00	0.40	0.07*	0.92	0.46	0.38	0.21	0.02
		20	0.99	0.71	1.00	1.00	1.00	0.98	0.38	0.02*	0.92	0.38	0.45	0.36	0.01
		40	0.98	0.60	1.00	0.99	1.00	0.93	0.24	0.008*	0.94	0.28	0.42	0.39	0.005
		60	0.61	0.57	1.00	0.98	0.99	0.91	0.05	0.004*	0.92	0.23	0.24	0.13	0.003
		100	0.55	0.64	1.00	0.96	0.97	0.79	0.01	0.003*	0.92	0.2	0.26	0.10	0.002
	Type 2	10	0.58	0.71	0.53	1.00	1.00	0.93	0.15	0.0256*	0.025	0.26	0.43	0.10	0.021
		20	0.55	0.61	0.55	0.98	1.00	0.70	0.04	0.016*	0.01	0.28	0.43	0.05	0.01
		40	0.56	0.58	0.51	0.88	0.58	0.56	0.07	0.0068*	0.006	0.20	0.05	0.03	0.005
		60	0.53	0.55	0.53	0.64	0.59	0.55	0.008	0.0058*	0.005	0.01	0.003	0.009	0.003
		100	0.52	0.59	0.5	0.59	0.56	0.60	0.01	0.0024*	0.002	0.003	0.004	0.02	0.002
ShapesAll	Type 1	10	1.00	1.00	1.00	1.00	1.00	1.00	0.60	0.09*	0.79	0.66	0.7	0.55	0.02
		20	1.00	0.86	1.00	1.00	1.00	0.99	0.31	0.03*	0.74	0.56	0.66	0.51	0.011
		40	0.85	0.65	1.00	1.00	1.00	1.00	0.20	0.008*	0.88	0.45	0.74	0.76	0.005
		60	0.83	0.65	1.00	1.00	1.00	0.96	0.50	0.005*	0.65	0.44	0.72	0.79	0.003
		100	0.70	0.57	1.00	0.98	1.00	0.85	0.002	0.003*	0.83	0.31	0.49	0.48	0.002
	Type 2	10	0.60	0.82	0.54	1.00	1.00	0.93	0.02	0.0467*	0.04	0.63	0.50	0.32	0.02
		20	0.54	0.57	0.52	1.00	1.00	0.89	0.04	0.0132*	0.013	0.50	0.73	0.40	0.01
		40	0.59	0.60	0.52	0.90	0.72	0.73	0.02	0.005*	0.005	0.40	0.20	0.36	0.005
		60	0.57	0.59	0.51	0.65	0.61	0.72	0.06	0.0037*	0.003	0.22	0.34	0.46	0.003
		100	0.52	0.59	0.50	0.55	0.58	0.55	0.04	0.0027*	0.002	0.005	0.02	0.05	0.002
Rank			3.95	3.9	3	1.65	1.6	2.85	3.85	4.45	3	2.6	2.15	2.75	

Table 3: C-acc and Dr-acc averaged accuracy for 10 runs over synthetic datasets.

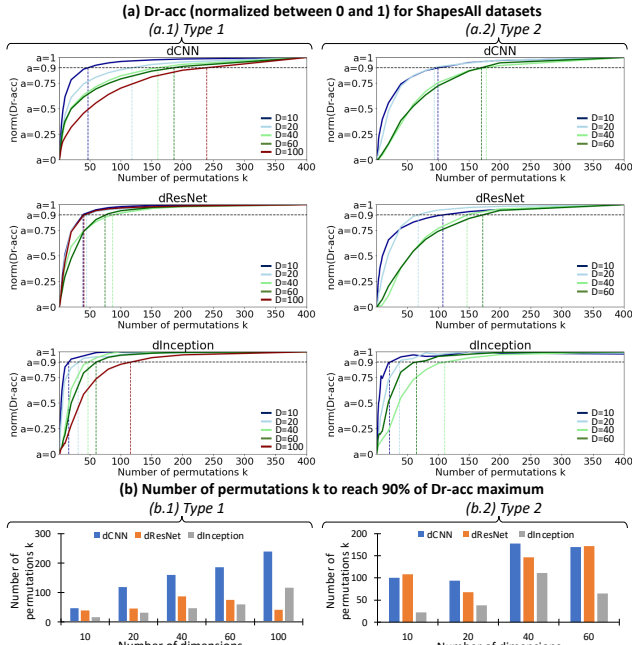


Figure 10: Influence of k on $Dr\text{-acc}$ for ShapesAll datasets.

influences convergence speed, and that convergence speed reduces as the number of dimensions increases. Figure 10(b) shows that the number of permutations needed to reach 90 percents of the best $Dr\text{-acc}$ is greater when D is higher. The latter holds for *Type 1* (Figure 10(b.1)) and *Type 2* datasets (Figure 10(b.2)). Overall, we notice that the dCAM computation with dResNet and dInceptionTime converges faster than dCNN. Studying deep neural network architectures that could reduce the number of permutations needed to reach the maximum $Dr\text{-acc}$ is an open research problem.

5.6 C-acc versus Dr-acc

In this section, we first analyze the relation between $C\text{-acc}$ and $Dr\text{-acc}$. We then evaluate the impact that $C\text{-acc}$ has on the number of permutations that have been correctly classified n_g . We finally evaluate the impact that n_g has on $Dr\text{-acc}$.

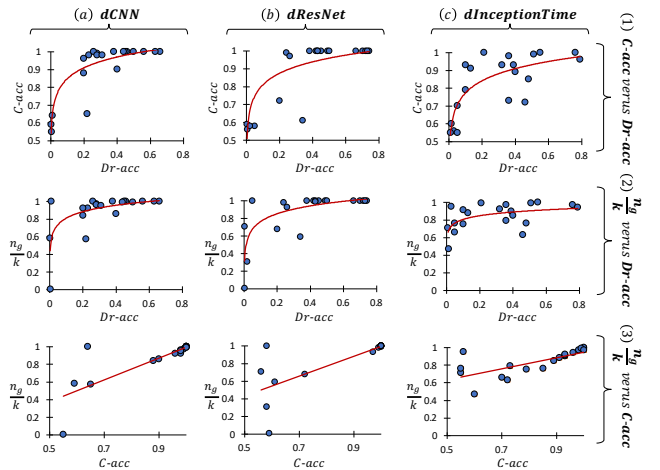


Figure 11: Evaluation of $C\text{-acc}$, $Dr\text{-acc}$, and ratio between number of permutations k and number of permutations correctly classified n_g , for dCNN, dResNet and dInceptionTime.

Figure 11(1) depicts the relation between $C\text{-acc}$ and $Dr\text{-acc}$ for dCNN (Figure 11(a.1)), dResNet (Figure 11(b.1)) and dInceptionTime (Figure 11(c.1)) for all synthetic datasets. Note that all methods have a logarithmic relation (dotted red line) between $Dr\text{-acc}$ (x-axis) and $C\text{-acc}$ (y-axis). This confirms that the accuracy of the trained model has a significant impact on discriminant feature identification.

Figure 11(3) depicts on the y-axis the ratio of correctly classified permutations (n_g) among all permutations (k) versus the $C\text{-acc}$ (on the x-axis). In this case, for all of dCNN (Figure 11(a.3)), dResNet (Figure 11(b.3)) and dInceptionTime (Figure 11(c.3)), we observe that there exists a linear relationship for $C\text{-acc}$ between 0.7 and 1. This means that n_g will be greater when the model is more accurate. Nevertheless, for $C\text{-acc}$ between 0.5 and 0.7, we observe a high variance for n_g/k . Thus, an inaccurate model may still lead to a high n_g . Finally, Figure 11(2) depicts the relation between n_g/k (on the y-axis) and $Dr\text{-acc}$ (on the x-axis). We observe a similar relationship between $C\text{-acc}$ and $Dr\text{-acc}$, which means that a low n_g may lead to inaccurate discriminant features identification.

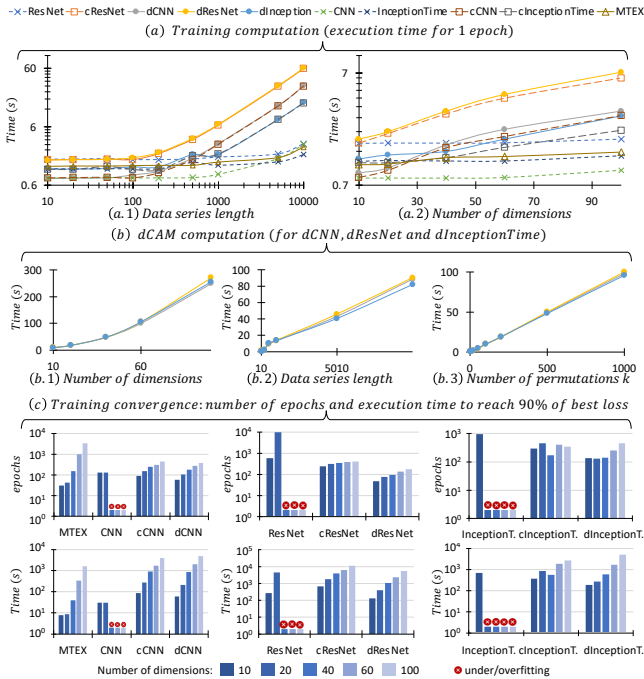


Figure 12: Execution time (seconds) for (a) training computations when we vary (a.1) data series length and (a.2) number of dimensions. (b) Execution time for dCAM computation when we vary (b.1) number of dimensions, (b.2) data series length, (b.3) number of permutations k , (c) training time to reach convergence (90% of the best loss).

As hypothesized in Section 4, the experimental results confirm that an inaccurate model (for all of dCNN, dResNet, and dInceptionTime) cannot be used to identify discriminant features. Moreover, since in a real use-case it is not possible to measure $Dr-acc$, we can use n_g/k to estimate the discriminant feature identification accuracy. Even though Figure 11(2) demonstrates that a high n_g/k does not always lead to a high $Dr-acc$, in practice, we can safely assume that a low n_g/k will most probably correspond to a low $Dr-acc$. Therefore, such measure can be used as a proxy for the estimation of the explanation quality.

5.7 Execution time evaluation

In this section, we evaluate the execution time of our proposed approaches and the baselines. Figure 12(a) depicts the training execution time (for one epoch) when we vary the data series length with a constant number of dimensions fixed to 10 (Figure 12(a.1)), and when we vary the number of dimensions with a constant data series length fixed to 100 (Figure 12(a.2)). In these two experiments, we use a batch size of 4 for all models. Overall, CNN and InceptionTime-based architectures are faster than ResNet-based architectures, and CNN, ResNet, InceptionTime, and MTEX-CNN are faster when the number of dimensions and the data series length is increasing. Nevertheless, both dCNN/dResNet/dInceptionTime and cCNN/cResNet/cInceptionTime require the same training time.

We now evaluate the execution time and the number of epochs required to train our proposed approaches and the baselines.

Figure 12(c) depicts the time (in seconds) and the number of epochs to reach 90% of the best loss (on the test set) for *Type 1* ShapesAll datasets varying the number of dimensions between 10 and 100. We use for all models a batch size of 16. The red dot indicates that a model is either overfitted or underfitted (i.e., the loss for the first epoch is approximately equal to the best loss). We observe that cCNN/cResNet/cInceptionTime and dCNN/dResNet/dInceptionTime require the same amount of time to be trained, but traditional baselines require more epochs than the proposed d-methods. Thus, training time for ResNet is longer than dResNet for $D = 10$ and $D = 20$.

Finally, we measure the execution time required to compute dCAM (for dCNN, dResNet, and dInceptionTime), when we vary the number of dimensions with a constant data series length fixed to 400 (Figure 12(b.1)), when we vary the data series length with a constant number of dimensions fixed to 10 (Figure 12(b.2)), and when we vary the number of permutations (Figure 12(b.3)). Note that the dCAM execution times are very similar for the three types of architectures. Moreover, the execution time increases super-linearly with the number of dimensions but is linear to the data series length and the number of permutations k .

5.8 Use Case: Surgeon skills explanation

We now illustrate the applicability of our method to a real-world use case. In this use case, we train our dCNN network on the JIGSAWS dataset [21] to identify novice surgeons, based on kinematic data series when performing *surgical suturing* tasks (i.e., wound stitching) using robotic arms and surgical grippers.

[Dataset] The data series are recorded from the *DaVinciSurgicalSystem*. The multivariate data series are composed of 76 dimensions (an example of multivariate data series is depicted in Figure 13(a)). Each dimension corresponds to a sensor (with an acquisition rate of 30 Hz). The sensors are divided into four groups: patient-side manipulators (left and right PSMs: green rectangle in Figure 13(a) top left), and left and right master tool manipulators (left and right MTMs: blue rectangle in Figure 13(a) bottom left). Each group contains 19 sensors. These sensors are: 3 variables for the Cartesian position of the manipulator, 9 variables for the rotation matrix, 6 variables for the linear and angular velocity of the manipulator, and 1 variable for the gripper angle.

To perform a suture, the surgeons perform different gestures (11 in total). For example, $G1$ refers to reaching for the needle with the right hand, while $G11$ refers to dropping the suture at the end and moving to end points. Each gesture corresponds to a specific time segment of the dataset, involving all sensors. For example, the dotted red rectangle in Figure 13(a) represents gesture $G6$: pulling the suture with the left hand. Surgeons that reported having more than 100 hours of experience are considered experts, surgeons with 10-100 hours are considered intermediate, and surgeons with less than 10 hours are labeled as novices. We have 19 multivariate data series in the novice class, denoted as C_N , 10 in the intermediate, C_i , 10 multivariate data series in the expert class, C_E . More information on this dataset can be found in [21].

[Training] For the training procedure, we use 80% of the dataset (randomly selected from the three classes) for training. The rest 20% of the dataset is used for validation and early stopping. Since the

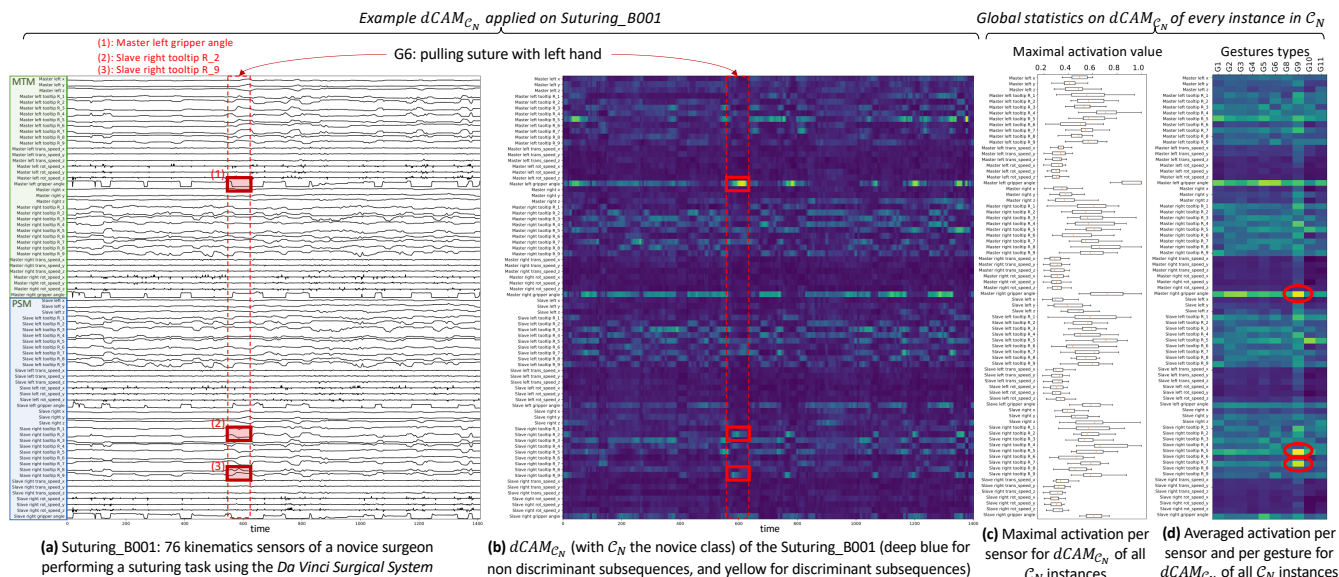


Figure 13: Example of the result of dCAM on a multivariate data series of the JIGSAWS dataset. (a) the depicted data series corresponds to a novice surgeon performing a suture operation (joined with the corresponding dCAM in (b)). General statistical results over the entire novice class C_N are depicted such as (c) box-plots of the maximal activation value per sensor and (d) the averaged activation per sensor per gesture performed.

instances do not have the same length, we use batches composed of one instance when training the models in the GPU.

[Evaluation] Similar to what has been reported in previous work [26], we achieve 100% accuracy on the train and test datasets (for ten different randomly selected train and test sets). We proceed to compute the $dCAM_{C_N}$ for every instance of the novice class C_N . The $dCAM_{C_N}$ of the multivariate data series named *Suturing_B001* (Figure 13(a)) is displayed in Figure 13(b). In the latter, the deep blue color indicates low activated subsequences (i.e., non-discriminant of belonging to the novice class C_N), while the yellow color is pointing to highly activated sensors. First, we note that some groups of sensors (dimensions) are more activated than others. In Figure 13(b), the left and right "MTM gripper angles" are the most activated sensors. Figure 13(c), which depicts the box-plot of the maximal activated values per sensor, confirms that in the general case, MTM gripper angles, as well as the MTM and PSM tooltip rotation matrices (three of these sensors are highlighted in red in Figure 13(a)), are the most discriminant sensors. On the contrary, linear and angular speeds are not discriminant and hence cannot explain the novice class C_N .

As explained in Section 4.6, we now extract global explanations at the scale of the dataset. We compute the dCAM for each instance, and we then extract global statistics on the sensors, i.e., aggregated over all instances. Figure 13(d) depicts the averaged activation per sensor per gesture. Overall, $dCAM_{C_N}$ identifies gesture *G9* (using the right hand to help tighten the suture) as a discriminant gesture, because of the discriminant subsequences present in the sensors "right MTM gripper angle", "5th element", and "7th element" (marked with red ovals in Figure 13(d)). These three identified sensors (dimensions) are relevant to the right PSM tooltip rotation matrix and are important for the suturing process. Similarly, we

observe that gesture *G6* (i.e., pulling suture with left hand) is discriminant, and activated the most by the "left NTM gripper angle" sensor. This result is consistent with a previous study [26], which also identified gesture *G6* as discriminant of belonging to the novice class. Nevertheless, this previous study was using CAM to only highlight the *time interval* corresponding to gesture *G6*. On the contrary, dCAM provides more accurate (and useful) information: it does not only identify the discriminant gesture *G6*, but also the *discriminant sensors*. This allows the analysts to recognize exactly what aspects of the particular gesture are problematic.

[Summary] The application of dCAM in the robot-assisted surgeon training use case demonstrated its effectiveness. Our approach was able to provide meaningful explanations for the classification decisions, based on specific gestures (subsequences), and specific sensors (dimensions) that describe particular aspects of these gestures, i.e., the positioning and rotation angles of the tip of the stitch gripper. Such explanations can help surgeons to improve their skills.

6 CONCLUSIONS

Even though data series classification using deep learning has attracted a lot of attention, existing techniques for explaining the classification decisions fail for the case of multivariate data series. We described a novel approach, dCAM, based on CNNs, which detects discriminant subsequences within individual dimensions of a multivariate data series. The experimental evaluation with synthetic and real datasets demonstrates the superiority of our approach.

Acknowledgments Work supported by EDF R&D and ANRT French program, and HPC resources from GENCI-IDRIS (Grants 2020-101471 and 2021-101925), and NVIDIA Corporation for the Titan Xp GPU donation used in this research.

REFERENCES

- [1] 2022. *dCAM Webpage*. <https://helios2.mi.parisdescartes.fr/~themisp/dCAM/>
- [2] Roy Assaf, Ioana Giurgiu, Frank Bagehorn, and Anika Schumann. 2019. MTEX-CNN: Multivariate Time Series EXplanations for Predictions with Convolutional Neural Networks. In *2019 IEEE International Conference on Data Mining (ICDM)*. 952–957. <https://doi.org/10.1109/ICDM.2019.00106>
- [3] Anthony Bagnall, Jason Lines, Aaron Bostrom, James Large, and Eamonn Keogh. 2016. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Min. Knowl. Discov.* 31 (2016).
- [4] Anthony Bagnall, Jason Lines, Jon Hills, and Aaron Bostrom. 2015. Time-Series Classification with COTE: The Collective of Transformation-Based Ensembles. *IEEE TKDE* 27 (2015).
- [5] Anthony J. Bagnall, Richard L. Cole, Themis Palpanas, and Konstantinos Zoumpatianos. 2019. Data Series Management (Dagstuhl Seminar 19282). *Dagstuhl Reports* 9, 7 (2019), 24–39.
- [6] Paul Boniol, Michele Linardi, Federico Roncallo, Themis Palpanas, Mohammed Meftah, and Emmanuel Remy. 2021. Unsupervised and Scalable Subsequence Anomaly Detection in Large Data Series. *VLDBJ* (2021).
- [7] Paul Boniol and Themis Palpanas. 2020. Series2Graph: Graph-based Subsequence Anomaly Detection for Time Series. *Proc. VLDB Endow.* 13, 11 (2020), 1821–1834.
- [8] Paul Boniol, John Paparrizos, Themis Palpanas, and Michael J. Franklin. 2021. SAND: Streaming Subsequence Anomaly Detection. *Proc. VLDB Endow.* 14, 10 (2021), 1717–1729.
- [9] Huanhuan Chen, Fengzhen Tang, Peter Tino, and Xin Yao. 2013. Model-Based Kernel for Efficient Time Series Analysis. In *ACM SIGKDD*.
- [10] Kyunghyun Cho, Bart van Merriënboer, Çaglar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP*. ACL, 1724–1734.
- [11] Zhicheng Cui, Wenlin Chen, and Yixin Chen. 2016. Multi-Scale Convolutional Neural Networks for Time Series Classification. *CoRR* (2016).
- [12] H. A. Dau, A. Bagnall, K. Kamgar, C. M. Yeh, Y. Zhu, S. Gharghabi, C. A. Ratanamahatana, and E. Keogh. 2019. The UCR time series archive. *IEEE/CAA J. Automatic.* 6, 6 (2019).
- [13] Jesse Davis and Mark Goadrich. 2006. The Relationship between Precision-Recall and ROC Curves. In *Proceedings of the 23rd International Conference on Machine Learning (Pittsburgh, Pennsylvania, USA) (ICML '06)*. Association for Computing Machinery, New York, NY, USA, 233–240. <https://doi.org/10.1145/1143844.1143874>
- [14] Angus Dempster, François Petitjean, and Geoffrey I. Webb. 2020. ROCKET: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Min. Knowl. Discov.* 34, 5 (2020), 1454–1495.
- [15] Karima Echihabi, Panagiota Fatourou, Kostas Zoumpatianos, Themis Palpanas, and Houda Benbrahim. 2022. Hercules Against Data Series Similarity Search. *PVLDB* (2022).
- [16] Karima Echihabi, Kostas Zoumpatianos, Themis Palpanas, and Houda Benbrahim. 2018. The Lernaean Hydra of Data Series Similarity Search: An Experimental Evaluation of the State of the Art. *Proc. VLDB Endow.* 12, 2 (2018), 112–127.
- [17] Karima Echihabi, Kostas Zoumpatianos, Themis Palpanas, and Houda Benbrahim. 2019. Return of the Lernaean Hydra: Experimental Evaluation of Data Series Approximate Similarity Search. *Proc. VLDB Endow.* 13, 3 (2019), 403–420.
- [18] Philippe Esling and Carlos Agon. 2012. Time-series data mining. *CSUR* (2012).
- [19] Yifeng Gao and Jessica Lin. 2019. HIME: discovering variable-length motifs in large-scale time series. *Knowl. Inf. Syst.* 61, 1 (2019), 513–542.
- [20] Yifeng Gao, Jessica Lin, and Constantin Brif. 2020. Ensemble Grammar Induction For Detecting Anomalies in Time Series. In *Proceedings of the 23rd International Conference on Extending Database Technology, EDBT 2020, Copenhagen, Denmark, March 30 - April 02, 2020*. OpenProceedings.org, 85–96.
- [21] Yixin Gao, S. Vedula, Carol E. Reiley, N. Ahmidi, B. Varadarajan, Henry C. Lin, L. Tao, L. Zappella, B. Béjar, D. Yuh, C. C. Chen, R. Vidal, S. Khudanpur, and Gregory Hager. 2014. JHU-ISI Gesture and Skill Assessment Working Set (JIGSAWS) : A Surgical Activity Dataset for Human Motion Modeling.
- [22] Anna Gogolou, Theophanis Tsandilas, Karima Echihabi, Anastasia Bezerianos, and Themis Palpanas. 2020. Data Series Progressive Similarity Search with Probabilistic Quality Guarantees. In *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*. ACM, 1857–1873.
- [23] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [24] Tsung-Yu Hsieh, Suhang Wang, Yiwei Sun, and Vasant Honavar. 2021. Explainable Multivariate Time Series Classification: A Deep Neural Network Which Learns to Attend to Important Variables As Well As Time Intervals. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (Virtual Event, Israel) (WSDM '21)*. Association for Computing Machinery, New York, NY, USA, 607–615. <https://doi.org/10.1145/3437963.3441815>
- [25] Aya Abdelsalam Ismail, Mohamed K. Gunady, Héctor Corrada Bravo, and Soheil Feizi. 2020. Benchmarking Deep Learning Interpretability in Time Series Predictions. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/47a3893cc405396a5c30d91320572d6d-Abstract.html>
- [26] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2018. Evaluating Surgical Skills from Kinematic Data Using Convolutional Neural Networks. In *MICCAI*.
- [27] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep Learning for Time Series Classification: A Review. *Data Min. Knowl. Discov.* 33, 4 (2019).
- [28] Hassan Ismail Fawaz, Benjamin Lucas, Germain Forestier, Charlotte Pelletier, Daniel F. Schmidt, Jonathan Weber, Geoffrey I. Webb, Lhassane Idoumghar, Pierre Alain Muller, and François Petitjean. 2020. InceptionTime: finding AlexNet for time series classification. *Data Mining and Knowledge Discovery* 34 (7 Sept. 2020), 1936–1962.
- [29] Luka Jakovljevic, Dimitre Kostadinov, Armen Aghasaryan, and Themis Palpanas. 2022. Towards Building a Digital Twin of Complex System Using Causal Modelling. In *Complex Networks & Their Applications X*. Springer International Publishing, Cham, 475–486.
- [30] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*
- [31] A. Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet classification with deep convolutional neural networks. *Commun. ACM* 60 (2012).
- [32] Arthur Le Guennec, Simon Malinowski, and Romain Tavenard. 2016. Data Augmentation for Time Series Classification using Convolutional Neural Networks. In *ECML/PKDD on AALTD Workshop*.
- [33] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521 (2015).
- [34] Oleksandra Levchenko, Boyan Kolev, Djamel Edine Yagoubi, Reza Akbarinia, Florent Masseglia, Themis Palpanas, Dennis E. Shasha, and Patrick Valduriez. 2021. BestNeighbor: efficient evaluation of kNN queries on large time series databases. *Knowl. Inf. Syst.* 63, 2 (2021), 349–378.
- [35] Xiaosheng Li, Jessica Lin, and Liang Zhao. 2019. Linear Time Complexity Time Series Clustering with Symbolic Pattern Forest. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, Sarit Kraus (Ed.). ijcai.org, 2930–2936.
- [36] Michele Linardi, Yan Zhu, Themis Palpanas, and Eamonn J. Keogh. 2020. Matrix profile goes MAD: variable-length motif and discord discovery in data series. *Data Min. Knowl. Discov.* 34, 4 (2020), 1022–1071.
- [37] J. Lines, S. Taylor, and A. Bagnall. 2016. HIVE-COTE: The Hierarchical Vote Collective of Transformation-Based Ensembles for Time Series Classification. In *IEEE ICDM*.
- [38] Benjamin Lucas, Ahmed Shifaz, Charlotte Pelletier, Lachlan O'Neill, Nayyar A. Zaidi, Bart Goethals, François Petitjean, and Geoffrey I. Webb. 2019. Proximity Forest: an effective and scalable distance-based classifier for time series. *Data Min. Knowl. Discov.* 33, 3 (2019), 607–635.
- [39] Vinod Nair and Geoffrey E. Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *ICML*.
- [40] Themis Palpanas. 2015. Data Series Management: The Road to Big Sequence Analytics. *SIGMOD Rec.* 44, 2 (2015), 47–52.
- [41] Themis Palpanas. 2016. Big Sequence Management: A glimpse of the Past, the Present, and the Future. In *SOFSEM (Lecture Notes in Computer Science)*, Rusins Martins Freivalds, Gregor Engels, and Barbara Catania (Eds.), Vol. 9587. Springer, 63–80.
- [42] Themis Palpanas. 2019. Evolution of a Data Series Index. In *ISIP*.
- [43] Themis Palpanas and Volker Beckmann. 2019. Report on the First and Second Interdisciplinary Time Series Analysis Workshop (ITISA). *SIGMOD Rec.* 48, 3 (2019), 36–40.
- [44] John Paparrizos and Michael J. Franklin. 2019. GRAIL: Efficient Time-Series Representation Learning. *Proc. VLDB Endow.* 12, 11 (2019), 1762–1777.
- [45] John Paparrizos and Luis Gravano. 2016. k-Shape: Efficient and Accurate Clustering of Time Series. *SIGMOD Rec.* 45, 1 (2016), 69–76.
- [46] John Paparrizos and Luis Gravano. 2017. Fast and Accurate Time-Series Clustering. *ACM Trans. Database Syst.* 42, 2 (2017), 8:1–8:49.
- [47] John Paparrizos, Yuhao Kang, Paul Boniol, Ruy S. Tsay, Themis Palpanas, and Michael J. Franklin. 2022. TSB-UAD: An End-to-End Benchmark Suite for Univariate Time-Series Anomaly Detection. *PVLDB* (2022).
- [48] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*, Vol. 32.
- [49] Botao Peng, Panagiota Fatourou, and Themis Palpanas. 2021. Fast Data Series Indexing for In-Memory Data. *VLDBJ* (2021).

- [50] Botao Peng, Panagiota Fatourou, and Themis Palpanas. 2021. ParIS+: Data Series Indexing on Multi-Core Architectures. *IEEE Trans. Knowl. Data Eng.* 33, 5 (2021), 2151–2164.
- [51] Botao Peng, Panagiota Fatourou, and Themis Palpanas. 2021. SING: Sequence Indexing Using GPUs. In *ICDE*.
- [52] Alejandro Pasos Ruiz, Michael Flynn, James Large, Matthew Middlehurst, and Anthony J. Bagnall. 2021. The great multivariate time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Min. Knowl. Discov.* 35, 2 (2021), 401–449. <https://doi.org/10.1007/s10618-020-00727-3>
- [53] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. 1986. Learning representations by back-propagating errors. *Nature* 323 (1986), 533–536.
- [54] Patrick Schäfer and Ulf Leser. 2017. Fast and Accurate Time Series Classification with WEASEL. In *Proceedings of the 2017 ACM Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 - 10, 2017*. ACM, 637–646.
- [55] Patrick Schäfer and Ulf Leser. 2020. TEASER: early and accurate time series classification. *Data Min. Knowl. Discov.* 34, 5 (2020), 1336–1362.
- [56] J. Schneider, P. Wenig, and T. Papenbrock. 2021. Distributed detection of sequential anomalies in univariate time series. *VLDBJ* 30 (2021).
- [57] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 618–626. <https://doi.org/10.1109/ICCV.2017.74>
- [58] J. Serrà, S. Pascual, and Alexandros Karatzoglou. 2018. Towards a universal neural network encoder for time series. In *CCIA*.
- [59] Ahmed Shifaz, Charlotte Pelletier, François Petitjean, and Geoffrey I. Webb. 2020. TS-CHIEF: a scalable and accurate forest algorithm for time series classification. *Data Min. Knowl. Discov.* 34, 3 (2020), 742–775.
- [60] Denis Smirnov and Engelbert Mephu Nguifo. 2018. Time Series Classification with Recurrent Neural Networks. In *ECML/PKDD Workshop on Advanced Analytics and Learning on Temporal Data (aalt18)*.
- [61] Chang Wei Tan, François Petitjean, and Geoffrey I. Webb. 2020. FastEE: Fast Ensembles of Elastic Distances for time series classification. *Data Min. Knowl. Discov.* 34, 1 (2020), 231–272.
- [62] Liudmila Ulanova, Nurjahan Begum, and Eamonn J. Keogh. 2015. Scalable Clustering of Time Series with U-Shapelets. In *Proceedings of the 2015 SIAM International Conference on Data Mining, Vancouver, BC, Canada, April 30 - May 2, 2015*, Suresh Venkatasubramanian and Jieping Ye (Eds.). SIAM, 900–908.
- [63] J. Wang, Z. Wang, Jianfeng Li, and J. Wu. 2018. Multilevel Wavelet Decomposition Network for Interpretable Time Series Analysis. *ACM SIGKDD* (2018).
- [64] Qitong Wang and Themis Palpanas. 2021. Deep Learning Embeddings for Data Series Similarity Search. In *SIGKDD*.
- [65] Z. Wang, W. Yan, and T. Oates. 2017. Time series classification from scratch with deep neural networks: A strong baseline. In *IJCNN*.
- [66] B. Xu, Naiyan Wang, T. Chen, and Mu Li. 2015. Empirical Evaluation of Rectified Activations in Convolutional Network. *Deep Learning Workshop, ICML* (2015).
- [67] Qiang Yang and Xindong Wu. 2006. 10 Challenging Problems in Data Mining Research. *IJITDM* 05, 04 (2006).
- [68] Dragomir Yankov, Eamonn J. Keogh, and Umaa Rebbapragada. 2008. Disk aware discord discovery: finding unusual time series in terabyte sized datasets. *Knowl. Inf. Syst.* 17, 2 (2008), 241–262.
- [69] Lexiang Ye and Eamonn J. Keogh. 2011. Time series shapelets: a novel technique that allows accurate, interpretable and fast classification. *Data Min. Knowl. Discov.* 22, 1-2 (2011), 149–182.
- [70] Chuxu Zhang, Dongjin Song, Yuncong Chen, Xinyang Feng, Cristian Lumezanu, Wei Cheng, Jingchao Ni, Bo Zong, Haifeng Chen, and Nitesh V. Chawla. 2019. A Deep Neural Network for Unsupervised Anomaly Detection and Diagnosis in Multivariate Time Series Data. *AAAI* 33, 01 (2019).
- [71] Xuchao Zhang, Yifeng Gao, Jessica Lin, and Chang-Tien Lu. 2020. TapNet: Multivariate Time Series Classification with Attentional Prototypical Network. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 04 (Apr. 2020), 6845–6852. <https://doi.org/10.1609/aaai.v34i04.6165>
- [72] B. Zhao, H. Lu, S. Chen, J. Liu, and D. Wu. 2017. Convolutional neural networks for time series classification. *J. Syst. Eng. Electron.* 28, 1 (2017).
- [73] Yudi Zheng, Qi Liu, Enhong Chen, Yong Ge, and J. Leon Zhao. 2014. Time Series Classification Using Multi-Channels Deep Convolutional Neural Networks. In *WAIM*.
- [74] B. Zhou, A. Khosla, Lapedriza, A., A. Oliva, and A. Torralba. 2016. Learning Deep Features for Discriminative Localization. *CVPR* (2016).
- [75] Yan Zhu, Abdullah Mueen, and Eamonn J. Keogh. 2021. Matrix Profile IX: Admissible Time Series Motif Discovery With Missing Data. *IEEE Trans. Knowl. Data Eng.* 33, 6 (2021), 2616–2626.