

# Classifying Heart Sounds using Multiresolution Time Series Motifs: an exploratory study

Elsa Ferreira Gomes

GECAD- Knowledge Engineering and Decision Support Research Center  
ISEP/IPP-School of Engineering, Polytechnic of Porto, Portugal  
efg@isep.ipp.pt

Alípio M. Jorge  
LIAAD-INESC TEC

DCC-FCUP, Universidade do Porto, Portugal  
amjorge@fc.up.pt

Paulo J. Azevedo  
HASLab / INESC TEC  
Department of Informatics  
pja@di.uminho.pt

## ABSTRACT

The aim of this work is to describe an exploratory study on the use of a SAX-based Multiresolution Motif Discovery method for Heart Sound Classification. The idea of our work is to discover relevant frequent motifs in the audio signals and use the discovered motifs and their frequency as characterizing attributes. We also describe different configurations of motif discovery for defining attributes and compare the use of a decision tree based algorithm with random forests on this kind of data. Experiments were performed with a dataset obtained from a clinic trial in hospitals using the digital stethoscope DigiScope. This exploratory study suggests that motifs contain valuable information that can be further exploited for Heart Sound Classification.

## Keywords

Heart sound classification, motif discovery, time series analysis, SAX.

## 1. INTRODUCTION

The general problem we address is the identification of cardiac pathologies by analyzing the features of heartbeat audio recordings collected from a digital stethoscope or from mobile devices such as smart phones. Our aim is to provide algorithms and methods that are able to perform the first level of screening of cardiac pathologies both in a hospital environment by a doctor (using a digital stethoscope) and at home by the patient (using a mobile device). Such algorithms do what is called Heart Sound Classification.

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*C<sup>3</sup>S<sup>2</sup>E* – 13, 2013 July 10-12, Porto [Portugal]

Editor: Bipin C. DESAI

Copyright 2013 ACM 978-1-4503-1976-8/12/06 ...\$15.00.

The main components of heart sound signal of a normal heart are the first heart sound, *S1* (or lub), corresponding to the systolic period, and the second heart sound, *S2* (or dub), the diastolic period [11]. The problem of Heart Sound Classification is commonly addressed in two steps. The first one is segmentation, where *S1* and *S2* sound segments are located within audio data. The second step is classification and uses the segments discovered in the first step. Data is gathered in real-world situations and frequently contains background noise of every conceivable type. The differences between heart sounds corresponding to different heart symptoms can also be extremely subtle and challenging to separate. Success in classifying this form of data requires extremely robust classifiers. Despite its medical significance, to date this is a relatively unexplored application for machine learning [2].

In our previous work [8], we proposed an approach that followed these 2 steps. The approach has been tested on a dataset obtained from a clinic trial in hospitals using the digital stethoscope DigiScope [18]. This dataset has been used as Dataset B for the PASCAL Classifying Heart Sounds Challenge [2]. These audio files are of varying lengths, between 1 second and 30 seconds (some have been clipped to reduce excessive noise and provide the salient fragment of the sound). The sounds of the dataset are pre-classified into three categories (Normal, Murmur and Extrasystole).

In this paper we describe an exploratory study in the use of SAX-based Multiresolution Motif Discovery for Heart Sound Classification. This technique has never been used in this problem. A time series motif is a segment that is frequently observed. This is the case of heartbeat sounds in a cardiac audio. Different kinds of sounds should correspond to different motifs. The idea of our work is to discover relevant frequent motifs in the audio signals and use the discovered motifs and their frequency as characterizing attributes. We proceed by trying different configurations of motif discovery for defining attributes. We also compare the use of classification algorithms (Decision Trees, Logistic Regression, Rotation Forest and Random Forest) on this kind of data. Our experiments are conducted on the dataset

described above. This exploratory study enables us to conclude that motifs contain valuable information that can be further exploited for Heart Sound Classification.

## 2. BACKGROUND


### 2.1 Heart Sound Segmentation and Classification

The strategy of segmenting the heart sound signal before classification is broadly used. The segmentation of a heart sound signal is quite difficult because the number of sound components may be different and the existence of anomalies like arrhythmia or murmurs are unpredictable. The periods of heart beat cycles are inconsistent in a PCG signal. Attempts to segment phonocardiographic (PCG) signals have been reported in literature. The majority of them exploit electrocardiogram (ECG) signals or/and carotid pulse data. Groch presented a solution where the segmentation was based on the time domain characteristics of the signal [9]. On the other hand, Strunic [19] extracted signals on certain band to reduce anomalies and then set a amplitude threshold to pick out the spikes and realize the segmentation [8]. To achieve classification, Karraz extracted the QRS complex from the signal as features and used them in a Neural Network Classifier based on a Bayesian framework [14]. Strunic integrated all the segmented heart cycles into one average heart cycle and used it to train the Artificial Neural Network (ANN) to classify heartbeat into categories. Babbei [1] also used ANN after using Discrete Wavelet Transform (DWT) Kampouraki [12] used Support Vector Machines (SVMs) to classify ECG recordings [8]. Kao [13] used SVM to classify PCG signals, after segmentation and features extraction with Fourier transform and 2-D discrete Fourier transform. Liang [16] chose Chebyshev type I low-pass filter combined with Shannon energy to attenuate noise and make the findings of low intensity sounds, namely heartbeats, easier [16]. This methodology was also followed by some other authors like Kumar [15] and Gupta [10].

### 2.2 Multiresolution Motif Discovery

The extraction of frequent patterns (motifs) from a time series database is an important data mining task. These patterns provide useful insights to the domain expert about the problem at hand [7] and help summarize the time series database. Motif discovery has been used in different areas namely in health and medicine and in particular in EEG time series a motif may be a pattern that usually precedes a seizure [21]. One recent trend in time series analysis is to use SAX [17]. SAX is a symbolic approach for time series that represents the continuous series as a discretized one. It allows for dimensionality reduction and indexing. In classic data mining tasks such as clustering or classification SAX is as good as well-known representations such as Discrete Wavelet Transform (DWT) and Discrete Fourier Transform (DFT), while requiring less storage space. The representation allows researchers to avail of the wealth of data structures and algorithms in bioinformatics or text mining, and also provides solutions to data mining tasks, such as motif discovery [17].

As a symbolic approximation,  $SAX(T, w, a)$  converts the original real time series  $T$  of length  $n$  into a sequence of  $w$  (word size) symbols belonging to an alphabet of size  $a$ . In Figure 1 we see how a time series segment can be discretized



**Figure 1: SAX discretization of a time series segment - image obtained with the iMotifs tool [5]**

into a sequence of SAX symbols  $\{1, 1, 3, 3, 1, 1, 1, 1\}$ , using the iMotifs tool [5] (a GUI for the MrMotif algorithm). The alphabet size is called resolution. The SAX algorithm starts by reducing the dimensionality of the time series by dividing it into  $w$  segments (word length) with the same length using the Piecewise Aggregate Approximation (PAA) algorithm. This algorithm assigns to each segment its average value. Then, the amplitude of the time series is divided into intervals and then a symbol can be assigned to each interval. To generate equiprobable intervals it uses  $a - 1$  breakpoints, producing the same area under the Normal curve, and symbols are obtained from the intervals. The segments below the smallest breakpoint are assigned the 0 symbol, the segments between the first and second breakpoints the symbol 1, and so forth.

The iSAX representation extends classic SAX by allowing different simultaneous resolutions for the same word. The main idea of the MrMotif algorithm is to start from a low iSAX resolution and then expand to higher resolutions. As this expansion is performed, the number of instances of a given cluster reduces as each cluster is split into several of the next resolution. At the highest resolution, a cluster is formed only if the subsequences in that cluster are very similar, as each iSAX symbol covers only a narrow interval in the amplitude of the time series. The minimum possible resolution  $g_{min}$  in iSAX is 2 and the maximum resolution  $g_{max}$  is assigned to 64 (it uses 2, 4, 8, 16, 32 and 64 resolutions).

Castro et al. [6] have proposed a Multiresolution Motif Discovery in Time Series algorithm (MrMotif) based on the iSAX representation. This algorithm solves the motif discovery problem as an approximate Top-K frequent subsequence discovery problem. The aim of MrMotif algorithm is to find the solution for the problem: given a time series database  $D$ , a motif length  $m$  and a  $K$  parameter, for each resolution in  $(g_{min}, g_{min} \times 2, \dots, g_{max})$  find the top- $K$  frequent motifs. In this work we will use MrMotif for generating motif-based features for heart sound classification.

## 3. DATASET

The dataset we use in this study (PASCAL challenge Dataset B) consists of 312 auscultations gathered using the DigiScope (Figure 2) Collector system [18] deployed in the Maternal and Fetal Cardiology Unit of the Real Hospital Português (RHP) in Recife Brazil, led by Dra. Sandra Mattos who coordinated this data collection task. Each auscultation consists of 6 to 10 seconds recorded for each of the four standard cardiac auscultation spots in children. Relevant

patient and auscultation information was also annotated by clinicians from RHP using the DigiScope Collector system, including the presence of abnormal sounds such as murmurs. Each individual heartbeat was manually segmented by Fabio Hedayioglu, one of the researchers of the DigiScope project.



**Figure 2: The DigiScope Collector system that was deployed in Real Hospital Português, Recife, Brazil, gathering more than 200 auscultations from real patients during routine clinical practice.**


As referred before the dataset has 3 classes: Normal, Murmur and Extrasystole. Normal is denoted as N in this paper and has a count of 200. Murmur (M) counts 66 cases and Extrasystole (E) 46. The relative class imbalance makes the problem harder.

## 4. PREVIOUS PEAK DETECTION BASED APPROACH


### 4.1 Heart Sound Segmentation

In the previous approach we aimed to produce a method for determining the location of S1 and S2 sounds within audio data (top plot in Figure 4), segmenting the Normal audio files in the the dataset. The recorded signals were pre-processed before performing segmentation. In the first step the signals were down sampled and filtered (second plot in Figure 4). In the second step, we performed the segmentation. Our algorithm was based on the envelope calculated using the normalized average Shannon energy [16] (third plot in Figure 4). After obtaining the normalized average Shannon energy curve we identified the peaks. For that, we adapted the open source function `peakdet` ([3]), written in Matlab (fourth plot in Figure 4). This function finds the local maxima (and minima) using the strategy that a point is considered a maximum peak if it has a locally maximal value, and was preceded (to the left) by a value lower than a given delta. Using this strategy, we segmented almost all heart sounds. However, we also need to distinguish between S1 and S2. Our current approach for S1/S2 discrimination is still unsatisfactory. First, we tried to perform the detection of S1 and S2 sounds based on the fact that S2 is longer than S1, for normal heart rates [15]. Bearing this in mind we tried to pick each heart cycle and the corresponding systolic interval. The duration of S1-S2 sounds, or the distance between S1 and S2, was calculated and compared for every segment[11]. The longest interval between two sounds was considered to correspond to the diastolic period and the sound at the right side was assigned as S1 and the sound at

the left side was assigned as S2. Unfortunately, we find that those intervals vary widely from file to file, in our dataset (Figure 3). This happens because these heart sounds have been collected from children of different ages with different heart rates. We also tried to use a similar process to the energy, assuming S2 has higher frequency, but until now, without success.



**Figure 3: Two samples of the heart sounds dataset**



**Figure 4: Peak detection on heart sound signal (x-Time/y-Amplitude)**

### 4.2 Heart Sound Classification

This phase involves feature construction and selection. The aim of this phase of the challenge is to label correctly the sounds in the dataset. After the pre-processing and segmentation of the heart sound signal, some features were extracted. Six features were used. Four of them were extracted from the distances between S1 and S2 (peaks). Assuming that  $sS1$  corresponds to smaller segments and  $sS2$  to the others, the first feature is the ratio of the standard deviation of  $sS1$  over the whole standard deviation. The second is similar for  $sS2$ . The third and fourth features are the ratio of the mean of  $sS1$  ( $sS2$  respectively) over total mean. The fifth feature,  $R_{median}$ , is the ratio of the median of the (three) largest segments in the sample over total mean. The sixth feature,  $R2$ , is the  $r$  square of the array of the sorted

segments of the sample (a measure of linearity). After obtaining the features were applied two different methods from the Weka data mining suite [20]: J48, which generates decision trees, and MLP, the Multi Layer Perceptron.

### 4.3 Results

This approach yielded interesting results and was the best of the three final competitors of the PASCAL Classifying Heart Sounds Challenge. The evaluation was performed on a test set which was hidden to participants. However, there were some limitations in the identification of non-normal heart sounds.

## 5. MULTIREOLUTION MOTIF DISCOVERY IN TIME SERIES APPROACH

Considering the limitations of the previous approach in identifying the non-normal heartbeats, we have explored a very different path. The general idea is to find frequent motifs in the cardiac audio time series using a frequent pattern mining algorithm. Such discovered motifs are regarded as features. Our hypothesis is that these features contain valuable information for discrimination tasks. We believe that cardiac signals from different conformations vary in the kind and/or frequency of the motifs found. To test this hypothesis we first identify relevant motifs in the original dataset and build a new dataset where each relevant motif is an attribute. Then, we compare the predictive ability of motif based modeling with the previous peak based modeling.

### 5.1 Methodology

A motif in a time series is a frequently repeated subsequence (frequent pattern). Figure 1 shows an example of a motif. There are different approaches for this task. We have followed the Multiresolution Motif Discovery (in Time Series) algorithm [5] to detect the common (and relevant) patterns. This algorithm uses the SAX methodology to discretize the continuous signals and looks for patterns in the resulting discrete sequences. In particular, we have used the MrMotif algorithm as implemented by its authors.

For our experimental exploration, we followed the steps below.

1. Apply to the original audio dataset the pre-processing (filters and normalized average Shannon energy) used in the previous approach.
2. Apply the MrMotif algorithm to the resulting time series. In this step we have to chose concrete values to MrMotif’s parameters. We have tried different combinations.
3. Build a dataset where each line corresponds to a line of the original dataset. Each relevant motif found in the previous step is a candidate attribute for this new dataset. The value of such attribute is the frequency of the motif in the corresponding time series.
4. Run a machine learning algorithm on the resulting dataset and estimate the predictive ability of the obtained classifier.

The parameters of the MrMotif implementation are the following: motif length  $m$  - the length of the sliding window that contains the section to discretize in the original time

series; number of motifs generated  $K$  - these are the top- $K$  relevant motifs for each resolution from 4 to 64; word size  $w$  - this is the number of discrete symbols of the iSAX word; and overlap  $o$  the extent to which two windows can overlap. In Figure 5 we can see an example of the processing of one audio from the dataset. The figure shows the starting location of motifs (identified by number) of length 40. We can see that motifs 0 and 15 appear frequently in the time series.




Figure 5: All motifs of length 40 for one audio file (x-Time/y-Amplitude)


In Figure 6 we observe with more detail the location of motif 15 throughout the same series. It is represented in the original series although motif discovery is made on the discretized series. We show motif 15 in alternating colors so that we can better observe separations. Notice that a 10% overlap was allowed. As we can see motif 15 has a specific characteristic that has been automatically discovered and which is clearly visible.

In Figure 7 we can see the location and form of one other motif encountered (motif 0).

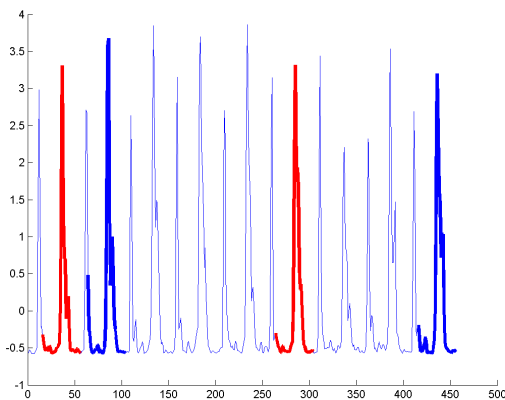
In Figure 1 a sample of one of the dataset is described. Each attribute  $M_i$  is a top-10 motif of resolution 4. Values are frequencies of motifs. The last column is class.

| M1 | M2 | M3 | M4 | M5 | Class |
|----|----|----|----|----|-------|
| 1  | 0  | 5  | 2  | 0  | N     |
| 2  | 2  | 7  | 4  | 3  | N     |
| 4  | 0  | 2  | 2  | 1  | N     |
| 7  | 6  | 4  | 4  | 5  | M     |
| 5  | 4  | 5  | 7  | 5  | M     |
| 2  | 1  | 0  | 1  | 0  | M     |
| 0  | 1  | 3  | 2  | 1  | E     |
| 6  | 4  | 5  | 3  | 3  | E     |
| 2  | 0  | 2  | 2  | 0  | E     |

Table 1: Sample of a resulting dataset of resolution 4, Top-5.



**Figure 6:** Motif 15 shown in thicker line (alternating blue and red) overimposed on the original time series (x-Time/y-Amplitude).



**Figure 7:** Motif 0 shown in thicker line (alternating blue and red) overimposed on the original time series (x-Time/y-Amplitude).

Classifiers were evaluated using the Weka data mining suite [20]. We proceed by describing the conducted experiments.

## 6. EXPERIMENTS

### 6.1 Varying number of motifs

In this first set of experiments we have varied the parameter  $K$  corresponding to the number of motifs selected for attributes. The motif resolution  $R$  was fixed with value 4. For window size we have used 40, and for overlap 10. The size of the SAX word is 8 symbols. In the experiments described below these are the default values. For obtaining the classification models we have first tried the J48 decision tree inducer since it had given good results on the PASCAL Challenge hidden test set. Different classifiers were obtained by its minimal leaf size parameter. As a reference we show the results obtained with the previous peaks approach. We have measured Accuracy, Precision, Recall and ROC Area

Under the Curve (AUC) using 10 fold cross validation. Precision, Recall and AUC are obtained by weighted average over the classes.

**Table 2: Varying number of motifs and minimal leaf size with J48. Comparison with the peaks approach.**

| Dataset | Alg. | Params  | Acc.         | AUC   |
|---------|------|---------|--------------|-------|
| Peaks   | J48  | leaf=2  | 63.46        | 0.589 |
| Peaks   | LR   | -       | 65.06        | 0.667 |
| K10R4   | J48  | leaf=2  | 66.99        | 0.602 |
| K10R4   | J48  | leaf=15 | 69.55        | 0.602 |
| K20R4   | J48  | leaf=1  | 66.03        | 0.628 |
| K20R4   | J48  | leaf=2  | 65.06        | 0.622 |
| K20R4   | J48  | leaf=15 | 70.83        | 0.644 |
| K30R4   | J48  | leaf=5  | 64.42        | 0.601 |
| K30R4   | J48  | leaf=15 | <b>70.51</b> | 0.647 |
| K40R4   | J48  | leaf=10 | 67.63        | 0.625 |
| K40R4   | J48  | leaf=15 | <b>70.51</b> | 0.647 |

As can be observed in Table 2, the motif based representation clearly increased the predictive accuracy of the J48 algorithm over the previous peak based approach. The motif based approach seems to be favoured by a higher value of minimal leaf size, contrarily to the peaks approach. This J48 parameter limits the size of the tree and helps to control overfitting. Since the peak based approach is less prone to overfitting it is less sensitive to this parameter. Despite the relatively good results that the peak approach had obtained with J48 on the Challenge, it performed poorly under cross validation. Therefore, we have tried Logistic Regression. Results have improved for the peak approach.

We have also tried using motif resolutions other than 4. We have not improved the results, neither by using other resolutions (8 and 16) nor by combining attributes corresponding to motifs of resolution 4 with other of resolutions 8 and 16. A possible explanation is that a higher resolution implies lower motif frequencies. This has a negative impact since it leads to more specific models.

### 6.2 Using Random Forests

One shortcoming of the motif based approach is that, for higher resolutions, very similar motifs can be counted as different. These variations are not explored by J48. We have thus followed the path of multiple models (ensembles) and conducted experiments with random forests [4]. This algorithm builds multiple decision trees with different attribute sets sampled from the original dataset. This results in complementary models that together have an effect similar to a probabilistic disjunctive model. The important parameters for the random forest algorithm are the number of models  $I$  and the number of attributes sampled for each model (in the tables we use the letter  $K$  for this parameter as originally used by the Random Forest (RdF) weka implementation, albeit the collision with the motif  $K$  parameter). We have also made experiments with Rotation Forest (RtF), another ensemble approach but the results were not superior to Random Forests.

Best results (Table 3) were obtained with Random Forest (RdF), using 50 models and 2 attributes on a dataset with the top 40 relevant motifs. We have also run the Random Forest algorithm on the peaks approach but results were much worse than with J48 or Logistic Regression. This

**Table 3: Experiments with Random Forest and Rotation Forest**

| Dataset | Alg. | Params     | Acc.         | AUC   |
|---------|------|------------|--------------|-------|
| Peaks   | RdF  | -I 25-K 2  | 60.90        | 0.623 |
| K10R4   | RdF  | -I 20 -K 2 | 67.63        | 0.650 |
| K20R4   | RdF  | -I 25 -K 2 | 69.23        | 0.658 |
| K20R4   | RtF  | -I 25      | 71.15        | 0.692 |
| K30R4   | RdF  | -I 25 -K 2 | 70.83        | 0.671 |
| K30R4   | RdF  | -I 40 -K 2 | 70.83        | 0.674 |
| K30R4   | RtF  | -I 30      | 71.47        | 0.698 |
| K40R4   | RdF  | -I 50 -K 3 | <b>72.44</b> | 0.708 |

was expected since our use of the Random Forest approach was inspired by the specific characteristics of the motif attributes.

### 6.3 Extending the attribute set

We have explored the motif attributes with different resolutions, different number of top motifs and also to vary other MrMotif parameters. The choice of the window size for the discretization process is relevant. If it does not have the adequate scale it does not capture informative motifs. In all experiments above we have used window size 40 and word size 8. Here we describe the experiments done with datasets that combine attribute sets. We combine our most successful conformations of  $K$  and  $R$  with the Peaks attribute set. We also show results for the combination of attribute sets obtained with different resolutions (4 and 8). We use Random Forest (RdF) in all experiments. The number of models ( $-I$ ) is related to the number of attributes.

**Table 4: Experiments with combination of attribute sets**

| Dataset     | Alg. | Params     | Acc.         | AUC   |
|-------------|------|------------|--------------|-------|
| K10R4+Peaks | J48  | leaf=10    | 66.03        | 0.620 |
| K10R4+Peaks | RdF  | -I 16 -K 2 | 69.55        | 0.673 |
| K30R4+Peaks | RdF  | -I 30 -K 2 | 70.51        | 0.700 |
| K30R4+K10R8 | RdF  | -I 70 -K 2 | <b>72.76</b> | 0.682 |
| K40R4+Peaks | RdF  | -I 30 -K 2 | 72.44        | 0.705 |

In Table 4 we can see that the combination of the motif attribute sets and the peaks attribute set yield some gain in accuracy in the case of  $K = 10$ . This means that it may be worthwhile to explore the combination of information coming from both origins. More surprising is the result obtained with the combination of resolutions 4 and 8 (approach K30R4+K10R8). This is slightly better than K40R4 with resolution 4 alone but still it is the best absolute result overall. It can be the case that a higher resolution is too detailed for being used independently but it adds features that are useful for discriminating a few cases. The combination of motifs obtained with windows of size 40 and 80 did not result in visible improvements (results not shown).

### 6.4 Comparing discriminative power

We now compare the discriminative power of the best peak-based approaches (J48 and Logistic Regression, LR) with the best motif based one (Random Forest -I 70 -K 2 on K30R4+K10R8 dataset). In Table 5 and Table 6 we show values of precision and recall per class for the three methods

(respectively). It can be observed that with motifs there is an increase in precision and recall for the two most representative classes (Normal and Murmur) over both peak based approaches. For class Extrasystole the motif approach simply ignores it, as it is the case with the peak-based Logistic Regression. These results are not entirely satisfactory since the two abnormal classes are important to detect. However, the proposed method deals well with both normal and murmur classes.

**Table 5: Precision per class for the best peak based and motif based approaches**

| Class                | Precision |         |       |
|----------------------|-----------|---------|-------|
|                      | peak-J48  | peak-LR | motif |
| N-Normal (64%)       | 0.670     | 0.671   | 0.710 |
| M-Murmur (21%)       | 0.553     | 0.500   | 0.861 |
| E-Extrasystole (15%) | 0.154     | 0.000   | 0.000 |

**Table 6: Recall per class for the best peak based and motif based approaches**

| Class                | Recall   |         |       |
|----------------------|----------|---------|-------|
|                      | peak-J48 | peak-LR | motif |
| N-Normal (64%)       | 0.875    | 0.950   | 0.980 |
| M-Murmur (21%)       | 0.318    | 0.197   | 0.470 |
| E-Extrasystole (15%) | 0.043    | 0.000   | 0.000 |

## 6.5 Summary of results

Table 7 contains a selection of our main results. Our exploratory study with dataset B of the Pascal challenge leads us to believe that motifs can contribute to the problem of automatic classification of cardiac sounds. This approach yields a significant improvement when compared with the result obtained by the previous approach. Moreover, we have obtained better results with Random Forests, which seem particularly appropriate for the motif based approach. The combination of resolutions 4 and 8 also gave interesting results. Looking at the discriminative power per class we can verify that the use of motifs improves the identification of the normal class and of the most frequent abnormal class (M).

**Table 7: Summary of Experiments**

| Dataset     | Alg. | Params  | Acc.         | AUC          |
|-------------|------|---------|--------------|--------------|
| Peaks       | J48  | leaf=2  | 63.46        | 0.589        |
| Peaks       | LgR  | -       | 65.06        | 0.667        |
| K30R4       | J48  | leaf=15 | 70.51        | 0.647        |
| K40R4       | J48  | leaf=15 | 70.51        | 0.647        |
| Peaks       | RdF  | -I25-K2 | 60.90        | 0.623        |
| K40R4       | RdF  | -I50-K3 | 72.44        | <b>0.708</b> |
| K30R4+K10R8 | RdF  | -I70-K2 | <b>72.76</b> | 0.682        |

## 7. CONCLUSIONS

In this paper we have presented results on the novel application of SAX-based motif discovery in time series to the problem of heart sound classification. Our main conclusion is that there is a strong indication that SAX based motifs

contain valid information which is enough to beat a challenge winning previous approach based on peak detection. In fact, we may postulate that motif discovery is a generalization of the peak detection process. Thus, motif based attributes yield a thinner characterization of heart sounds. Another interesting observation is the predictive ability of the Random Forest - a multiple model technique - when compared to single decision trees for the datasets with motif attributes. One possible explanation is that the multiple model technique is able to disjunctively combine information conveyed by different motifs. The combination of different resolutions can yield better results. The combination of motif attributes and peak based attributes also resulted in slightly better performance.

Finally, the best motif based configuration improves both precision and recall of the normal class and the most frequent abnormal class with respect to the best peak based model.

## 7.1 Future Work

Our next step will be to perform a wider evaluation using multiple datasets which will be soon available from the DigiScope project. We will also make a thorough sensitivity analysis of the approach with respect to the value of some key motif generation parameters such as window length, overlap, resolution and classification algorithm parameters.

There are also important methodological improvements to be made. In this exploratory approach we have applied motif discovery to the initial dataset and produced a complete new dataset on which we have done the cross validation evaluation. However, in a real application, only training data for motif is available and we will have to identify the discovered motifs in the test/deployment data. From our results we can see that this was clearly the case for relevant motifs. Nevertheless, it is something to be considered.

To improve this work, we will explore a second order approach that is able to relate multiple motifs. Possible paths are the clustering of motifs for smoothing spurious differences between motifs. Also, motifs of motifs for identifying frequent motif combinations is an idea that deserves to be explored.

## 8. ACKNOWLEDGMENTS

This work is part-funded by the ERDF - European Regional Development Fund through COMPETE Programme (operational programme for competitiveness), by the Portuguese Funds through the FCT (Portuguese Foundation for Science and Technology) within project FCOMP - 01-0124-FEDER-022701.

This work is also supported by FEDER Funds through the "Programa Operacional Factores de Competitividade - COMPETE" program and by National Funds through FCT "Fundação para a Ciência e a Tecnologia" under the project: FCOMP-01-0124-FEDER-PEst-OE/EEI/UI0760/2011.

## 9. REFERENCES

- [1] S. Babaei and A. Geranmayeh. Heart sound reproduction based on neural network classification of cardiac valve disorders using wavelet transforms of pcg signals. *Comp. in Bio. and Med.*, 39(1):8–15, 2009.
- [2] P. Bentley, G. Nordehn, M. Coimbra, and S. Mannor. The PASCAL Classifying Heart Sounds Challenge 2011 (CHSC2011) Results. [www.peterjbentley.com/heartchallenge](http://www.peterjbentley.com/heartchallenge), 2011.
- [3] Billauer. [www.billauer.co.il/peakdet.html](http://www.billauer.co.il/peakdet.html), 2011.
- [4] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- [5] N. Castro. Multiresolution motif discovery in time series website. <http://www.di.uminho.pt/castro/mrmotif>.
- [6] N. Castro and P. J. Azevedo. Multiresolution Motif Discovery in Time Series. In *SDM*, pages 665–676, 2010.
- [7] P. G. Ferreira, P. J. Azevedo, C. G. Silva, and R. M. M. Brito. Mining approximate motifs in time series. In *Discovery Science*, pages 89–101, 2006.
- [8] E. Gomes, P. Bentley, M. Coimbra, E. Pereira, and Y. Deng. Classifying Heart Sounds: approaches and results for the PASCAL Challenge. In *6th International Conference on Health Informatics, HealthInf 2013 6th International Conference on Health Informatics, Barcelona, Spain*, 2013.
- [9] M. W. Groch, J. R. Domnanovich, and W. D. Erwin. A new heart-sounds gating device for medical imaging. *IEEE Transactions on Biomedical Engineering*, 39:307–310, 1992.
- [10] C. N. Gupta, R. Palaniappan, S. Rajan, S. Swaminathan, and S. M. Krishnan. Segmentation and classification of heart sounds. In *Canadian Conference on Electrical and Computer Engineering*, pages 1674–1677, 2005.
- [11] C. N. Gupta, R. Palaniappan, S. Swaminathan, and S. M. Krishnan. Neural network classification of homomorphic segmented heart sounds. *Applied Soft Computing*, 7:286–297, 2007.
- [12] A. Kampouraki, G. Manis, and C. Nikou. Heartbeat Time Series Classification With Support Vector Machines. *IEEE Transactions on Information Technology in Biomedicine*, 13:512–518, 2009.
- [13] W.-C. Kao, L.-W. Cheng, C.-Y. Chien, and W.-K. Lin. Robust brightness measurement and exposure control in real-time video recording. *IEEE T. Instrumentation and Measurement*, 60(4):1206–1216, 2011.
- [14] G. Karraz and G. Magenes. Automatic Classification of Heartbeats using Neural Network Classifier based on a Bayesian Framework. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 4016–4019, 2006.
- [15] D. Kumar, R. Carvalho, M. Antunes, R. Gil, J. Henriques, and L. Eugenio. A New Algorithm for Detection of S1 and S2 Heart Sounds. In *International Conference on Acoustics, Speech, and Signal Processing*, volume 2, 2006.
- [16] H. Liang, S. Lukkarinen, and I. Hartimo. Heart sound segmentation algorithm based on heart sound envelopogram. *Proceedings of The IEEE*, pages 105–108, 1997.
- [17] J. Lin, E. Keogh, S. Lonardi, and P. Patel. Finding motifs in time series. In *Proceedings of the 2nd Workshop on Temporal Data Mining*, pages 53–68, 2002.
- [18] D. Pereira, F. Hedayioglu, R. Correia, T. Silva, I. Dutra, F. Almeida, S. Mattos, and M. Coimbra. DigiScope - unobtrusive collection and annotating of auscultations in real hospital environments. In *Engineering in Medicine and Biology Society, EMBC*,

*2011 Annual International Conference of the IEEE*,  
pages 1193–1196, 30 2011-sept. 3 2011.

- [19] S. L. Strunic, F. Rios-gutiérrez, R. Alba-flores, G. Nordehn, and S. Burns. Detection and Classification of Cardiac Murmurs Using Segmentation Techniques and Artificial Neural Networks. In *IEEE Symposium on Computational Intelligence and Data Mining*, pages 128–133, 2007.
- [20] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. 2005.
- [21] D. Yankov, E. J. Keogh, J. Medina, B. Y. chi Chiu, and V. B. Zordan. Detecting time series motifs under uniform scaling. In P. Berkhin, R. Caruana, and X. Wu, editors, *KDD*, pages 844–853. ACM, 2007.