

Morphological Classification of Heartbeats Using Similarity Features and a Two-Phase Decision Tree

F Chiarugi¹, D Emmanouilidou^{1,2}, I Tsamardinos^{1,2}, IG Tollis^{1,2}

¹Foundation for Research and Technology - Hellas, Institute of Computer Science, Heraklion, Crete, Greece

²University of Crete, Computer Science Department, Heraklion, Crete, Greece

Abstract

Significant clinical information can be obtained from the analysis of the dominant beat morphology. In such respect, the identification of the dominant beats and their averaging can be very helpful, allowing clinicians to perform the measurement of amplitudes and intervals on a beat much cleaner from noise than a generic beat selected from the entire ECG recording.

In this paper an algorithm for the morphological classification of heartbeats based on a two-phase decision tree is described. Similarity features extracted from every beat are used in the decision trees for the identification of different morphological classes for the beats of the ECG signal.

The results, in terms of dominant beat discrimination, have been evaluated on all annotated beats of the MIT-BIH Arrhythmia Database with Sensitivity = 99.05%, Specificity = 93.94%, Positive Predictive Value (PPV) = 99.32% and Negative Predictive Value (NPV) = 91.69%. Satisfactory results have been also obtained on all the detected beats of the same database using an already published QRS detector developed by the same authors and obtaining Sensitivity = 98.71%, Specificity = 93.81%, PPV = 99.30% and NPV = 89.11%.

1. Introduction

HEARTFAID is a Research and Development project aimed at devising, developing and validating an innovative knowledge based platform of services, able to improve early diagnosis and to make more effective the medical-clinical management of heart failure (HF) within elderly population. ECG is one of the very basic examinations performed in the evaluation and assessment of HF. During the initial design of the HEARTFAID platform it was not defined which ECG devices with the capability of transferring raw data to a host computer in an interoperable way (functional and semantic interoperability) would be available at the clinical validation sites. However, there was the confidence that

some ECG device with raw data interoperability would have been available at some clinical site. For this reason it was decided to start with the design of some generic and robust ECG processing algorithms for QRS detection and classification that had the possibility to be adapted to different sample frequencies, ECG lead number and signal resolution and, being not available any ECG record from the clinical sites, it was decided to start the design using, as learning set, ECG recordings from freely available annotated databases.

At an advanced stage of the design and implementation, after interviews with clinicians and an accurate inventory of the medical devices available at the validation sites, a significant operative scenario was identified, where ECGs acquired with an interoperable non-interpretive electrocardiograph are transferred to the hospital gateway and from there processed in order to detect the QRS complexes, identify the dominant beats and evaluate the averaged dominant beat for each lead.

The evaluation of heartbeats is usually performed considering for each beat its morphology and its occurrence compared to the previous and following beats. A complete rhythm analysis would require the evaluation of both atrial and ventricular rhythm, but this is beyond the scope of this paper focused on the morphological classification of heartbeats, that is a fundamental part of the overall rhythm analysis. It has to be noticed that significant clinical information can be obtained from the analysis of the dominant beat morphology. In such respect, the identification of the dominant beats and their averaging can be very helpful allowing clinicians to perform the measurement of amplitudes and intervals on a beat much cleaner from noise than a generic one selected from the entire ECG recording.

In this paper we describe an algorithm for the morphological classification of heartbeats based on a two-phase decision tree. In the first phase, a first classification of all heartbeats is performed, while, in the second phase, all non-dominant classes are re-processed and re-evaluated. Similarity features are extracted from every beat and used in the decision trees for the

identification of the different morphological classes and the assignment of each beat to the proper class. The noise associated to each heartbeat, which cannot be completely removed in the pre-filtering phase, has been taken into account, slackening some decision tree thresholds in case of significant residual noise in the current heartbeat.

The algorithm has been evaluated on the entire MIT-BIH Arrhythmia Database [1] and then also tested on a short-term ECG database provided by the Department of Cardiology, University Magna Graecia, Catanzaro, Italy in the context of the HEARTFAID EU project with satisfactory results.

2. Methods

Real data (surface ECGs) have been used from the MIT-BIH Arrhythmia Database [1]. The database consists of 48 half-hour excerpts of two-channel ambulatory ECG recordings obtained from 47 subjects. The recordings are digitized at 360 Hz with 11-bit resolution over a 10 mV range. Reference annotation for every beat of each record is available, identifying the QRS position and type.

A signal pre-filtering has to be applied to the raw data before the classification takes place, aiming at the removal of baseline wandering and residual noise that may affect the heartbeat morphology and hamper a proper classification. Each channel is first sub-sampled at 1/5 of the original frequency in order to pass it, with satisfactory execution time, through a chain of median filters of 200 msec and 700 msec respectively [2]. The first median filter has the goal of suppressing the QRSs while the second eliminates the T waves. The output of the filter chain is a signal representing the course of the baseline. This signal, low-pass (1 Hz) filtered and re-sampled at the original frequency, is subtracted from the original one providing the original signal where the baseline wandering has been cancelled. Then, the signal is processed with a band-pass filter (1-40 Hz) in order to reduce the residual noise. Figure 1 shows ECG signals from a record of the MIT-BIH Arrhythmia Database before and after the application of the described pre-filtering. It can be easily noticed how the baseline wandering has been removed in the filtered signals.

The classification of the pre-filtered signal is a typical problem of unsupervised clustering since information about the number of the existing classes and the morphology of the dominant QRST complexes is not available at the start of the algorithm. The implemented unsupervised classification algorithm was based on a two-phase decision tree: in the first phase an initial classification of all beats was performed, while in the second phase the created classes were re-estimated and, if necessary, redefined.

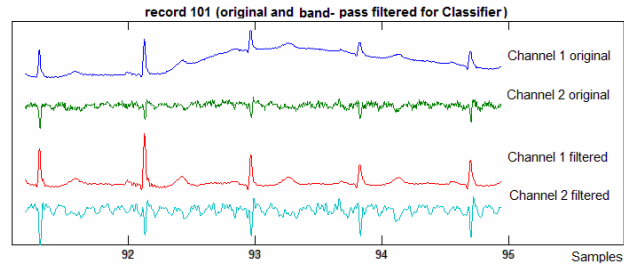


Figure 1: Excerpts from record 101 of the annotated database. The original signal is in dark blue (lead 1) and in green (lead 2); the filtered signal is in red (lead 1) and in turquoise blue (lead 2). Different scales are used for the original and filtered signals.


In order to begin the first phase of the clustering algorithm, the positions (samples) of all R peaks of the signal have to be provided. For the MIT-BIH Arrhythmia Database all R peaks have been annotated by referees and can be used as input.

In the first phase of the algorithm, a rough centroid (template) is calculated for all beats of the signal, one for each lead, with duration 0.6 s. Each sample of this template is formed by the mean values of all beats for the specified lead (excluding the distribution tails) and, once this template is built, it is compared with every beat of the signal in order to obtain the L1 (or cityblock) distance, normalized with a $(p2pX_{beat}/p2pX_{centr})$ factor; where $p2pX_{beat}$ stands for the peak-to-peak value of the beat under consideration and the $p2pX_{centr}$ stands for the peak-to-peak value of the centroid. However, each single QRST complex in the recording can have a shorter duration due to the premature occurrence of the next QRST complex. Thus, the duration of each QRST complex is estimated based on the occurrence of the next R peak and the comparison with the centroid (in terms of normalised L1 distance) is then limited to the number of samples the beat under consideration has (the minimum between 0.6 s and the estimated QRST duration). The most similar beats according to the L1 distance from both leads are extracted and used for the calculation of a new centroid, which represents a more accurate estimation of the average dominant QRST complex.

Once this centroid has been obtained, the normalized L1 distance of each complex from the new centroid is re-evaluated. This time horizontal and vertical wiggling is used in order to have a better alignment between the QRST complex under investigation and the estimated centroid and to produce more reliable values. The alignment is obtained in the position that minimizes the L1 distance from the centroid. Along with L1, the L2 distance (both normalized with the $(p2pX_{beat}/p2pX_{centr})$ factor) is also calculated and these two features, combined with the centroid-to-beat correlation

coefficient, are the main features for the “similarity” evaluation of the QRST complex under investigation. Using these 3 features and the peak-to-peak value of the beat, it is determined with a decision tree whether a complex should be considered dominant.

Then, the same algorithm is applied again to the non-dominant beats (remaining beats). A rough centroid is formed and compared to the remaining beats in order to find their main class. Using the L1 distance a more accurate centroid is computed and each beat is compared to it and classified. The beats that do not fall in the main class are the new remaining beats and are reprocessed in the same way. The algorithm stops when all beats have been classified or no further classification (of the residual beats) can be made and thus, a final group with all residual unclassified beats is additionally formed. Figure 2 shows the decision tree related to phase 1.



$L1d$ =matrix (nbeatsx2) of the L1 distances of each beat from the estimated centroid (with best L1 alignment)
 $L2d$ =matrix (nbeatsx2) of the L2 distances of each beat from the estimated centroid (with best L1 alignment)
 ccv =matrix (nbeatsx2) of the correlation coefficient value of each beat with the estimated centroid (with best L1 alignment)
 ppv =matrix (nbeatsx2) of the peak-to-peak value of each beat (after best L1 alignment with the estimated centroid)
 NSR =matrix (nbeatsx2) of the Noise-to-signal ratio for each beat

$T1=((ccv(1) \geq thrccv1(1)) \&\& (ccv(2) \geq thrccv1(2))) \vee ((ccv(1) \geq thrccv2(1)) \&\& (ccv(2) > 0.2)) \vee ((ccv(2) \geq thrccv2(2)) \&\& (ccv(1) > 0.2)) \vee ((ccv(1) > 0.96) \vee (ccv(2) > 0.96));$
 $T2=((L1d(1) \leq thr1d1(1)) \&\& (L1d(2) \leq thr1d1(2))) \vee ((L2d(1) < thr1d2(1)) \vee (L2d(2) < thr1d2(2)));$

Figure 2. The decision tree of phase 1.

In the second phase, all non-dominant groups are re-

processed. The beats are not examined one by one as in the first phase, mainly for reasons of time performance and also because a first grouping has been already performed. The groups containing non-dominant beats (according to the first phase) that are large in number are split into smaller ones, according to their L1 distance from the dominant group, and their status is re-evaluated. The centroid of each reformed group is compared to the centroid of the dominant group using the same features of the first phase. Only the beats of the groups whose classification criteria are satisfied are put into the dominant class. The structure of the classification criteria is similar to the one used in the first phase, but the thresholds are different. The last group of phase 1, with the remaining unclassified beats, is usually more a collection of leftover complexes than a group of similar complexes. For this group each beat is separately compared to the centroids of all groups formed so far. If a satisfactory likeness with any other class is not obtained, then a new group is formed consisting of that beat. With this last step the final separation between dominant and non-dominant beats is ultimately obtained.

Figure 3 illustrates a paradigm of the execution process during the two-phase algorithm while processing record number 100 of the MIT-BIH Arrhythmia Database. At the end of the first phase of the algorithm, 4+1 classes had been identified (the last class is an agglomerate of unclassified beats and its centroid has no real significance). Classes from 2 to 5 were re-evaluated in the second phase, resulting to the formation of 3 classes as shown in the lower figure. In this case only the beat of group 2 was misclassified as non-dominant, due to its residual significant noise in both channels overlapped mainly to the T-wave.

Both channels are processed for each beat in order to evaluate its morphological class during both phases of the algorithm. This means the similarity features extracted for each beat should satisfy the criteria’s thresholds for both channels but noise can occur differently in the two channels. More specifically, there are cases where noise appears in the same segments in both channels, cases where it appears in segments of channel 1 and not in the same segments of channel 2, or vice versa. Due to this fact, the thresholds used to evaluate the assignment or not of a beat to a group have been lightened, according to the beat’s noise levels estimated using the same algorithm described in the QRS detector developed in [3]. That is, if the noise value of the beat currently processed is large, the criteria are lightened much more than in case of a beat having a small noise value. The general meaning is that more consideration and trust should be given to the channel with less noise. In particular, the noise value is used for the processing of the criteria during phase 1, and for the processing of the last group during phase 2.

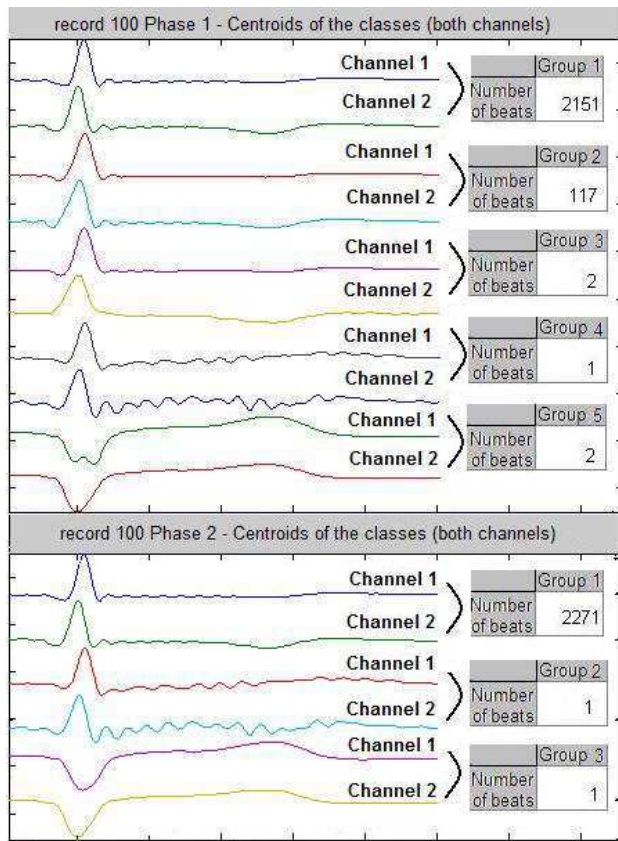


Figure 3. Results during the two-phase classifier of record 100. Upper figure: phase 1 gives 4+1 classes. Lower figure: the final classes formed in phase 2. In both figures, centroids of all classes, for both leads, and the number of beats belonging to each class are shown.

3. Results

The performances of the classifier have been evaluated in terms of discrimination between dominant and non-dominant beats. On the MIT-BIH Arrhythmia Database the results have been very satisfying on all the annotated QRSs (109494 beats) with Sensitivity = 99.05%, Specificity = 93.94%, Positive Predictive Value (PPV) = 99.32% and Negative Predictive Value (NPV) = 91.69%. Evaluating the algorithm on QRSs detected by the QRS detector presented in [3], a slight reduction of the performances can be noticed. This is mainly due to the misalignment in the fiducial point detected by the detection algorithm that cannot be always recovered without increasing too much the execution time. However, the results are still very satisfactory on all the detected QRSs (109439 beats) with Sensitivity = 98.71%, Specificity = 93.81%, PPV = 99.30% and NPV = 89.11%. In such evaluation all the False Positives (FP) of the QRS detector were treated as non-dominant beats.

The algorithm has also been tested on a small testing set provided by the Department of Cardiology, University Magna Graecia, Catanzaro, Italy and consisting of 63 short-term 12-lead ECG files sampled at 500 Hz and acquired from HF patients. Results on all annotated QRSs (685 beats) have been Sensitivity = 99.54%, Specificity = 100.00%, PPV = 100.00% and NPV = 91.67% and on all detected QRSs (690 beats) Sensitivity = 99.54%, Specificity = 97.37%, PPV = 99.85% and NPV = 92.50%. Once again, all the FP of the QRS detector were treated as non-dominant beats.

4. Discussion and conclusions

The algorithm presented has been designed and implemented considering easy usability and extensibility as requirements and trying to keep both execution time and complexity to an acceptable level. The algorithm produces very satisfactory results. Much more analysis would be necessary in order to improve the performances, but the good results on the small testing set probably do not justify such an effort at this moment.

Acknowledgements

This work was partially supported by the European Project HEARTFAID "A knowledge based platform of services for supporting medical-clinical management of the Heart Failure within the elderly population" (IST-2005-027107).

The authors would like to thank Dr. A. Sciacqua of the Department of Cardiology, University Magna Graecia, Catanzaro, Italy and the entire HEARTFAID consortium for their valuable collaboration and help.

References

- [1] MIT-BIH Arrhythmia Database, <http://www.physionet.org/physiobank/database/mitdb/>.
- [2] de Chazal P, O'Dwyer M, Reilly RB. Automatic classification of heartbeats using ECG morphology and heartbeat interval features. *IEEE Trans Biomed Eng.* 2004; 51: 1196-1206.
- [3] Chiarugi F, Sakkalis V, Emmanouilidou D, Krontiris T, Varanini M, Tollis IG. Adaptive Threshold QRS Detector with Best Channel Selection Based on a Noise Rating System. *Computers in Cardiology* 2007;34: 157-160.

Address for correspondence

Franco Chiarugi
 Foundation for Research and Technology - Hellas,
 Institute of Computer Science,
 Biomedical Informatics Laboratory,
 P.O. Box 1385, N. Plastira 100, Vassilika Vouton
 GR 70013, Heraklion, Crete, Greece

E-mail address: chiarugi@ics.forth.gr