

Time Series Kernel Similarities for Predicting Paroxysmal Atrial Fibrillation from ECGs

Filippo Maria Bianchi^{†*}, Lorenzo Livi[§], Alberto Ferrante[†], Jelena Milosevic[‡], Mirosław Malek[†]

[†]*ALaRI, Faculty of Informatics, Università della Svizzera italiana, Lugano, Switzerland*

^{*}*Machine Learning Group, UiT the Arctic University of Norway, Tromsø, Norway*

[‡]*Institute of Telecommunications, TU Wien, Vienna, Austria*

[§]*Dept. of Computer Science, University of Exeter, Exeter, UK*

*Email: *filippo.m.bianchi@uit.no, §l.livi@exeter.ac.uk,*

[†]{alberto.ferrante,mirosław.malek}@usi.ch, [‡]jelena.milosevic@tuwien.ac.at

Abstract—We tackle the problem of classifying Electrocardiography (ECG) signals with the aim of predicting the onset of Paroxysmal Atrial Fibrillation (PAF). Atrial fibrillation is the most common type of arrhythmia, but in many cases PAF episodes are asymptomatic. Therefore, in order to help diagnosing PAF, it is important to design procedures for detecting and, more importantly, predicting PAF episodes. We propose a method for predicting PAF events whose first step consists of a feature extraction procedure that represents each ECG as a multi-variate time series. Successively, we design a classification framework based on kernel similarities for multi-variate time series, capable of handling missing data. We consider different approaches to perform classification in the original space of the multi-variate time series and in an embedding space, defined by the kernel similarity measure. We achieve a classification accuracy comparable with state of the art methods, with the additional advantage of detecting the PAF onset up to 15 minutes in advance.

Index Terms—Atrial fibrillation; time series; kernel methods; classification; prediction; feature selection.

I. INTRODUCTION

Wearable devices are nowadays able to capture bio-signals, such as electrocardiograms (ECGs), for an extended period of time. Data recorded by these devices is of paramount clinical importance in the assessment of numerous heart-related conditions. Among them, the prediction of Paroxysmal Atrial Fibrillation (PAF) episodes and the risk stratification of PAF-prone patients is of high importance [1]–[3]. Atrial fibrillation is the most common type of arrhythmia. Its symptoms include fatigue or decreased exercise tolerance, palpitations, dyspnea on exertion, and generalized weakness, even though in many cases PAF episodes are asymptomatic [4]. The risk of PAF progressing to permanent atrial fibrillation increases with time.

The problem of detecting and predicting PAF events from ECG recordings received attention from researchers from different fields [5]. Many approaches have been proposed to tackle such problems, including the use of entropy descriptors [6], [7], P-wave characterization [8]–[10], the number of premature atrial complexes from R-R intervals [11] (i.e., the interval between R peaks, shown in Fig. 2), heart variability rate [12], the average number of f-waves in a TQ interval [13], and hybrid complex networks / time series analysis techniques [14], [15].

In this work, we propose a new framework based on time series analysis methods to predict the onset of PAF. The overall procedure consists of different steps. First, from each ECG we extract numerical descriptors (features) for characterizing the original signals. Then, we perform feature selection to extract a reduced subset of features according to their mutual dependencies. The values assumed over time by the selected features are represented as a multivariate time series (MTS); that might contain missing values. Finally, in order to predict the PAF onset, we train a classifier based on special kernel similarity measures designed for MTS with missing data [16], [17]. The classifier is trained either in the MTS input space, using directly the relationships among data defined by the kernel similarity, or in an embedding space where each MTS is represented by its similarity values with respect to the training data.

To evaluate the effectiveness of the proposed framework, we take into account the data from the PhysioNet PAF prediction challenge [18]–[20]. Results show performance higher or on par with respect to the state-of-the-art results to correctly identify ECGs associated with PAF events. We also have the option to predict the PAF onset up to 15 minutes before the event, thus providing the ability to alert the patient of the possible occurrence of a PAF episode.

Our work provides the following key contributions:

- As methodological contribution, we propose a novel method for predicting PAF events that is based on a classifier trained either on the input or in the embedding space induced by a kernel similarity measure. In such a space, samples are described only by their similarities with respect to the elements in the training set.
- Kernel similarities for MTS with missing data [16], [17] are a recent development in time series analysis and they have never been considered so far in the context of PAF prediction, which represents thus a new field of application.
- We analyze a new type of PAF prediction problem, by studying how much time ahead the PAF events can be accurately detected by considering only the earliest measurements in the ECGs.
- The proposed framework is characterized by a high

degree of parallelization. This provides the opportunity of significantly speeding up the computation if the procedure is implemented on highly scalable systems, such as embedded devices with parallel computing support.

The remaining part of the paper is organized as follows: Section II describes the dataset of the PAF prediction challenge that we have used to develop our method; Section III describes the details of our framework, and Section IV presents the results that we have obtained.

II. DATASET DESCRIPTION

We consider the ECG data from the Physionet PAF prediction challenge [18]–[20]. Half of the records present in the database are acquired from patients prone to PAF (either immediately before fibrillation or after some time a fibrillation occurred), the other half from healthy subjects. No ECG of patients with an ongoing PAFs is considered. The database provides a training set, consisting of 50 records of PAF patients and 50 records of healthy subjects and a test set containing 50 records. Excerpts are 30 minutes long, concurrently acquired from two sensors.

The PAF prediction challenge [18] aims at predicting the onset of PAF events before their occurrence. Therefore, the task we tackle here is not a simple detection problem. The dataset under analysis consists of 106 recordings organized as follows. For the test set of the first PAF event (PAF detection), there are 22 recordings classified as “control” and 28 as “PAF-prone”. For generating the test set of the PAF prediction challenge (the one considered in this paper), only the 28 recordings of PAF-prone patients are considered. These 28 PAF-prone recordings are further divided into two classes: the “close-to-onset” and “far-from-onset” class. The former, represents records preceding a PAF event; the latter, instead, represents the records that do not precede a PAF event. Therefore, for generating the test set of the PAF prediction challenge, we end up with 56 recordings equally distributed between the two different classes. The training set is constructed by considering again pairs of recordings, this time taken from the 25 PAF-prone subjects belonging to the records used for training – for a total of 50 records to be used for training.

III. METHODOLOGY

We propose a classification framework, which implements the procedure summarized in Fig. 1. The procedure consists of three main steps: i) feature extraction, ii) computation of the kernel similarity and iii) training of the classifier. The details of each phase are given in the remainder of this section.

A. Feature extraction

The objective of this step is to extract from each ECG a set of descriptors, which are useful for the classification. Each selected descriptor is represented by a variable in an MTS that, in turn, becomes the new representation of the original ECG. Therefore, the ECG dataset will be represented by N MTS $\mathbf{X}_1, \dots, \mathbf{X}_N$ with $\mathbf{X} \in \mathbb{R}^{V \times T}$, where V is the number of variables, T_m is the number of time steps of the shortest



Fig. 1: Schematic depiction of the proposed classification procedure. The original dataset of ECG signals is processed by a feature extraction block, to represent them as MTS. Specific kernel measures are then used to evaluate the similarities between the MTS, which are stored in a kernel matrix. Such kernel matrix is finally exploited to train a classifier for predicting the PAF onset for each original ECG signal.

MTS, and T_M the number of time steps of the longest one. For the PAF dataset under analysis, we have $N = 106$, $T_m = 992$, and $T_M = 3364$. The number of variables V is determined by the feature extraction procedure, which consists of two main steps: *feature generation* and *feature selection*.

1) *Feature generation*: The ECG signals are initially processed for identifying the interval between two successive heartbeats (the R-R interval shown in Fig. 2). The variation in the time interval between heartbeats provides important clinical descriptors, which are used to characterize heart failure and other circulatory diseases [21].

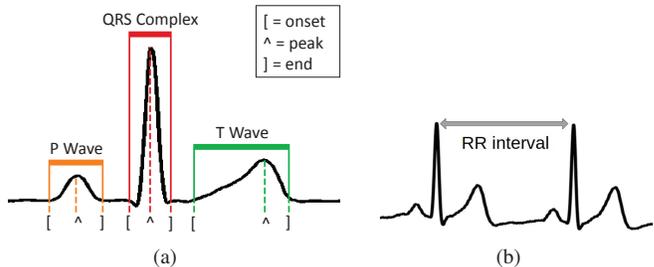


Fig. 2: In panel (a) the fiducial points of an ECG heartbeat [3]. In panel (b), an example of an R-R interval, corresponding to the time span between the R peaks of two consecutive QRS complex.

To compute the R-R intervals, we first pre-process the ECG signal using the morphological filtering technique described in [22]. Morphological filtering allows to retain useful information from acquired signals, while effectively eliminating noise originating from multiple sources such as low-frequency baseline wandering (caused by respiration and perspiration) and higher-frequency components (caused by muscular activity). A further step to enhance the quality of the acquired data is to combine signals from different inputs (leads) before the feature extraction step. In this study, we employ a Root Mean Square (RMS) combination of the two signals provided in the PAF-prediction database. Delineation is performed both on the two individual signals provided by the PAF prediction database and on their root-mean-square combination.

After pre-processing is completed, an algorithm based on the digital wavelet transform [3], [22] is used to retrieve R-R interval as well as the the fiducial points of each heartbeat:

the start, peak, and end of each characteristic wave (see Fig. 2).

There may be cases in which fiducial points cannot be extracted due to the low quality of the ECG signal. Low quality may be due to different reasons, e. g. imprecise/faulty measurements registered by the leads. In those cases, there will be a missing value in the descriptors, in correspondence of the time steps relative to portion of ECG measured with low quality.

While other techniques with better performance are available for extracting the fiducial points, the ones used in this work are suitable for being used in resources-constrained devices, such as the wearable ones. Thus, they satisfy the conditions in which our methodology is designed to operate.

The output of this procedure produces, for each original ECG signal, 95 time series of features related to morphological wave characteristics (for details on such features, see [3]).

2) *Feature selection*: After the feature generation procedure, each original ECG is represented by an MTS with $V = 95$ variates. Initially, we considered all the available features. However, preliminary results showed that training a classifier using all variates yielded worse results. Indeed, most of the original features are highly correlated. Furthermore, by including all features, the computational time to evaluate the kernel similarities and the training of the classifier increase considerably. Therefore, we considered different subsets of features in our experiments by performing feature selection on the original 95 variates. For guiding the selection of the features, we evaluated their linear Pearson correlation. In Fig. 3, we show the sample correlation matrix \mathbf{C} between the 95 variables, which is computed as follows:

- 1) $c_n^{i,j} = \text{corr}(\mathbf{x}_n^i, \mathbf{x}_n^j)$;
- 2) $\bar{c}^{i,j} = \frac{1}{N} \sum_n c_n^{i,j}$;
- 3) $\mathbf{C} = \begin{cases} \bar{c}^{i,j} = 0 & \text{if } |\bar{c}^{i,j}| < \theta_c, \\ \bar{c}^{i,j} = 1 & \text{otherwise.} \end{cases}$

where \mathbf{x}_n^i is the time series relative to the i -th variable in the n -th MTS \mathbf{X}_n and θ_c is a hyperparameter that thresholds the maximum degree of correlation between the variables.

In particular, we avoided to select at the same time two attributes if they were highly correlated (depicted as white pixels in the correlation matrix in Fig. 3). As a suitable threshold for the average correlation matrix \mathbf{C} , we used $\theta_c = 0.4$. This choice allows to get rid of highly correlated variables and, at the same time, significantly reduce the memory requirements of the procedure. After this feature selection procedure, we reduced the number of variables to $V = 22$. Those, describe features that are mostly related to R-R intervals and are summarized in Tab. I.

B. Kernel similarities for multivariate time series

Dynamic Time Warping [23] is the most widely used approach to assess the similarity between time series, but cannot treat MTS in its original formulation. More complicated variants have been proposed to deal with multiple variables [24], but they are characterized by difficult hyperparameter tuning

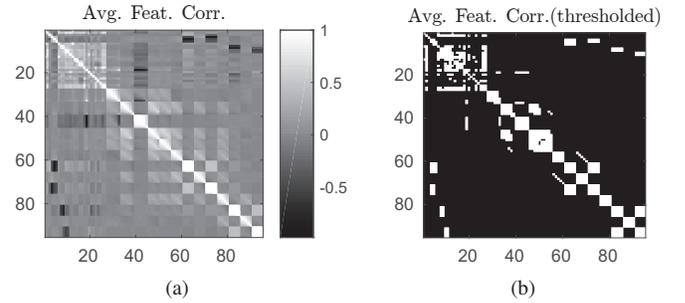


Fig. 3: In (a), the average correlation matrix of the ECG features for the MTS in the whole dataset. In (b), the same matrix where the values have been binarized using $\theta_c = 0.4$.

TABLE I: Variables identified by the feature selection procedure; see Fig. 2 for reference.

Name	Description
R_Position	Position of the R peak
IdentifiedStructures	Number of structures in the beat (P, R, and Q waves)
R_Max	Peak value of the R wave
R_Area	Area under the R wave
R_Width	Width of the R wave
R_ini	Time distance from the beginning of the R peak until the maximum value
P_Position	Position of the P wave
P_Max	Peak value of the P wave
P_Area	Area under the P wave
P_Width	Width of the P wave
P_ini	Time distance from the beginning of the P peak until the maximum value
PR_Interval	Interval between the P and the R waves
T_Position	Position of the T peak
T_Max	Peak value of the T wave
T_Area	Area under the T wave
T_Width	Width of the T wave
T_ini	Time distance from the beginning of the T peak until the maximum value
RT_Interval	Interval between the R and T wave
RR_Interval	Interval between two R peaks
PR_Segment	Time distance between P and R waves
RR_Interval_50HB_Mean	Interval between two R peaks, mean over the last 50 samples
RR_Interval_5HB_Mean	Interval between two R peaks, mean over the last 5 samples

and high computational complexity. Furthermore, there is not a straightforward way to account for missing values in the MTS when computing the similarities.

In this work, we adopt two different kernel similarities for MTS, the *Learned Pattern Similarity* (LPS) [16] and the *Time series Cluster Kernel* (TCK) [17], whose details are provided in the following two subsections. Those methods benefit from high parallelization capabilities and learn a model that, once the training is over, can quickly process new unseen data. Both procedures for computing LPS and TCK require all the MTS to have equal length T . Therefore, we followed the approach proposed in [25] and, by means of linear interpolation, we

transformed all MTS so that they have the same number of time steps T . The value of T is determined by $T = \left\lceil \frac{T_M}{\frac{T_M}{25}} \right\rceil$, where T_M is the length of the longest MTS in the dataset and $\lceil \cdot \rceil$ is the ceiling operator. Before computing TCK we also standardized each MTS, i.e., from each variable v in the dataset we subtract its mean value \bar{v} and divide by its standard deviation σ_v ; both \bar{v} and σ_v are computed on the training set. On the other hand, since decision trees are scale invariant, we did not apply this transformation for LPS (in accordance with [16]).

The training procedure returns a kernel matrix $\mathbf{K}_{tr} \in \mathbb{R}^{N_{tr} \times N_{tr}}$, whose components are the similarities among the elements of the training set. Once LPS and TCK are fitted, we process the test set and generate an additional matrix $\mathbf{K}_{te} \in \mathbb{R}^{N_{te} \times N_{tr}}$, whose components are the similarities of the elements in the training set with respect to those in the test set. N_{te} and N_{tr} represent the size of test and training set respectively, with $N = N_{te} + N_{tr}$.

TCK and LPS account for the missing patterns in the data to compute their similarities, rather than relying on imputation methods that may introduce strong biases. Imputation techniques replace the missing values with predefined or computed values, such as 0s, the empirical mean, or the last seen values. The choice is often arbitrary and, in presence of large amounts of missing values, the resulting data may end up being significantly different as well as the result of the analysis.

1) *Learned Pattern Similarity*: The learned pattern similarity (LPS) algorithm [16] is based on the identification of segments-occurrence within time series. It generalizes naturally to MTS by means of regression trees where a bag-of-words type compressed representation is created, which in turn is used to compute the similarity score. LPS is considered the state-of-the-art for MTS, inherits the decision tree approach to handle missing data, and is based on an ensemble strategy. Specifically, one extracts from MTS all possible segments of length L ($L < T$) starting from each time index $t = 1, 2, \dots, T - L + 1$ and fit them to nT randomly initialized regression trees. Since LPS naturally deals with missing data, in general it is not necessary to replace the missing entries.

LPS depends on two hyperparameters, which are the maximum number of regression trees considered, nT , and the length of the segments, denoted as L .

2) *Time Series Cluster Kernel*: The *Time series Cluster Kernel* (TCK) [17] computes an unsupervised kernel similarity for MTS with missing data. The method is based on an ensemble learning approach wherein the robustness to hyperparameters is ensured by joining the clustering results of many Gaussian mixture models (GMM) to form the final kernel. Hence, no critical hyperparameters have to be tuned by the user.

In order to deal with missing data, the GMMs are extended using informative prior distributions [26]. To generate partitions with different resolutions that capture both local and global structures in the input data, the TCK similarity matrix is

built by fitting GMMs to the set of time series for a number of mixtures, each one with a different number of components. To provide diversity in the ensemble, each partition is evaluated on a random subset of attributes and segments, using random initializations and randomly chosen hyperparameters. This also helps to provide robustness for what concerns hyperparameters selection. TCK is then built by summing (for each partition) the inner products between pairs of posterior distributions corresponding to different time series.

The hyperparameters that needs to be specified in TCK are the number of different random initializations Q and the maximum number of Gaussian mixtures C . It is sufficient to set those hyperparameters to reasonable high values to obtain good performance in many tasks [27], [28].

C. Classification

Once the kernel matrices are generated, we train a classifier based on the similarities yielded by \mathbf{K}_{tr} and \mathbf{K}_{te} . While in principle we can use any classifier operating on real-valued vectors, we consider two different ones, which are the k -Nearest Neighbors classifier (k NN) and the Support Vector Machine classifier (SVM). To perform classification we follow two different approaches: *classification in the input space* and *classification in the similarity-induced embedding space*.

1) *Classification in input space*: In the first approach, classification is performed directly in the MTS input space, considering as similarity between the samples the values contained in the kernel matrix \mathbf{K}_{te} . In k NN, for each element \mathbf{X}_i in the test set we select the relative row $\mathbf{k}_i \in \mathbf{K}_{te}$ and then we select the indices relative to the k highest values in \mathbf{k}_i , which identify the k MTS of the training set that are most similar to \mathbf{X}_i . The estimated \mathbf{X}_i label \hat{y}_i is the most frequent one among those k MTS.

On the other hand, in SVM is possible to train the classifier using \mathbf{K}_{tr} as precomputed kernel, which defines the kernel space where the optimal separating hyperplane is computed. In this case, the label \hat{y}_i is assigned according to the region (delimited by such hyperplane) that contains \mathbf{X}_i .

2) *Classification in embedding similarity space*: In this second approach, rather than assuming as input for the classification the original MTS, whose similarity is described by the kernel matrix, we train the classifier on the rows of \mathbf{K}_{tr} and \mathbf{K}_{te} . In particular, each row $\mathbf{k}_i \in \mathbb{R}^{N_{tr}}$ is considered as a vectorial embedding representation of the original MTS \mathbf{X}_i in the similarity space.

Therefore, in this approach to train the classifier it is necessary to compute an additional vector similarity among the rows of both \mathbf{K}_{tr} and \mathbf{K}_{te} . In k NN, we consider Euclidean, cosine, Cityblock and Pearson correlation similarities to identify, for each vector in \mathbf{K}_{te} , the k most similar vectors in \mathbf{K}_{tr} . Analogously, in SVM we compute a kernel similarity among the rows of \mathbf{K}_{tr} and \mathbf{K}_{te} using radial basis functions (rbf), whose analytic expression reads

$$d_{ij} = \exp \left\{ \gamma \| \mathbf{k}_i - \mathbf{k}_j \|^2 \right\}, \quad (1)$$

where γ defines the bandwidth of the kernel. By applying Eq. 1 to all the rows $i \in \mathbf{K}_{tr}$ and $j \in \mathbf{K}_{te}$, the embedding vectors are mapped into a new kernel space, where the optimal separating hyperplane is computed.

Compared to directly performing the classification in the input space, this approach requires an additional computation to evaluate the vector similarities among the representations (rows/columns of the kernel matrix) of the original MTS. The increment in the training time scales with the dimensionality of the vectors, which are equal to N_{tr} , i.e. the size of the training set. The procedure is more computational demanding in the case a SVM classifier is adopted, since the rbf kernels must be evaluated for each pair of vectors.

On the other hand, the proposed approach based on similarity-induced embedding provides a stronger form of regularization, which usually improves the generalization capability of the classifier. Additionally, the similarity embedding can model better the relationships among classes, thanks to its robustness to instance level noise [29]. Finally, the similarity embedding can highly synthesize the information when the MTS are characterized by a large number of variables and time steps, i.e. when $T \times V \gg N_{tr}$. However, in those cases where the classifier is required to capture a higher degree of variance among the data, the embedding approach might not be beneficial.

IV. RESULTS

In this section, we describe the experiments that we have performed and we report the results that we have obtained. The whole classification framework, including TCK and LPS kernels, are implemented in MATLAB¹. In each experiment, for LPS we set the number of regression trees to $nT = 200$ and the length of the segments to $L = 10$. Instead, for TCK we set the maximum number of Gaussian mixtures to $C = 40$ and the number of random initializations to $Q = 30$.

The two kernel matrices yielded by TCK and LPS are depicted in Fig. 4. The higher the values at the position (i, j) in such matrices, the higher the similarity between the corresponding MTS \mathbf{X}_i and \mathbf{X}_j . In the depiction, brighter colors correspond to higher similarities. It is possible to observe a block structure in the matrices, which denotes that the intra-class similarities are higher than the inter-class similarities. The blocks representing the MTS relative to *far-from-onset* and *close-to-onset* patients are visible more clearly in the kernel matrix yielded by TCK, which suggests its superior capability to discriminate between the two classes.

In the following, we present two different experiments. First, in Sec. IV-A we report the classification results on the original PAF dataset relative to the second challenge (see Sec. II for details), obtained by means of the proposed framework. Then, in Sec. IV-B we analyze how such results change as we vary the number of time steps taken into account for the signals; accordingly, this results in evaluating how the performance

¹TCK: <https://github.com/kmi010/Time-series-cluster-kernel-TCK->, LPS: <http://www.mustafabaydogan.com/learned-pattern-similarity-lps.html>

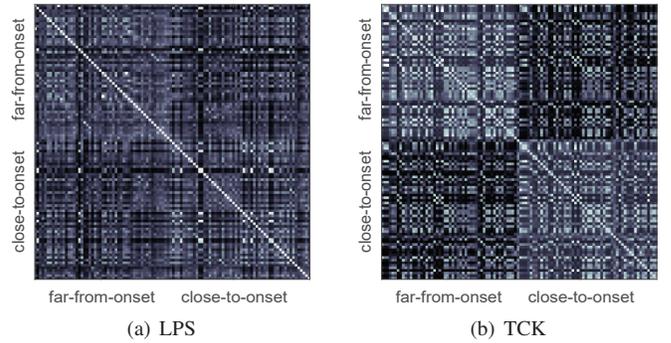


Fig. 4: Kernel matrix obtained with the LPS and TCK algorithms.

deteriorates when we omit data closer to the actual PAF onset. This last experiment is especially interesting as it shows how our methodology can be used to predict a PAF episode before its occurrence.

A. Classification results

Here, we evaluate the results obtained with the proposed approach to classify the MTS into the *far-from-onset* and *close-to-onset* classes, which is the objective of the second PAF challenge (PAF2) [18]. For comparison, we also report the three best official results, in terms of accuracy, obtained on the PAF2 challenge [30]. As discussed in Sec. III-C, it is possible to use the TCK and LPS kernel matrices directly as the pre-computed kernel in a kernel-based classifier, such as SVM, or to exploit the pairwise similarity scores in a k NN classifier operating in input space. As a second approach, we consider each row of the kernel matrix as an N_{tr} -dimensional embedding vector. Compared to their original representations, the samples are now highly compressed since we move from $\mathbb{R}^{V \times T}$ (size of the original MTS) to $\mathbb{R}^{N_{tr}}$ (length of the embedding vector), with $V \times T = 319580$ and $N_{tr} = 50$ in our dataset. In the embedding-based approach, several similarity measures have been used for evaluating the relationships of the embedding vectors \mathbf{k}_i .

To evaluate classification performance, we consider accuracy (ACC), recall on the “close-to-onset” class (REC), and finally the F1 score. To select the optimal hyperparameters of the classifiers, we perform k -fold cross-validation. For the k NN classifier, we consider a range of values for the number of neighbors k in $[1, 49]$ and, in case of classification in the embedding space, we consider as possible vector similarities Euclidean distance, cosine similarity, CityBlock distance and Pearson correlation. In presence of ties (when k assumes even values), we resolve them by deterministically assigning label of class “far-from-to-onset” to the samples. For the SVM classifier, we search the margin of the hyperplane in $C \in [2^{-20}, 2^{10}]$ and the kernel bandwidth in $\gamma \in [2^{-5}, 2^5]$. The classification results obtained from the kernel similarities yielded by LPS and TCK are shown in Tab. II, along with the optimal hyperparameters found. We refer to k NN-i, SVM-i

and k NN-e, SVM-e as the k NN and the SVM applied in the input (-i) and embedding (-e) space, respectively.

TABLE II: Classification results in the input space (k NN-i and SVM-i) and in the embedding space (k NN-e and SVM-e), using the two time series similarities, LPS and TCK. We report classification accuracy (ACC), recall (REC) for the close-to-onset class and the F1 measure (F1) relative to the configuration achieving the best ACC value in the validation procedure. We also report the optimal hyperparameters for such a configuration. Best results for each method are shown in bold. The last entries in the table are the best official results reported so far for the challenge [30]. Note that only ACC is reported as performance measure.

Method	Similarity	ACC	REC	F1	k	Diss	C	γ
k NN-i	LPS	0.714	0.928	0.764	41	-	-	-
	TCK	0.696	0.75	0.712	33	-	-	-
k NN-e	LPS	0.66	0.678	0.666	49	cosine	-	-
	TCK	0.678	0.683	0.675	29	cityblock	-	-
SVM-i	LPS	0.625	0.678	0.644	-	-	7.029	-
	TCK	0.696	0.571	0.653	-	-	3.031	-
SVM-e	LPS	0.625	0.39	0.517	-	-	20.072	1.045
	TCK	0.632	0.643	0.631	-	-	1.749	0.075
[31]	-	0.71	-	-	-	-	-	-
[32]	-	0.68	-	-	-	-	-	-
[33]	-	0.68	-	-	-	-	-	-

We observe that k NN and SVM when configured with TCK yield most of the time better results. However, when k NN is configured with LPS and it operates in the input space, it provides the best results, which are comparable with the state-of-the-art. Indeed, the highest official accuracy obtained for this problem is 0.71 and k NN-i+LPS obtains 0.714.

Interestingly, we notice that the validation procedure for the k NN always identifies as optimal high values of the number of neighbors (k). A high value of k improves the robustness to noise of the classifier. This is particularly important in real-world datasets, such as the one considered in this work, especially when the boundaries between different classes are not very distinct [34]. With respect to the remaining hyperparameters, we notice that the optimal configuration found with cross-validation varies from case to case.

B. Results on windows of increasing length

In this second experiment, we study how the classification outcomes change as we take into account an increasing number of time steps for the input MTS. In particular, we expand an initially small time window located far from the end of recording by iteratively including measurements that are closer to end and, hence, closer to the PAF onset. This allows us to evaluate how much time ahead the PAF events can be predicted accurately in our framework.

By applying a window of increasing size to the original data we obtain 10 different classification problems, c_1, c_2, \dots, c_{10} . Fig. 5 graphically illustrates the procedure. The first problem c_1 is created by considering the smallest window of heartbeats,

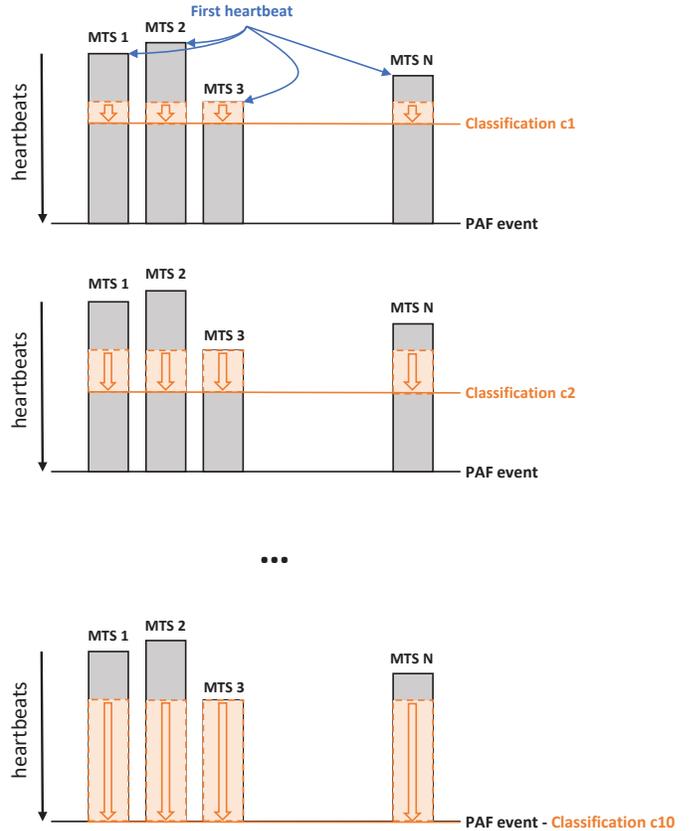


Fig. 5: Classification on windows of increasing size. The first classification problem, c_1 , takes into account a small time-window of heartbeats, which is far from the actual PAF event. The last classification problem, c_{10} , takes into account the maximum number of time steps, by considering a window with a size equal to the shortest MTS in terms of heartbeats (MTS_3 in this case).

positioned as far as possible from the PAF event (i.e., at the beginning of the recording). Note that the first heartbeat considered is the first of the shortest time series (MTS_3 in the example shown in Fig. 5). Therefore, the window of maximum width considered in the last classification problem c_{10} has length equal to the number of heartbeats in the shortest MTS of the PAF dataset (960 in this case). The details of each one of the 10 classification problems are reported in Tab. III. Note that, assuming that an adult has approximately 60-100 heartbeats per minute, the last column constitutes an approximation of the minutes between the end of the window and the occurrence of the PAF event.

The classification results for each problem by the different classification methods are reported in Fig. 6, in terms of accuracy (ACC) and recall (REC). Each classifier uses the same hyperparameter configurations that are described in Tab. II. We notice that the best performance for this task in terms of recall (REC) are obtained by the k NN classifier when is configured with LPS and operating in the input space. On the other hand, when the k NN is operating in input space and

TABLE III: Details of the classification problems generated by considering windows of increasing length. HB column denotes the number of heartbeats.

Problem	Win. size	HBs before PAF	Mins before PAF
c1	96	864	≈ 9 - 15
c2	192	768	≈ 8 - 13
c3	288	672	≈ 7 - 11
c4	384	576	≈ 6 - 10
c5	480	480	≈ 5 - 8
c6	576	384	≈ 4 - 7
c7	672	288	≈ 3 - 5
c8	768	192	≈ 2 - 3
c9	864	96	≈ 1 - 2
c10	960	0	0

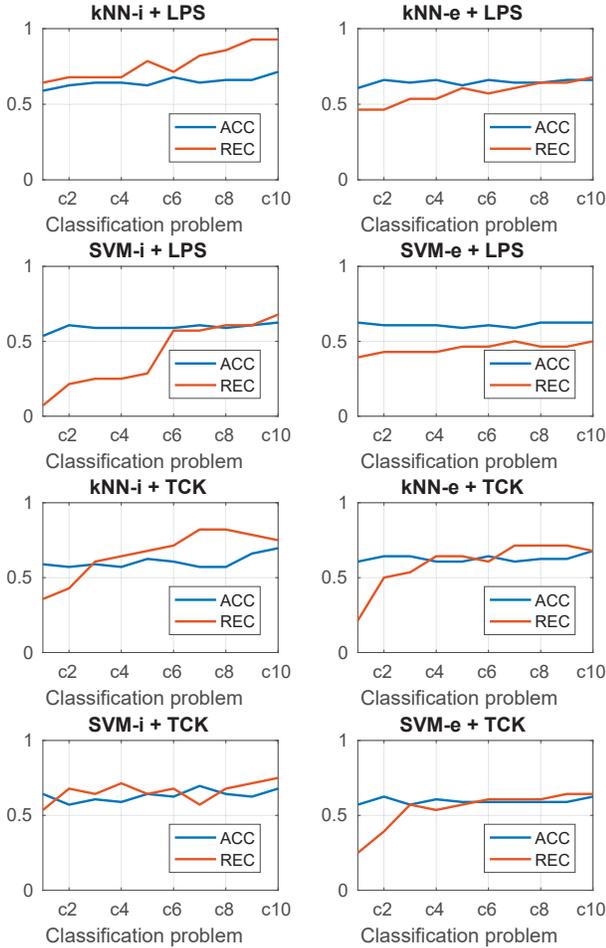


Fig. 6: Classification accuracy (ACC) and recall (REC) for the close-to-onset class obtained by the different classifiers when considering 10 different windows of increasing size as reported in Table III.

it is configured with TCK, we obtain the best performance in terms of accuracy (ACC). We also note that kNN-e configured with LPS achieves stable results for the different lengths of the time windows, since the performance of the classification does not deteriorate significantly when only a few data items before the event are considered. This suggests that such a method can

be the most suitable for an early detection of the PAF onset. We note that, in this latter case, high accuracy and recall are obtained in the problem c5 where the considered time window terminates roughly 8-10 minutes before the actual PAF onset.

Interestingly, in each experiment we observe that the recall is in general much lower when short time windows are considered. On the other hand, the accuracy of the classification is very stable in each classification problem. This demonstrates that the proposed method is capable of detecting the onset of PAF with a classification accuracy comparable with state of the art methods, with the additional advantage of detecting the PAF event several minutes before its occurrence.

V. CONCLUSIONS

In this paper, we proposed a methodology to classify ECG with the aim of predicting the onset Paroxysmal Atrial Fibrillation. We represented the ECGs as multi-variate time series of measurements of different descriptors over time. We have considered kernel functions for multi-variate time series, which allowed us to compute the pair-wise similarity between the ECGs. Such similarities are then exploited for training a classifier, defined either in the original input space or in a similarity-induced embedding space.

Our classification results are higher or on par with other state-of-the-art approaches. In addition, we show that with our proposed framework, it is possible to predict with a good precision the onset of PAF several minutes before the end of the available recordings. This result is of significant practical importance, as it makes possible to conceive an early-warning system for the onset of PAF, which can alert a patient before the occurrence of the event.

The proposed methodology is particularly efficient in terms of computational resources. In particular, the bottleneck, which is the computation of the kernel similarity matrix, can be implemented very efficiently by the TCK, which is an embarrassingly parallelizable algorithm and its computational complexity scales down linearly with the number of available computing units. Thanks to this property, the proposed framework is also suitable for implementations on highly parallel hardware devices and on multicore microprocessors, such as the ones that are being adopted in embedded systems. Future includes porting onto embedded systems of the proposed software pipeline.

ACKNOWLEDGEMENTS

This work was supported and funded by the Hasler Foundation under the Project “HSD: A Personal Device for Automatic Evaluation of Health Status during Physical Training” (Grant No. 15048). The paper reflects only the view of the authors.

The authors would like to thank Elisabetta De Giovanni, Dr. Amir Aminifar, and Prof. Alonso David Atienza from École Polytechnique Fédérale de Lausanne for helping in the feature extraction process as well as for their precious feedback.

REFERENCES

- [1] R. Alcaraz and J. J. Rieta, "A review on sample entropy applications for the non-invasive analysis of atrial fibrillation electrocardiograms," *Biomedical Signal Processing and Control*, vol. 5, no. 1, pp. 1–14, 2010.
- [2] B. Hickey, C. Heneghan, and P. De Chazal, "Non-episode-dependent assessment of paroxysmal atrial fibrillation through measurement of RR interval dynamics and atrial premature contractions," *Annals of Biomedical Engineering*, vol. 32, no. 5, pp. 677–687, 2004.
- [3] J. Milosevic, A. Dittrich, A. Ferrante, M. Malek, C. R. Quiros, R. Braojos, G. Ansaloni, and D. Atienza, "Risk assessment of atrial fibrillation: a failure prediction approach," in *Computing in Cardiology Conference*. Cambridge, MA, USA: IEEE, Sep. 2014, pp. 801–804.
- [4] L. S. Wann, A. B. Curtis, C. T. January, K. A. Ellenbogen, J. E. Lowe, N. M. Estes, R. L. Page, M. D. Ezekowitz, D. J. Slotwiner, W. M. Jackman, W. G. Stevenson, and C. M. Tracy, "2011 accf/aha/hrs focused update on the management of patients with atrial fibrillation (updating the 2006 guideline)," *Circulation*, vol. 123, no. 1, pp. 104–123, 2011. [Online]. Available: <http://circ.ahajournals.org/content/123/1/104>
- [5] L. Y. Di Marco, D. Raine, J. P. Bourke, and P. Langley, "Characteristics of atrial fibrillation cycle length predict restoration of sinus rhythm by catheter ablation," *Heart Rhythm*, vol. 10, no. 9, pp. 1303–1310, 2013.
- [6] M.-T. Lo, Y.-C. Chang, C. Lin, H.-W. Young, Y.-H. Lin, Y.-L. Ho, C.-K. Peng, and K. Hu, "Outlier-resilient complexity analysis of heartbeat dynamics," *Scientific Reports*, vol. 5, 2015.
- [7] R. Cervigón, J. Moreno, R. B. Reilly, J. Millet, J. Pérez-Villacastín, and F. Castells, "Entropy measurements in paroxysmal and persistent atrial fibrillation," *Physiological Measurement*, vol. 31, no. 7, p. 1011, 2010.
- [8] A. Martínez, D. Abásolo, R. Alcaraz, and J. J. Rieta, "Alteration of the P-wave non-linear dynamics near the onset of paroxysmal atrial fibrillation," *Medical Engineering & Physics*, vol. 37, no. 7, pp. 692–697, 2015.
- [9] I. Dotsinsky, "Atrial wave detection algorithm for discovery of some rhythm abnormalities," *Physiological Measurement*, vol. 28, no. 5, p. 595, 2007.
- [10] A. Martínez, R. Alcaraz, and J. J. Rieta, "Morphological variability of the P-wave for premature envision of paroxysmal atrial fibrillation events," *Physiological Measurement*, vol. 35, no. 1, p. 1, 2013.
- [11] T. Thong, J. McNames, M. Aboy, and B. Goldstein, "Prediction of paroxysmal atrial fibrillation by analysis of atrial premature complexes," *IEEE Transactions on Biomedical Engineering*, vol. 51, no. 4, pp. 561–569, Apr. 2004.
- [12] J. Schlenker, V. Socha, L. Riedlbauchová, T. Nedělka, A. Schlenker, V. Potočková, Š. Malá, and P. Kutílek, "Recurrence plot of heart rate variability signal in patients with vasovagal syncope," *Biomedical Signal Processing and Control*, vol. 25, pp. 1–11, 2016.
- [13] X. Du, N. Rao, M. Qian, D. Liu, J. Li, W. Feng, L. Yin, and X. Chen, "A novel method for real-time atrial fibrillation detection in electrocardiograms using multiple parameters," *Annals of Noninvasive Electrocardiology*, vol. 19, no. 3, pp. 217–225, 2014.
- [14] X. Li, D. Yang, X. Liu, and X.-M. Wu, "Bridging time series dynamics and complex network theory with application to electrocardiogram analysis," *IEEE Circuits and Systems Magazine*, vol. 12, no. 4, pp. 33–46, Nov. 2012.
- [15] X. Li and Z. Dong, "Detection and prediction of the onset of human ventricular fibrillation: An approach based on complex network theory," *Physical Review E*, vol. 84, no. 6, p. 062901, 2011.
- [16] M. G. Baydogan and G. Runger, "Time series representation and similarity based on local autopatterns," *Data Mining and Knowledge Discovery*, vol. 30, no. 2, pp. 476–509, 2016.
- [17] K. Ø. Mikalsen, F. M. Bianchi, C. Soguero-Ruiz, and R. Jenssen, "Time series cluster kernel for learning similarities between multivariate time series with missing data," *Pattern Recognition*, vol. 76, pp. 569 – 581, 2018.
- [18] "The PAF Prediction Challenge Database," last accessed: 5-Jul-2016. [Online]. Available: <https://physionet.org/physiobank/database/afpdb/>
- [19] G. B. Moody, A. L. Goldberger, S. McClennen, and S. P. Swiryn, "Predicting the onset of paroxysmal atrial fibrillation: The computers in cardiology challenge 2001," in *Computers in Cardiology*. Rotterdam, Netherlands: IEEE, Sep. 2001, pp. 113–116.
- [20] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. Ch. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "Physiobank, physiotoolkit, and physionet components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. 215–220, 2000.
- [21] M. Malik, P. Färbom, V. Batchvarov, K. Hnatkova, and A. Camm, "Relation between qt and rr intervals is highly individual among healthy subjects: implications for heart rate correction of the qt interval," *Heart*, vol. 87, no. 3, pp. 220–228, 2002.
- [22] R. Braojos Lopez, G. Ansaloni, D. Atienza Alonso, R. Vallejos, and F. Javier, "Embedded Real-Time ECG Delineation Methods: a Comparative Evaluation," in *IEEE 12th International Conference on BioInformatics and BioEngineering (BIBE 2012)*, vol. 1, 2012, pp. 99–104.
- [23] D. J. Berndt and J. Clifford, "Using dynamic time warping to find patterns in time series," in *KDD workshop*, vol. 10, no. 16. Seattle, WA, 1994, pp. 359–370.
- [24] M. Shokoohi-Yekta, B. Hu, H. Jin, J. Wang, and E. Keogh, "Generalizing dtw to the multi-dimensional case requires an adaptive approach," *Data Mining and Knowledge Discovery*, vol. 31, no. 1, pp. 1–31, 2017.
- [25] L. Wang, Z. Wang, and S. Liu, "An effective multivariate time series classification approach using echo state network and adaptive differential evolution algorithm," *Expert Systems with Applications*, vol. 43, pp. 237 – 249, 2016.
- [26] B. M. Marlin, D. C. Kale, R. G. Khemani, and R. C. Wetzel, "Unsupervised pattern discovery in electronic health care data using probabilistic clustering models," in *Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium*, Miami, FL, USA, Jan. 2012, pp. 389–398.
- [27] K. Ø. Mikalsen, F. M. Bianchi, C. Soguero-Ruiz, S. O. Skrivseth, R.-O. Lindsetmo, A. Revhaug, and R. Jenssen, "Learning similarities between irregularly sampled short multivariate time series from EHRs," in *Proc. 3rd International Workshop on Pattern Recognition for Healthcare Analytics at ICPR 2016*, 2016.
- [28] K. Øyvind Mikalsen, C. Soguero-Ruiz, F. M. Bianchi, A. Revhaug, and R. Jenssen, "An Unsupervised Multivariate Time Series Kernel Approach for Identifying Patients with Surgical Site Infection from Blood Samples," *ArXiv e-prints*, 2018.
- [29] K. Q. Weinberger and O. Chapelle, "Large margin taxonomy embedding for document categorization," in *Advances in Neural Information Processing Systems*, 2009, pp. 1737–1744.
- [30] "Paf challenge scores." [Online]. Available: <https://physionet.org/challenge/2001/top-scores.shtml>
- [31] G. Schreier, P. Kastner, and W. Marko, "An automatic eeg processing algorithm to identify patients prone to paroxysmal atrial fibrillation," in *Computers in Cardiology 2001*. IEEE, 2001, pp. 133–135.
- [32] P. de Chazal and C. Heneghan, "Automated prediction of the onset of paroxysmal atrial fibrillation from surface electrocardiogram recordings," in *Computer in Cardiology*, 2001.
- [33] C. Maier, M. Bauch, and H. Dickhaus, "Screening and prediction of paroxysmal atrial fibrillation by analysis of heart rate variability parameters," in *Computers in Cardiology 2001*. IEEE, 2001, pp. 129–132.
- [34] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. New York, NY: Wiley & Sons, 2001.