

Tensor-based ECG Signal Processing Applied to Atrial Fibrillation Detection

Simon Geirnaert[†], Griet Goovaerts^{†*}, Sibasankar Padhy^{†*}, Martijn Boussé[†], Lieven De Lathauwer^{†‡},
Sabine Van Huffel^{†*}

[†]Department of Electrical Engineering (ESAT), KU Leuven, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium;
*imec, Leuven, Belgium;

[‡]Group Science, Engineering and Technology, KU Leuven Kulak, E. Sabbelaan 53, B-8500 Kortrijk, Belgium.

Abstract—Atrial fibrillation (AF) is the most common cardiac arrhythmia, increasing the risk of a stroke substantially. Hence, early and accurate detection of AF is paramount. We present a matrix- and tensor-based method for AF detection in single- and multi-lead electrocardiogram (ECG) signals. First, the recordings are compressed into one heartbeat via the singular value decomposition (SVD). These representative heartbeats, single-lead, are collected in a matrix with modes time and recordings. In the multi-lead case, we obtain a tensor with modes lead, time and recording. By modeling the matrix (tensor) with a (multilinear) SVD, each recording, as well as new recordings, can be expressed by a coefficient vector. The comparison of a new coefficient vector with those of the model set results in morphological features, which are combined with heart rate variability information in a Support Vector Machine classifier to detect AF. The SVD-based method is tested on the 2017 PhysioNet/CinC Challenge dataset, resulting in an F_1 -score of 0.77. The multilinear SVD-based method is applied on the MIT-BIH AFIB and AFTDB dataset, resulting in a perfect separation. An advantage of our methods is the interpretability of the features, which is a key element in the application of automatic methods in clinical practice.

I. INTRODUCTION

Atrial fibrillation (AF) affects approximately one percent of the general population, making it the most widespread cardiac arrhythmia [1]. During AF, electrical chaos originates in the atria, resulting in uncoordinated and unsynchronized muscle contractions, which cause the atria to flutter or fibrillate. The atrioventricular node filters these additional impulses so the ventricles still contract normally, although at a highly irregular rate. Since the cardiac pump function is mainly driven by ventricular contraction, it remains unaltered, but the accumulation of blood in the atria can cause the formation of blood clots which possibly leads to pulmonary embolism or strokes [1].

AF affects both heartbeat morphology and heart rate, allowing one to effectively diagnose the arrhythmia via the

This research is funded by an Aspirant Grant from the Research Foundation - Flanders (FWO), Fonds de la Recherche Scientifique – FNRS and Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS Project no 30468160 (SeLMA), Research Council KU Leuven: C1 project C16/15/059-nD, EU: The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013)/ERC Advanced Grant: BIOTENSORS (no. 339804). This paper reflects only the authors’ views and the Union is not liable for any use that may be made of the contained information.

electrocardiogram (ECG). More specifically, AF replaces the contraction of the atria with fibrillations, eliminating the P-wave. Early detection is essential to start treatment which is often done with medication [2]. The need for an early and accurate detection of AF motivates the development of automatic methods to assist cardiologists in the processing of large amounts of data which originate, for example, from mobile health technologies.

In this paper, we present a method that combines information from the heart rhythm and the morphology of heartbeats in order to detect AF in short-term single- and multi-lead ECG signals. While heart rate variability (HRV) features are well-known in literature, see for example [3], we focus on the design of morphology-based features in this paper. More specifically, we use a similar method as proposed for irregular heartbeat detection [4].

In the preprocessing stage, ECG recordings of different lengths can be represented in a matrix by compressing them each in one representative heartbeat via the singular value decomposition (SVD). In the modeling stage, representative heartbeats, single-lead, of different recordings are collected into a matrix, which is then modeled by an SVD. Each representative heartbeat is represented in the model by one compressed coefficient vector. When considering a *new* ECG signal, for which we have to determine whether AF activity is present or not, we can apply a similar strategy. By comparing the obtained coefficient vector of the new signal with the database of *known* coefficient vectors, one obtains a similarity measure with each rhythm present in this model set. In the last step, we combine this morphological information with traditional HRV features in a Support Vector Machine (SVM) classifier to determine one final label.

This method can be extended and generalized in an intuitive way to tensor-based modeling using the multilinear singular value decomposition (MLSVD) [5] for multi-lead ECG signals. In multi-lead ECG, the electrical activity of the heart is registered from different directions, resulting in extra spatial information. The tensor allows to fully exploit the higher-order structure in the multi-lead ECG data.

In the remainder of this section, we describe notations and definitions. We present an SVD-based method for single-lead ECG in Section II and generalize to multi-lead ECG using tensors in Section III. We discuss experiments in Section IV.

A. Notations and definitions

Scalars, vectors, matrices and tensors are denoted by lowercase (t), bold lowercase (\mathbf{t}), bold uppercase (\mathbf{T}) and calligraphic letters (\mathcal{T}), respectively. A mode- n vector of a tensor $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is the vector obtained by fixing all indices but the n^{th} . The mode- n unfolding is the matrix $\mathbf{T}_{(n)} \in \mathbb{R}^{I_n \times J}$, with $J = \prod_{k \neq n}^N I_k$, with the mode- n vectors as its columns, ordered as in [6]. The vectorization of a tensor is a vector $\text{vec}(\mathcal{T}) \in \mathbb{R}^J$, with $J = \prod_{k=1}^N I_k$, constructed by filling it first by the first mode, then the second mode, \dots , until the N^{th} mode.

The Kronecker product of two matrices is denoted by \otimes . The mode- n product of a matrix $\mathbf{U} \in \mathbb{R}^{J \times I_n}$ and a tensor $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is a tensor $\mathcal{T} \cdot_n \mathbf{U} \in \mathbb{R}^{I_1 \times \dots \times I_{n-1} \times J \times I_{n+1} \times \dots \times I_N}$ in which each mode- n vector is multiplied with \mathbf{U} : $(\mathcal{T} \cdot_n \mathbf{U})_{(n)} = \mathbf{U} \mathbf{T}_{(n)}$.

The multilinear rank of a tensor is the tuple of mode- n ranks, where the mode- n rank is equal to the rank of the mode- n unfolding of the tensor.

B. Multilinear singular value decomposition

The multilinear singular value decomposition (MLSVD) is the higher-order generalization of the SVD for tensors [5].

Definition. The multilinear singular value decomposition of a third-order tensor $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ is equal to:

$$\mathcal{T} = \mathcal{S} \cdot_1 \mathbf{U}^{(1)} \cdot_2 \mathbf{U}^{(2)} \cdot_3 \mathbf{U}^{(3)},$$

with orthogonal factor matrices $\mathbf{U}^{(n)} \in \mathbb{R}^{I_n \times I_n}$ and all-orthogonal and ordered core tensor $\mathcal{S} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$.

The MLSVD has been used successfully in various signal processing and machine learning applications such as compression and dimensionality reduction [6]–[8].

II. SVD-BASED DETECTION IN SINGLE-LEAD ECG

In this section, we explain that AF detection in single-lead ECG can be accomplished by first compressing different ECG recordings in a single representative heartbeat. Next, a model matrix of representative heartbeats is constructed and modeled by means of an SVD. New recordings can be characterized and classified by computing the new coefficient vector with respect to the model. In the next section, we generalize our method to multi-lead ECG by using tensors.

A. Compression in one representative heartbeat

Assuming preprocessed signals, we compress the single-lead ECG recording into one representative heartbeat to accommodate for variable-length recordings. This enables us to collect different recordings into one matrix. Furthermore, this compression results in a denoising effect, removing irrelevant variations across heartbeats.

First, we use the Pan-Tompkins algorithm [9] to detect the R-peaks, allowing us to segment the signal into fixed-length windows of 580 ms, starting 250 ms before and ending 330 ms after each R-peak. We use the aforementioned values such that both the P-wave and T-wave are included in each segment,

see [4]. Next, the different windows are optimally aligned by maximizing the cross-correlation between different windows. This enables full exploration of the structure in the heartbeats.

The aligned heartbeats are collected in a $K \times N$ matrix \mathbf{X} , with K the number of heartbeats and N the number of samples in a heartbeat. Next, we compress all heartbeats into one *representative* heartbeat $\mathbf{x}_r \in \mathbb{R}^N$ using the SVD. The Eckart-Young theorem states that the best rank-1 approximation of a matrix \mathbf{X} can be calculated by means of an SVD: $\mathbf{X} \approx \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T$. Vector $\mathbf{v}_1 \in \mathbb{R}^N$ serves as a basis vector of the heartbeats space. This vector is rescaled to norm one and represents the representative heartbeat for the signal. Approximating the aligned heartbeats with a rank-1 matrix is meaningful because the variation between beats within a short-term signal is typically quite small, which corresponds to a single dominant singular value.

B. Model construction

In the modeling stage, a model set of representative heartbeats from different recordings is used to construct a basis for the recording space. The labels of the signals in the model set, corresponding to the underlying rhythm in the signal, are considered to be known.


$\mathbf{D}_{\text{model}} \in \mathbb{R}^{N \times M}$ is the matrix containing all M representative heartbeats $\mathbf{x}_r \in \mathbb{R}^N$ from the model set of M ECG recordings. We compute a low-rank model via the truncated SVD:

$$\mathbf{D}_{\text{model}} \approx \hat{\mathbf{U}}_{\text{time}} \hat{\mathbf{S}} \hat{\mathbf{U}}_{\text{recording}}^T, \quad (1)$$

with $\hat{\mathbf{U}}_{\text{time}} \in \mathbb{R}^{N \times r}$ an orthonormal basis for a time subspace, $\hat{\mathbf{U}}_{\text{recording}} \in \mathbb{R}^{M \times r}$ an orthonormal basis for a recordings subspace and $\hat{\mathbf{S}} \in \mathbb{R}^{r \times r}$ the core diagonal matrix.

The rank of the truncated SVD in (1), denoted by r , is crucial to tune the model in order to obtain good generalization properties, balancing between under- and overfitting.

Each representative heartbeat $\mathbf{x}_r^{(i)}$, $1 \leq i \leq M$ (i.e., each column from $\mathbf{D}_{\text{model}}$) is expressed by (1) as:



with $\mathbf{c}_{\text{rec}}^{(i)T} \in \mathbb{R}^r$ the i^{th} row from $\hat{\mathbf{U}}_{\text{recording}}$. This is the coefficient vector that expresses the corresponding signal in the basis spanned by \mathbf{B} . These column vectors are unique up to sign and scaling and will therefore be normalized. All normalized coefficient vectors $\tilde{\mathbf{c}}_{\text{rec}}^{(i)} = \text{sign}(\mathbf{c}_{\text{rec},1}^{(i)}) \frac{\mathbf{c}_{\text{rec}}^{(i)}}{\|\mathbf{c}_{\text{rec}}^{(i)}\|_2}$ then form a database. Each vector of this database has a known label and is afterwards used for classification of new signals.

C. Classification of new recordings

A new ECG signal is first compressed, as explained in Subsection II-A. We then project the new representative heartbeat $\mathbf{x}_r^{(\text{new})}$ onto the space spanned by \mathbf{B} , resulting in a new coefficient vector $\mathbf{c}_{\text{rec}}^{(\text{new})}$. This boils down to the following computation:

$$\mathbf{c}_{\text{rec}}^{(\text{new})} = \hat{\mathbf{S}}^{-1} \hat{\mathbf{U}}_{\text{time}}^T \mathbf{x}_r^{(\text{new})}.$$

Comparison between the normalized new coefficient vector $\tilde{\mathbf{c}}_{\text{rec}}^{(\text{new})}$ and the coefficient vectors in the database is done by pairwise calculation of the inner product, resulting in a similarity score between the new signal and all signals in the database:

$$s_i = \tilde{\mathbf{c}}_{\text{rec}}^{(\text{new})T} \tilde{\mathbf{c}}_{\text{rec}}^{(i)}, \forall i : 1 \leq i \leq M.$$

We then compute a weighted average of the similarity scores per class c , resulting in a similarity score between the new signal and each rhythm present in the database:

$$f_c^{(\text{new})} = \sum_{i \in c} w_i s_i,$$

with c the set of recordings in a predefined class/rhythm.

These scores are collected into one vector $\mathbf{f}^{(\text{new})}$ of morphological features. A meaningful weight vector, typically only averaging over the W best scores, as well as W , can be empirically determined.

Finally, different classification algorithms can be used to obtain a classification from the feature vector $\mathbf{f}^{(\text{new})}$. The most straightforward method is to assign to the new signal the label of the class which has the highest score. This method is for example used during cross-validation to determine the optimal rank. In Subsection III-E, SVMs are used for classification to simplify the integration of the morphological feature vector with other HRV features. This is an important extension w.r.t. [4], where only the straightforward method is used. In [10], a similar extension to an extra machine learning classifier is adopted.

III. MLSVD-BASED DETECTION IN MULTI-LEAD ECG

Clinically, ECG is usually measured with multiple leads. This section describes a method to detect AF in multi-lead ECG signals using the same principles as the previous method.

A. Tensorization

We now use a model set of known multi-lead ECG signals as basis for the model. As in the previous method, the ECG signals are first preprocessed lead-by-lead to improve the signal quality. They are then compressed into one representative heartbeat per-lead, where it is ensured that the sign of the R-peak is positive. Since a multi-lead ECG signal naturally has a second-order structure, stacking the representative beats as frontal slices of a model tensor $\mathcal{D}_{\text{model}}$ automatically leads to a tensorization which preserves the structural information between all leads, see Figure 1. $\mathcal{D}_{\text{model}} \in \mathbb{R}^{n_k \times N \times M}$ has modes lead \times time \times recording, with n_k the number of leads, N the number of samples per representative heartbeat and M the number of recordings in the model set.




Fig. 1: Tensorization of the different multi-lead ECG signals of the model set, after per-lead computation of a representative heartbeat. The resulting third-order tensor has modes lead \times time \times recording.

B. Model construction

The truncated MLSVD provides a low multilinear rank approximation of our model tensor $\mathcal{D}_{\text{model}}$. Although it is known that the Eckart-Young theorem does not generalize to tensors [5], the approximation is *reasonably* good. We obtain:

$$\mathcal{D}_{\text{model}} \approx \hat{\mathbf{S}} \cdot_1 \hat{\mathbf{U}}_{\text{lead}} \cdot_2 \hat{\mathbf{U}}_{\text{time}} \cdot_3 \hat{\mathbf{U}}_{\text{recording}}, \quad (2)$$

where $\hat{\mathbf{U}}_{\text{lead}} \in \mathbb{R}^{n_k \times r_{\text{lead}}}$, $\hat{\mathbf{U}}_{\text{time}} \in \mathbb{R}^{N \times r_{\text{time}}}$ and $\hat{\mathbf{U}}_{\text{recording}} \in \mathbb{R}^{M \times r_{\text{recording}}}$ form orthonormal bases for the corresponding subspaces. $\hat{\mathbf{S}} \in \mathbb{R}^{r_{\text{lead}} \times r_{\text{time}} \times r_{\text{recording}}}$ explains the interaction between the different modes.

Each multi-lead ECG signal with n_k leads and thus n_k representative heartbeats $\mathbf{X}_r^{(i)} \in \mathbb{R}^{n_k \times N}$ can be written as

$$\text{vec} \left(\mathbf{X}_r^{(i)} \right) \approx \mathbf{B} \mathbf{c}_{\text{rec}}^{(i)},$$

with $\mathbf{B} = \mathbf{B}_{(3)}^T \in \mathbb{R}^{n_k N \times r_{\text{recording}}}$ the mode-3 unfolding of the tensor $\mathcal{B} = \hat{\mathbf{S}} \cdot_1 \hat{\mathbf{U}}_{\text{lead}} \cdot_2 \hat{\mathbf{U}}_{\text{time}}$:

$$\mathbf{B} = (\hat{\mathbf{U}}_{\text{time}} \otimes \hat{\mathbf{U}}_{\text{lead}}) \hat{\mathbf{S}}_{(3)}^T.$$

The columns of \mathbf{B} then form a basis, not orthogonal anymore, for all leads. The coefficient vector expresses a multi-lead ECG signal in all leads simultaneously.

In a similar fashion as for the matrix case, the normalized coefficient vectors $\tilde{\mathbf{c}}_{\text{rec}}^{(i)}$ constitute a labeled database which can be employed for classification.

C. Rank estimation

The optimal multilinear rank in (2) can be determined using sequential cross-validation. We try to overcome the heavy complexity of standard cross-validation, in which the number of multilinear rank values to be assessed scales as a third power with the size of the search interval. This sequential cross-validation improves on the initial guess of the multilinear rank by fixing in each iteration two rank values, while varying the other one:

$$\dots \rightarrow \left. \begin{matrix} r_{\text{lead}}^{(i)} \\ r_{\text{time}}^{(i)} \\ r_{\text{rec}}^{(i)} \end{matrix} \right\} \xrightarrow{\text{Vary } r_{\text{time}}^{(i)}} \left. \begin{matrix} r_{\text{lead}}^{(i)} \\ r_{\text{time}}^{(i+1)} \\ r_{\text{rec}}^{(i)} \end{matrix} \right\} \xrightarrow{\text{Vary } r_{\text{rec}}^{(i)}} \left. \begin{matrix} r_{\text{lead}}^{(i)} \\ r_{\text{time}}^{(i+1)} \\ r_{\text{rec}}^{(i+1)} \end{matrix} \right\} \xrightarrow{\text{Vary } r_{\text{lead}}^{(i)}} \left. \begin{matrix} r_{\text{lead}}^{(i+1)} \\ r_{\text{time}}^{(i+1)} \\ r_{\text{rec}}^{(i+1)} \end{matrix} \right\} \rightarrow \dots$$

TABLE I: Traditional HRV features from literature.

AVRR	Average RR interval
RMSSD	Root Mean Square of the differences between subsequent RR intervals [3]
pRR50	Fraction of subsequent RR intervals that differ more than 50ms [3]
Cov(Δ RR)	Coefficient of variation of the Δ RR intervals [12]
min(RR)	Minimal RR interval
HTI	HRV Triangular Index, characterizing the concentration of the RR interval histogram [13]
SD1/SD2 ratio	Characterizes the Poincaré plot of RR intervals [13]
AFEvidence	Characterizes the Poincaré plot of Δ RR intervals [14], [15]
ApEn	Approximate entropy, quantifying the complexity of the signal [16], [17]
Toeplitz distance	A new feature that quantifies the uniformity of RR intervals in a global way, see Algorithm 1

Note that the choice of the order in which the rank values are varied is important, as there will be a bias towards the rank value(s) optimized last.

D. Classification of new recordings

In a similar way as described in Subsection II-C new ECG signals can be classified by computing the coefficient vector in basis \mathbf{B} . Given the per-lead representative heartbeats $\mathbf{X}_r^{(\text{new})}$, the coefficient vector $\mathbf{c}_{\text{rec}}^{(\text{new})}$ can be found by solving:

$$\text{vec} \left(\mathbf{X}_r^{(\text{new})} \right) = \mathbf{B} \mathbf{c}_{\text{rec}}^{(\text{new})},$$

which can then be processed in a similar way as in Subsection II-C. Note that the underlying assumption is that the lead space does not change for new ECG recordings, because $\hat{\mathbf{U}}_{\text{lead}}$ is fixed. This corresponds to the reasonable assumption that the ECG is measured in a normalized fashion. The removal of this assumption results in a so-called Kronecker product equation [4], [11].

E. Combination of morphological and HRV features

In the final stage of the algorithm, the morphological features derived in Sections II and III are combined with traditional HRV features, which are shown in Table I. We combine the HRV and morphological features into one feature vector, which serves as an input for a classifier.

Algorithm 1 Computation Toeplitz distance

Input: Vector $\mathbf{q} \in \mathbb{N}^P$ of positions of R-peaks

Output: Toeplitz distance t_d

- 1: Construct symmetric distance matrix $\mathbf{D} \in \mathbb{R}^{P \times P}$, with $\mathbf{D}_{ij} = |q_j - q_i|, 1 \leq i, j \leq P$
 - 2: Compute $\mathbf{g} \in \mathbb{R}^{2P-1}$, with $\mathbf{g}_i = \text{average}([D_{P-i+1,1} \quad D_{P-i+2,2} \quad \dots \quad D_{P,i}])$ and $g_{2P-i} = g_i, 1 \leq i \leq P$
 - 3: Construct the Toeplitz matrix $\mathbf{D}_{\text{toep}} \in \mathbb{R}^{P \times P}$, with generating vector \mathbf{g}
 - 4: Compute the Toeplitz distance: $t_d = \frac{1}{P} \|\mathbf{D} - \mathbf{D}_{\text{toep}}\|_F$
 - 5: **return** t_d
-

TABLE II: The data profile of the model, training and test set of the single-lead PhysioNet/CinC Challenge 2017 dataset.

	Class			Total
	Normal	AF	Other	
Model set	3062	413	1471	4946
Training set	994	163	493	1650
Test set	994	162	492	1648
Total	5050	738	2456	8244

IV. RESULTS AND DISCUSSION

A. SVD-based detection in single-lead ECG

We illustrate our SVD-based detection in single-lead ECG on the 2017 PhysioNet/Computing in Cardiology challenge dataset [18], [19]. This dataset consists of 8244 publicly available single-lead ECG signals. We consider three classes: normal sinus rhythms (NSR), AF and other signals. This latter class contains signals with other abnormal rhythms or on which annotators disagreed. The signals have a length between 9 and 60 seconds. Table II shows how the signals are divided over the different classes.

We apply compression as explained in Subsection II-A, obtaining representative heartbeats of size $N = 176$. The model matrix $\mathbf{D}_{\text{model}} \in \mathbb{R}^{176 \times 4946}$ is modeled with a rank-22 approximation, obtained after 10-fold cross-validation in search interval [1, 30]. The search interval can be determined based on visual inspection of the singular values. The weight vector in Subsection II-C is empirically chosen equal to:

$$\begin{bmatrix} 0.3478 & 0.1739 & 0.1304 & 0.0870 & 0.0435 & 0.0217 & 0.0217 & 0.0217 \\ & & 0.0217 & 0.0217 & 0.0217 & 0.0217 & 0.0217 & 0.0217 & 0.0217 \end{bmatrix}$$

to be multiplied with the 15 best similarity scores. The highest score is twice as important as the second one. Note that for interpretability of the resulting features, the weight vector sums up to one.

We show the two-dimensional projections of the morphological features $\mathbf{f} \in \mathbb{R}^3$ for each combination of classes in Figure 2. There is a clear separation between the normal and AF classes on the one hand and the AF and other classes on the other hand. The classes of normal and other signals are however strongly overlapping: many signals in the other category show a large morphological resemblance to normal signals. This could be due to HRV information that is lost in the segmentation, or disagreement between annotators (although a signal is normal).

Figure 3 shows the results of an experiment to evaluate the choice of the size of the model set in Table II. This figure shows the F_1 score based on the straightforward classification method of assigning each test signal to the class with the largest morphological similarity. For each size of the model set, 20 random model sets are constructed of the corresponding size, while the testing happens on the remaining part of the dataset. Note that the largest increase in performance is seen for AF. AF signals show the largest variation and are represented the least in the total dataset. This explains why more examples in the model set lead to a better model. Note




Fig. 2: While the morphological features discriminate well between NSR and AF and between AF and Other, there is a large overlap between NSR and Other. We visualize this via the 2D projections of morphological feature vector $\mathbf{f} \in \mathbb{R}^3$ for the multi-lead dataset. The projections show the datapoints from each combination of classes together with the first bisector.




Fig. 3: The increase in performance when increasing the size of the model set is the largest for the AF class, given the large variation in morphology and the small representation in the total dataset. The size of the model set in the experiments of Table II is chosen well (dotted line), given that the performance does not significantly increase anymore for larger model sets.

that there is no increase in performance for NSR, given that this class is largely represented and corresponds to a very regular rhythm. It also supports the choice made in Table II, given that the total performance does not increase significantly anymore for larger model sets.

We choose an SVM classifier, using one-vs-one coding. Optimization of hyperparameters is done with ten-fold cross-validation, while different kernels are tested. Table III indicates that, without using HRV features, the SVM with a linear kernel obtains the best results with an F_1 score of 0.593. Other kernels give similar results, which is expected because of the distribution of the data (Figure 2). For the other feature sets, an RBF kernel was chosen after validation. *Only* using HRV features, optimally selected to discriminate AF signals, a global F_1 score of 0.732 can be obtained. The optimal F_1 score of 0.77, is achieved by combining the new morphological

TABLE III: Combining the morphological features, that have been obtained by our method, with the classical HRV information results in a higher performance than only using classical HRV information.

Method	$P(\%)$	F_{1n}, F_{1a}, F_{1o}	F_1
HRV	77.67	0.853, 0.760, 0.585	0.732
SVD	70.02	0.813, 0.570, 0.396	0.593
SVD + HRV	80.22	0.866, 0.796, 0.647	0.770

features and classical HRV features, and is comparable to other methods [18]. The top-performing algorithm of the competition achieves a F_1 score of 0.83. Care must be taken when comparing the obtained results with the results of the competition because for the latter an additional noisy class was added and a hidden test set was used. An important advantage of the designed method is the interpretability: the morphological similarity scores are readily interpretable by cardiologists. Furthermore, these features are computed based on well-founded algebraic principles. The decision-making in the case of a linear SVM is interpretable as well, by inspecting the individual coefficients of the SVM. This is different from for example deep learning methods, which are less interpretable because of the complete black-box modeling. This argument should not be underestimated in the attempt to design automatic methods in a clinical context. The normal class has the best classification performance for all classifiers. As expected, most misclassifications are seen between the normal and other class. It could be helpful to divide this other class further in more specific subclasses to obtain a better performance.

B. MLSVD-based detection in multi-lead ECG

The MLSVD-based detection method is tested on a multi-lead dataset, constructed by combining two publicly available datasets: the MIT-BIH Atrial Fibrillation dataset [19], [20] and the AF Termination Challenge dataset [19], [21].

The first dataset consists of 23 two-lead Holter signals, containing signals of patients with paroxysmal or intermittent AF. From these signals, 80 segments of one minute were

TABLE IV: The data profile of the model, training and test set of the MIT-BIH AFIB & AFTDB dataset.

	Class		
	Normal	AF	Total
Model set	30	30	60
Training set	20	20	40
Test set	30	30	60
Total	80	80	160

TABLE V: Even though the AUC does not increase by combining the morphological features, obtained via our method, with classical HRV information, the morphological features by themselves contain a high degree of information.

Method	$P(\%)$	AUC	F_{1n}, F_{1a}	F_1
HRV	98.33	1.000	0.984, 0.983	0.983
MLSVD	93.33	0.990	0.935, 0.931	0.933
MLSVD + HRV	100.00	1.000	1.000, 1.000	1.000

extracted at random, all containing portions of the signal with a normal sinus rhythm. The second dataset, the AF Termination Challenge dataset, contains 80 Holter ECG signals (two leads). All signals are AF signals with a fixed length of one minute.

The final dataset is a combination of all signals from the AFTDB dataset and the 80 normal segments of the MIT-BIH AFIB dataset. The signals from the MIT-BIH dataset were resampled to obtain the same sampling frequency. Table IV gives a summary of all signals in the final dataset. Note that the test set contains, for the AF class, independent subjects from the model and training set.

After preprocessing and per-lead compression, the model tensor $\mathcal{D}_{\text{model}} \in \mathbb{R}^{2 \times 76 \times 60}$ is modeled with a rank-(1, 23, 23), obtained after 20-fold cross-validation. The weight vector is chosen equal to:

$$\begin{bmatrix} 0.3902 & 0.1951 & 0.1463 & 0.0976 & 0.0488 \\ 0.0244 & 0.0244 & 0.0244 & 0.0244 & 0.0244 \end{bmatrix}$$

taking only the best ten scores into account. All computations were performed in MATLAB using Tensorlab [22].

The results in Table V indicate that only by using the morphological features $\mathbf{f} \in \mathbb{R}^2$ a good discrimination is possible between normal and AF signals. We again choose an SVM classifier and optimize the hyperparameters using ten-fold cross-validation. After validation, a linear kernel was used for the MLSVD method. The other two SVM's were trained with an RBF kernel. Figure 4 shows the morphological features for all signals in both the training and test set, independent from the model set, together with the SVM decision boundary. Two clusters can be clearly distinguished: the samples in the upper left part of the plot correspond with signals which show a large morphological similarity with AF signals ($f_{AF} > f_{NSR}$) while the points in the lower right part show a larger correspondence with normal signals ($f_{NSR} > f_{AF}$). The linear decision boundary is capable of separating most of these points.

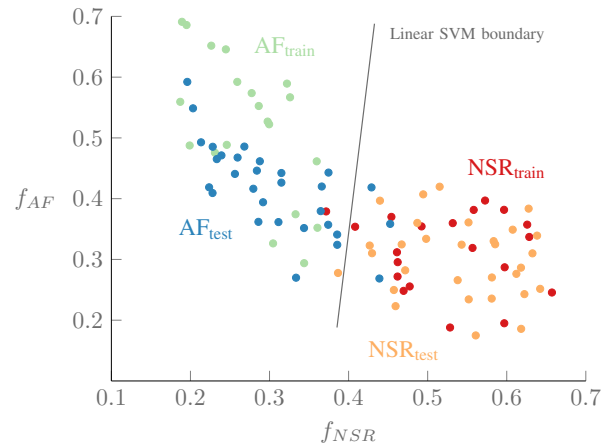


Fig. 4: Using the morphological features $\mathbf{f} \in \mathbb{R}^2$ for the multi-lead dataset, an almost perfect separation is possible: only three false negatives and one false positive are present in the test set. The linear SVM decision boundary is shown, trained and evaluated on the showed training and test sets.

Adding the HRV features results in a perfect classification: the two classes are completely separable. Note that only using HRV features we also obtain an AUC of one, meaning that this feature set is very informative to separate normal and AF signals. While the morphological features do not add extra performance per se, they do contain a lot of discriminative information.

V. CONCLUSION

We presented an SVD-based method to detect AF in single-lead ECG, extending and generalizing it to multi-lead ECG via an MLSVD-based method. ECG-signals are first compressed into one representative heartbeat. A model matrix or tensor is then modeled with the (ML)SVD, which leads to a compressed representation of each signal in the constructed basis. Each new signal can be expressed accordingly in the constructed basis, leading to a morphological similarity with predefined rhythms in the model set. These similarities can be combined with traditional HRV-features to train an SVM to classify new signals. The methods achieved high performance on different datasets, comparable with other methods in literature. In future work, the MLSVD-method could be tested in a more extensive way on a larger dataset, possibly also dividing the other class into more specific rhythms. A second extension is the usage of higher-order discriminant analysis to learn the subspace over the model set in a supervised way, while the current method constructs a basis in an unsupervised, PCA way. A similar method has already been applied to EEG classification [10]. Finally, research could be done to supply coupling between datasets, across different number of leads and technology. In such a method, the model set could be multi-lead ECG, originating from the hospital, while the training and test set could be single-lead, originating from for example mobile health technologies. Such coupling could be achieved by using

the coefficient vectors themselves as input for the SVM, rather than computing the similarity with the signals in the model set.

REFERENCES

- [1] G. Lip, L. Fauchier, S. Freedman, I. Van Gelder, A. Natale, C. Gianni, S. Nattel, T. Potpara, M. Rienstra, H.-F. Tse, and D. Lane, "Atrial fibrillation," *Nature Reviews Disease Primers*, vol. 2, pp. 1–26, 2016.
- [2] V. Fuster, L. E. Rydén, D. S. Cannom, H. J. Crijns, A. B. Curtis, K. A. Ellenbogen, J. L. Halperin, G. N. Kay, J.-Y. Le Huezey, J. E. Lowe, S. B. Olsson, E. N. Prystowsky, J. L. Tamargo, and L. S. Wann, "2011 ACCF/AHA/HRS focused updates incorporated into the ACC/AHA/ESC 2006 guidelines for the management of patients with atrial fibrillation," *Journal of the American College of Cardiology*, vol. 57, no. 11, pp. e101–e198, 2011.
- [3] F. Shaffer and J. Ginsberg, "An overview of heart rate variability metrics and norms," *Frontiers in Public Health*, vol. 5, p. 258, 2017.
- [4] M. Boussé, G. Goovaerts, N. Vervliet, O. Debals, S. Van Huffel, and L. De Lathauwer, "Irregular heartbeat classification using Kronecker product equations," in *Proc. of the 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2017, Jeju Island, South Korea)*, 2017, pp. 438–441.
- [5] L. De Lathauwer, B. De Moor, and J. Vandewalle, "A multilinear singular value decomposition," *SIAM Journal on Matrix Analysis and Applications*, vol. 21, no. 4, pp. 1253–1278, 2000.
- [6] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, 2009.
- [7] A. Cichocki, D. Mandic, L. De Lathauwer, G. Zhou, Q. Zhao, C. Caiafa, and H. A. Phan, "Tensor decompositions for signal processing applications: From two-way to multiway component analysis," *IEEE Signal Processing Magazine*, vol. 32, no. 2, pp. 145–163, 2015.
- [8] N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos, "Tensor decomposition for signal processing and machine learning," *IEEE Transactions on Signal Processing*, vol. 65, no. 13, pp. 3551–3582, 2017.
- [9] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Transactions on Biomedical Engineering*, vol. 32, no. 3, pp. 230–236, 1985.
- [10] S. Van Eyndhoven, M. Boussé, B. Hunyadi, L. De Lathauwer, and S. Van Huffel, "Single-channel EEG classification by multi-channel tensor subspace learning and regression," in *Proc. of the 2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP, 2018)*, 2018, pp. 1–6.
- [11] M. Boussé, N. Vervliet, I. Domanov, O. Debals, and L. De Lathauwer, "Linear systems with a canonical polyadic decomposition constrained solution: Algorithms and applications," *Numerical Linear Algebra with Applications*, vol. 25, no. 6, p. e2190, Aug. 2018.
- [12] K. Tateno and L. Glass, "Automatic detection of atrial fibrillation using the coefficient of variation and density histograms of RR and Δ RR intervals," *Medical and Biological Engineering and Computing*, vol. 39, no. 6, pp. 664–671, 2001.
- [13] A. H. Khandoker, C. Karmakar, M. Brennan, M. Palaniswami, and A. Voss, *Poincaré Plot Methods for Heart Rate Variability Analysis*. Boston, MA: Springer US, 2013.
- [14] S. Sarkar, D. Ritscher, and R. Mehra, "A detector for a chronic implantable atrial tachyarrhythmia monitor," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 3, pp. 1219–1224, 2008.
- [15] G. D. Clifford, R. Colloca, and J. Oster. `Comput_afev`, `comp_drr`, `metrics`, `bpcount`. Available online. [Online]. Available: <https://github.com/pratyusk/AFibClassification>
- [16] D. E. Lake and J. R. Moorman, "Accurate estimation of entropy in very short physiological time series: the problem of atrial fibrillation detection in implanted ventricular devices," *American journal of physiology. Heart and circulatory physiology*, vol. 300, no. 1, 2011.
- [17] K. Lee. (2012) `Apen`. Available online. [Online]. Available: <https://nl.mathworks.com/matlabcentral/fileexchange/32427-fast-approximate-entropy>
- [18] G. D. Clifford, C. Liu, B. Moody, L.-W. H. Lehman, I. Silva, Q. Li, A. E. Johnson, and R. G. Mark, "AF classification from a short single lead ECG recording: the PhysioNet/Computing in Cardiology challenge 2017," in *Proc. of Computing in Cardiology*, vol. 44, 2017, pp. 1–4.
- [19] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. 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. e215–e220, 2000 (June 13), circulation Electronic Pages: <http://circ.ahajournals.org/content/101/23/e215.full> PMID: 1085218; doi: 10.1161/01.CIR.101.23.e215.
- [20] G. Moody, "A new method for detecting atrial fibrillation using RR intervals," *Computers in Cardiology*, pp. 227–230, 1983.
- [21] —, "Spontaneous termination of atrial fibrillation: a challenge from PhysioNet and Computers in Cardiology 2004," *Computers in Cardiology*, vol. 31, pp. 101–104, 2004.
- [22] N. Vervliet, O. Debals, L. Sorber, M. Van Barel, and L. De Lathauwer. (2016, Mar.) `Tensorlab 3.0`. Available online. [Online]. Available: <https://www.tensorlab.net>