31st Annual International Conference of the IEEE EMBS
Minneapolis, Minnesota, USA, September 2-6, 2009

A Personalized Classification System for Holter Registers

Serkan Kiranyaz, Turker Ince, Jenni Pulkkinen and Moncef Gabbouj

Abstract— **In this paper we present a personalized long-term electrocardiogram (ECG) classification framework, which can be applied to any *Holter* register recorded from an individual patient. Due to the massive amount of ECG beats in a *Holter* register, visual inspection is quite difficult and cumbersome, if not impossible. Therefore the proposed system helps professionals to quickly and accurately diagnose any latent heart disease by examining only the representative beats (the so called master key-beats) each of which is automatically extracted from a time frame of homogeneous (similar) beats. We tested the system on a benchmark database where beats of each *Holter* register have been manually labeled by cardiologists. The selection of the right master key-beats is the key factor for achieving a highly accurate classification and thus we used exhaustive K-means clustering in order to find out (near-) optimal number of key-beats as well as the master key-beats. The classification process produced results that were consistent with the manual labels with over 99% average accuracy, which basically shows the efficiency and the robustness of the proposed system over massive data (feature) collections in high dimensions.**

I. Introduction

LONG-term continuous electrocardiogram (ECG) monitoring and recording, also known as *Holter* electrocardiogram or *Holter* register [4], is needed for detection of some diseases, such as *cardiac arrhythmias, transient ischemic episodes* and *silent myocardial ischemia*, and for arrhythmic risk assessment of patients [11]. Since visual analysis of long-term recordings of the heart activity, with more than 100,000 ECG beats in a single recording, is difficult to diagnose and can be highly error prone, automated computer analysis is of major importance. Most of the *Holter* classification techniques presented up-to-date mainly suffer from the usage of suboptimal clustering algorithms, such as Max-Min in [15], *k-medians* in [2] and SOMs in [8], some of which require *a priori* setting of some thresholds or parameters, such as $\theta = 50$ in [15]. Particularly, the performance of the approach in [8] is limited by the ability of a small number of Hermite expansion coefficients used for the approximation of the heartbeats. It is worth noting that although all these techniques claim to address the problem of long-term (*Holter*) ECG classification, *none* has really been applied to a real *Holter* register, probably due to such limitations.

In order to alleviate the problems of the aforementioned sub-optimum clustering schemes, in this paper we used exhaustive K-means clustering with the purpose of finding out the true (optimal) number of clusters and their centroids. To assess and find out the best K-means run among exhaustive number of trials, we performed cluster validity analysis, which is the assessment of the clustering method's output using a specific criterion for optimality, i.e. the so-called clustering validity index (CVI). Hence, we used a simple yet efficient CVI in order to assess the clustering performance of each K-means run with a given K value, which is also varied within a practical range. The particular K-means run with the best CVI score is then used for determining the representative beats, or the so-called key-beats. The proposed clustering approach is then applied over a real (benchmark) dataset, which contains 7 long-term electrocardiogram (ECG) recordings [13] to obtain semi-automatic classification (labeling). Such ambulatory ECG recordings with a typical duration of 24 to 48 hours, are particularly useful for estimating the risk of *ventricular arrhythmias*, such as *premature ventricular contractions (PVCs)*, in patients with heart disease, which may not be detected by a short-time ECG [11]. Yet any process that requires humans or even an expert cardiologist to examine more than a small amount of data can be highly error prone. A single record of a *Holter* register is usually more than 100,000 beats, which make the visual inspection almost infeasible, if not impossible. Therefore, the need for automatic techniques for analyzing such a massive data is imminent and in that, it is crucial not to leave out significant beats since the diagnosis may depend on just a few of them. However, the dynamic range and intra-signal variation in a typical *Holter* register are quite low and abnormal beats, which may indicate the presence of a potential disease, can be scattered along the signal. So based on the proposed exhaustive K-means clustering, a systematic approach is developed, which can summarize a long-term ECG record by discovering the so-called master key-beats that are the representative or the prototype beats from different clusters. With a great reduction in effort, the cardiologist can then perform a quick and accurate diagnosis by examining and labeling only the master key-beats, which in duration are no longer than 15 minutes of ECG record (for a Holter register of 24-48 hours). The expert labels over the master key-beats are then back-propagated over the entire ECG record to obtain a patient-specific, long-term ECG classification.

II. Exhaustive K-means Clustering

As the process of identifying natural groupings in a multidimensional data based on some distance metric (e.g. *Euclidean*), data clustering can be divided into two main

This work was supported by the Academy of Finland, project No. 213462 (Finnish Centre of Excellence Program (2006 - 2011)

categories: hierarchical and partitional [3]. Each category then has a wealth of sub-categories and different algorithmic approaches for finding the clusters. Clustering can also be performed in two different modes: hard (or crisp) and fuzzy. In the former mode, clusters are disjoint, non-overlapping and any data point belongs to a single cluster whereas in the latter case it can belong to all the clusters with some degree of membership [7]. *K-means* [17] is a well known and widely used clustering method, which first assigns each data point to one of the *K* cluster *centroids* and then updates them to the *mean* of their associated points. Starting from a random set of *K* centroids, this cycle is then iteratively performed until the convergence criteria, $\Delta_{Kmeans} < \varepsilon$ is reached where the objective function, Δ_{Kmeans} can be expressed as,

$$\Delta_{Kmeans} = \sum_{k=1}^{K} \sum_{x_p \in c_k} \left\| c_k - x_p \right\|^2 \qquad (1)$$

where c_k is the k^{th} cluster center, x_p is the p^{th} data point in cluster c_k and $\|.\|^2$ is the distance metric in the *Euclidean* space. As a hard clustering method, *K-means* is one of the fastest, i.e. *O(n)*, method but suffers from the following drawbacks:

- The number of clusters, *K*, needs to be set in advance.
- The performance of the method depends on the initial (random) centroid positions as the method converges to the closest local optima.
- The method is also dependent on the data distribution.

K-means, is the one of the simplest and the fastest clustering technique; however, it has severe drawbacks, particularly a single run of *K-means* where the centroids are randomly initialized over N-D data space, is bound to get trapped to the closest local optimum, and the optimal *K* is an unknown that should be determined within the process. Taking its speed advantage into account, we run K-means *exhaustively* (significant number of times, e.g. > 100 with random initializations) for each *K* within a certain range, i.e. $K_{\min} \le K \le K_{\max}$ in order to increase significantly the probability of converging to a (near-) optimal solution. Among all, we then use the "best" K-means run to find out the true (number of) clusters. For the assessment of the clustering performance, the following CVI is used to obtain computational simplicity with minimal or no parameter dependency,

$$f(K,Z) = Q_e K^{\alpha} \quad where \quad Q_e = \frac{1}{K} \sum_{j=1}^{K} \frac{\sum_{\forall z_p \in C_j} \left\| c_j - z_p \right\|^2}{N(C_j)} \qquad (2)$$

where Q_e is the quantization error (or the average intra-cluster distance) as the *Compactness* term, K^{α} is the *Separation* term, by simply penalizing higher cluster numbers with an exponential, $\alpha \ge 0$ and $N(C_j)$ is the number of items in cluster C_j. For $\alpha = 0$, CVI simply becomes the Q_e and using $\alpha = 1$ yields the simplest form (i.e. only the numerator of Q_e). On the other hand, (hard) clustering has some constraints. Let $C_j = \{c_j\}$ be the set of data points assigned to a (potential) cluster centroid. The partitions C_j, $\forall j \in [1, K]$ should maintain the following constraints:

- Each data point should be assigned to one cluster set, i.e. $\bigcup_{j=1}^{K} C_j = Z$
- Each cluster should contain at least one data point, i.e. $C_j \neq \{\phi\}$, $\forall j \in [1, K]$
- Two clusters should have no common data points, i.e. $C_i \cap C_j = \{\phi\}$, $i \neq j$ *and* $\forall i, j \in [1, K]$

As a hard clustering method, K-means is immune to the 1st and 3rd constraints; however, it may fail the 2nd constraint (the so-called under-clustering) especially in high dimensions. Therefore, if any K-means run violates this constraint, that run is simply discarded.

III. The Proposed Framework for Personalized Holter Classification

In this section we shall describe in detail the systematic approach for personalized classification of long-term ECG data. As the overview shown in Figure 1, the proposed system addresses the problem within the entire life-time of a long-term ECG signal recorded from an individual patient, i.e. starting with data acquisition and pre-processing, to the temporal segmentation, followed with a master key-beat extraction by two-pass exhaustive K-means clustering and finally, classification of the entire ECG data by back propagating the expert cardiologist labels over the master key-beats. As a *personalized* approach, the objective is to minimize as much as possible the amount of data from each individual patient by selecting the most relevant data, which will be subject to manual classification, so that the cardiologist can quickly and accurately diagnose any latent disease by examining only the representative beats (the master key-beats) each from a cluster of homogeneous (similar) beats. This justifies the application of the proposed clustering approach, which aims to extract the optimal (number of) clusters within a diverse dataset. Recall that *optimality* here can only be assessed according to the CVI, the *feature extraction* (data representation) and the *distance* (similarity) *metric* used. Therefore, the clustering performance can further be improved by using superior alternatives than the basic and simple ones intentionally used in the current work with the sole purpose of

demonstrating the basic performance level of the proposed approach. For both passes the clustering validity assessment is performed using the same CVI given in Eq. (2) with $\alpha = 0$. Recall that this is entirely parameter-free and in addition, L_2 *Minkowski* norm (*Euclidean*) is used as the distance metric in the feature space.

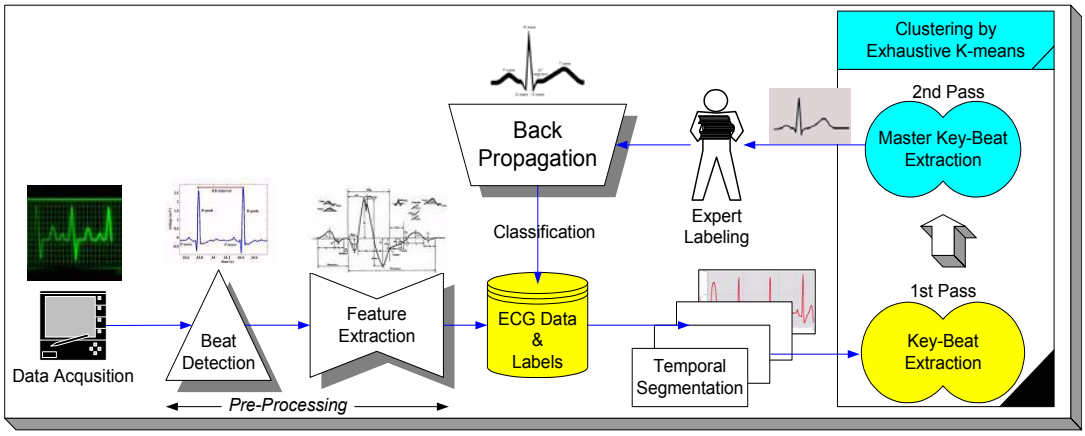

Figure 1. The overview of the proposed system for long-term ECG classification.

As shown in Figure 1, after the data acquisition is completed, the pre-processing stage basically contains beat detection and feature extraction of the sampled and quantized ECG signal. Before beat detection, all ECG signals are filtered to remove baseline wander, unwanted power-line interference and high-frequency noise from the original signal. This filtering unit can be utilized as part of the heartbeat detection process (for example, the detectors based on wavelet transforms [9]). For all records, we used the first-lead signals and the annotation information (provided with the MIT-BIH database [13]) to locate beats in ECG signals. Beat detection process is beyond the scope of this paper, as many beat detection algorithms achieving over 99% accuracy have been reported in the literature, e.g. [9] and [12]. Before feature extraction, the ECG signal is normalized to have a zero-mean and unit variance to eliminate the effect of dc offset and amplitude biases. Following the detection of each beat of the cardiac cycle within *quasiperiodic* ECG signals based on the R-peak detection and RR-intervals, morphological and temporal features are extracted as suggested in [1] and [6], and combined into a single characteristic feature vector to represent each heartbeat. The temporal features relating to heartbeat *fiducial* point intervals and morphology of the ECG signals are extracted by sampling the signals. They are calculated separately from the first-lead signals for each heartbeat. Since the detection of some arrhythmia (such as *Bradycardia, Tachycardia and premature ventricular contraction*) depends on the timing sequence of two or more ECG signal periods [16], four temporal interval features are considered in our study. They are extracted from heartbeat *fiducial* point intervals (RR-intervals), as follows:

1) *pre-RR-interval*: the RR-interval between a given heartbeat and the previous heartbeat,
2) *post-RR-interval*: the RR-interval between a given heartbeat and the following heartbeat,
3) *local average RR-interval*: the average of the ten RR-intervals surrounding a heartbeat,
4) *average-RR interval*: the mean of the RR-intervals for an ECG recording.

In addition to temporal features, ECG morphology features are extracted from two sampling windows in each heartbeat formation. The sampling windows are formed based on the heartbeat *fiducial* points (maximum of R-wave or minimum of S-wave). Specifically, the morphology of the QRS complex is extracted using a 150-ms window and 60-Hz sampling rate, resulting in nine ECG samples as features. The eight ECG samples representing the low-frequency T-wave morphology are extracted using a 350-ms window and 20-Hz sampling rate. The final feature vector for each heartbeat is then formed by combining 17 morphological and 4 temporal interval features.

Once the 21 dimensional (21-D) feature vectors composed from the temporal and morphological characteristics of ECG beats are extracted, the entire ECG data is temporally segmented into fixed size frames (segments) for achieving mainly two objectives. On one hand, the massive size of ECG data makes it almost infeasible to perform an efficient clustering and on the other hand, *outliers*, which are significantly different from the typical (normal) beats and thus may indicate the presence of an abnormal heart activity, may get lost due to their low frequency of occurrence. Therefore, we adopt a typical approach, which is frequently performed in audio processing, that is, temporally segmenting data into homogeneous frames. Due to the dynamic characteristics of an audio signal, the frame duration is typically chosen between 20ms and 50ms in order to get as a homogeneous signal as possible. Accordingly, for a 24 to 48 hours long *Holter* register, we have chosen ~5 minutes long (300 beats) duration for time segments since the intra-segment variation along the time axis is often quite low. So performing a

clustering operation within such homogeneous segments will yield only one or few clusters except perhaps the transition segments where a *change*, morphological or temporal, occurs on the normal form of the ECG signal. No matter how minor or insignificant duration this abnormal *change* might take, in such a limited time segment, the proposed exhaustive K-means clustering can separate those "different" beats from the normal ones and group them into a distinct cluster. One key-beat, which is the *closest* to the cluster centroid with respect to the distance metric used in 21-D feature space, is then chosen as the "prototype" to represent all beats in that cluster. Since the optimal number of clusters is extracted within each time segment, only necessary and sufficient number of key-beats is thus used to represent all 300 beats in a time segment. Note that the possibility of missing *outliers* is thus reduced significantly with this approach since one key-beat is equally selected either from an *outlier* or a typical cluster without considering their size. Yet redundancy among the key-beats of consecutive segments still exists since it is highly probable that similar key-beats shall occur among different segments. This is the main reason for having the 2nd pass, which performs the exhaustive K-means clustering over key-beats finally to extract the master key-beats. They are basically the "elite" prototypes representing all possible physiological heart activities occurring during a long-term ECG recording.

Since this is a personalized approach, each patient has, in general, normal beats with possibly one or few abnormal periods, indicating a potential heart disease or disorder. Therefore, ideally speaking only a few master key-beats would be expected at the end, each representing a cluster of similar beats from each type. For instance one cluster may contain *ventricular* beats arising from ventricular cavities in the heart and another may contain only *junctional* beats arising from atrioventricular junction of the heart. Yet due to the lack of discrimination power of the morphological or temporal features or the similarity (distance) metric used, the clustering operation may create more than one cluster for each anomaly. Furthermore, the normal beats have a broad range of morphological characteristics [15] and within a long time span of 24 hours or longer, it is obvious that the temporal characteristics of the normal beats may significantly vary too. Therefore, it is reasonable to represent normal beats with multiple clusters rather than only one. In short, several master key-beats may represent the same physiological type of heart activity. The presentation of the master key-beats to the expert cardiologist can be performed with any appropriate way as this is a *visualization* detail and hence beyond the scope of this work. Finally, the overall classification of the entire ECG data can be automatically accomplished by back propagating the master key-beats' labels in such a way that a beat closest to a particular master key-beat (using the same distance metric in 21-D feature space) is assigned to its label.

IV. Experimental Results

The systematic approach presented in Section 3 is applied to long-term ECG data in the Physionet MIT-BIH Long-Term database [13], which contains six two-channel ECG signals sampled at 128 Hz per channel with 12-bit resolution, and one three-channel ECG sampled at 128 Hz per channel with 10-bit resolution. The duration of the 7 recordings changes from 14 to 24 hours each and a total of 668,486 heartbeats in the whole database are used in this study. The database contains annotation for both timing information and beat class information manually reviewed by independent experts. The WFDB (Waveform Database) software package with library functions (from PhysioToolkit [14]) is used for reading digitized signals with annotations. In this study, for all records, we used the first lead signals and utilized the annotation to locate beats in ECG signals. The CVI, the feature extraction and the distance metric are already presented in Section 2.

Following the pre-processing that consists of the formation of heartbeats using the RR-intervals and the feature extraction thereafter, the patient's long-term ECG data is temporally segmented into homogenous frames of 300 beats (~5 minute duration) as described in Section 3. With 100 runs for each time frame, the exhaustive K-means clustering is then performed in 21-D feature space to extract the true number of clusters. We used $\alpha = 0$, to make CVI in Eq. (2) completely parameter-free, which then becomes the quantization error, Q_e and the range for K is set as $2 \le K \le 25$. The number of clusters, that is identical to the number of key-beats found automatically for each time frame depends on distinct physiological heartbeat types in each patient's ECG record. As a result, the proposed systematic approach by temporal segmentation and the dynamic clustering technique produces such key-beats that represent all possible physiological heart activities in a patient's ECG data. Therefore, finding the true number of clusters by the proposed systematic approach is the key factor that differentiates it from some earlier works such as [2] and [15], both of which iteratively determine this number by an empirical *threshold* parameter. In the proposed method no parameters or threshold values are used.

Table 1 shows the overall results of the proposed systematic approach over all patients from the MIT-BIH Long-Term ECG database. Labels manually annotated by the experts are used only for the master key-beats selected by the proposed system. The classification of the entire ECG data, or in other words, the labeling of all beats contained therein is then automatically accomplished by the back propagation of the master key-beat labels, as explained in Section 3. The performance results tabulated in Table 1 are calculated based on the differences between the labels generated by the proposed approach and the expert supplied

labels provided with the database. The *Association for the Advancement of Medical Instrumentation* (AAMI) provides standards and recommended practices for reporting performance results of automated arrhythmia detection algorithms [18]. In this study, according to the AAMI recommended practice, each ECG beat is classified into the following five heartbeat types: N (beats originating in the sinus mode), S (supraventricular ectopic beats), V (ventricular ectopic beats), F (fusion beats), and Q (unclassifiable beats). In the overall, the proposed systematic approach labeled heartbeats consistent with the cardiologist supplied annotations over 99% of the time within the entire benchmark dataset.

From the results in Table 1, the proposed systematic approach performed with very high accuracy for detection of normal (N) and ventricular (V) groups of beats. Specifically, accurate detection of premature ventricular contractions (PVCs) from the ventricular group (V) in long-term ECG data is essential for patients with heart disease since it may lead to possible life-threatening cardiac conditions [5]. On the other hand, for supraventricular ectopic (S) beats and some cases of fusion of ventricular (V) and fusion (F) beats, the proposed method did not form a separate cluster corresponding to each type of beat due to the fact that their morphological and temporal features are indeed quite similar to normal (N) beats. Therefore, we can conclude that a more accurate separation of both S and F beats from the N beats requires a superior feature extraction technique than the one used in the current work.

TABLE 1: OVERALL RESULTS FOR EACH PATIENT IN THE MIT-BIH LONG-TERM DATABASE USING THE PROPOSED SYSTEM. FOR EACH CLASS, THE NUMBER OF CORRECTLY DETECTED BEATS IS SHOWN RELATIVE TO THE TOTAL BEATS ORIGINALLY PRESENT.

Patient	N	S	V	F	Q	Accuracy
14046	105289/105405	1/1	9102/9765	73/95	0/0	99.31%
14134	38548/38766	3/29	9711/9835	744/994	0/0	98.75%
14149	144498/144534	0/0	243/264	0/0	0/0	99.96%
14157	82698/83412	104/244	4334/4368	57/63	0/0	98.99%
14172	57182/58315	401/1003	6526/6527	1/1	0/0	97.36%
14184	77606/78096	13/39	22479/23383	11/11	0/0	98.60%
15814	91129/91617	20/34	9706/9941	1601/1744	0/0	99.15%
Total	596950/600145	542/1350	62101/64083	2487/2908	0/0	**99.04%**

In addition to accuracy, which measures the overall system performance over all classes of beats, three other standard metrics found in the literature [6], sensitivity (*Sen*), specificity (*Spe*), and positive predictivity (*Ppr*), are used to quantify the performance of the proposed system with respect to detection of each class of beat. Their respective definitions using true positive (*TP*), true negative (*TN*), false positive (*FP*), and false negative (*FN*) are as follows: Sensitivity is the rate of correctly classified events among all events, $Sen = TP/(TP+FN)$; Specificity is the rate of correctly classified nonevents among all nonevents, $Spe = TN/(TN+FP)$; and Positive Predictivity is the rate of correctly classified events in all detected events, $Ppr = TP/(TP+FP)$. Accuracy is usually the most crucial metric for determining overall system performance, however due to large variation in the number of beats from different classes in the long-term ECG dataset, sensitivity, specificity, and positive predictivity can too be critical and relevant performance criteria for medical diagnosis. Table 2 presents performance results of the proposed system in these three areas for each class of beat. Overall, for normal (N) and ventricular (V) groups of beats the proposed system shows high performance, however its sensitivity and positive predictivity for supraventricular ectopic (S) beats are low as both morphology and temporal information features for S beats closely resemble F and N beats.

TABLE 2: PERFORMANCE OF THE PROPOSED SYSTEM FOR DETECTION OF EACH BEAT CLASS

	Sen	*Spe*	*Ppr*
N	99.47%	96.49%	99.60%
S	40.15%	99.78%	16.61%
V	96.89%	99.38%	99.46%
F	85.52%	99.49%	71.84%
Q	-	-	-

V. CONCLUSIONS

In this paper we proposed a long-term, personalized ECG classification system, which addresses the problem (presence of clinically relevant activity) within the entire life-time of a long-term ECG signal recorded from an individual patient and it is tested over a real (benchmark) database containing a total of 668,486 (manually) labeled heartbeats. To our knowledge this is the first work ever applied to a real full *Holter* database; since most of the earlier works tested only over regular half-hour excerpts from ambulatory ECG records with duration of 30 minutes

or even less, from the benchmark MIT-BIH arrhythmia database [10]. As a *personalized* approach with an expert labeling of only 5-15 minute long (clinically) distinctive ECG beats from each patient's long-term ECG recording, we achieved an average above 99% classification accuracy. In a typical 24-48 h long *Holter* register, selection of the right prototype beats, which can yield such a high accuracy level and a great reduction in effort, is mainly due to two key operations. The first one, the so called temporal segmentation, partitions the entire data into homogenous time segments that can be represented by *minimal* amount of key-beats. The following two-pass exhaustive K-means operations first extract the key-beats and then the master key-beats among them. In both operations, such delicate classification accuracy indicates that, the proposed approach successfully extracts the true (number of) clusters in a 21-D feature (data) space.

Moreover, such a systematic approach apparently promises a high level of *insensitivity* to the length (duration) of the data since the duration of the time segments is fixed and the number of clusters (master key-beats) found in the second pass is not related whatsoever with the number of key-beats in the first pass. Although the proposed system is intended and purposefully developed for analysis of long-term data sets by helping professionals focus on the most relevant patterns, it can also provide efficient and robust solutions for much shorter ECG data sets too. Besides classification, with some proper annotation, master key-beats can also be used for the *summarization* of any long-term ECG data for a fast and efficient visual inspection, and they can further be useful for indexing *Holter* databases, for a fast and accurate information retrieval. On the other hand, ~ 0.35% critical error rate, although may seem quite insignificant for a short ECG dataset, can still be high practically for *Holter* registers because it corresponds to several hundreds of misclassified beats, some of which might be important for a critical diagnosis. Yet recall that the optimality of the clustering algorithm depends on the CVI, the feature extraction method and the distance metric, in that, we purposefully use simple and typical ones so as to obtain a basic or unbiased performance level. Therefore, by using for instance a more efficient CVI and better alternatives for distance metric, as in [2] and [15], performance may be improved further. Instead of K-means, a better clustering method can also be used; however, note that the computational complexity may then become a serious drawback especially when used exhaustively as in the proposed approach. For better and more discriminative features, superior techniques can also be sought within computational biology and information theory, all of which are subject to our future work.

REFERENCES

[1] P. de Chazal, M. O'Dwyer, and R. B. Reilly, "Automatic classification of heartbeats using ECG morphology and heartbeat interval features," *IEEE Trans. Biomed. Eng.*, vol. 51, no. 7, pp. 1196–1206, Jul. 2004.

[2] D. Cuesta-Frau, J. C. Pe'rez-Corte's, and G. Andreu-Garci'a, "Clustering of electrocardiograph signals in computer-aided Holter analysis", *Computer Methods and Programs in Biomedicine*, vol. 72, no. 3, pp. 179-196, 2003.

[3] H. Frigui, R. Krishnapuram, "Clustering by competitive agglomeration", *Pattern Recognition*, vol. 30, pp. 1109-1119, 1997.

[4] N. J. Holter, "New methods for heart studies," *Science*, vol. 134, no. 3486, pp. 1214-1220, 1961.

[5] A. Iwasa, M. Hwa, A. Hassankhani, T. Liu, and S. M. Narayan, "Abnormal heart rate turbulence predicts the initiation of ventricular arrhythmias", Pace-Pacing and Clinical Electrophysiology. 28 (11), pp. 1189-1197, 2005.

[6] Y. Hu, S. Palreddy, and W. J. Tompkins, "A patient-adaptable ECG beat classifier using a mixture of experts approach," *IEEE Trans. Biomed. Eng.*, vol. 44, no. 9, pp. 891–900, Sep. 1997.

[7] A. K. Jain, M.N. Murthy and P.J. Flynn, "Data Clustering: A Review", *ACM Computing Reviews*, Nov 1999.

[8] M. Lagerholm, C. Peterson, G. Braccini, L. Edenbrandt, and L. Sörnmo, "Clustering ECG complexes using Hermite functions and self-organizing maps," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 7, pp. 838–848, Jul. 2000.

[9] C. Li, C. X. Zheng, and C. F. Tai, "Detection of ECG characteristic points using wavelet transforms," *IEEE Trans. Biomed. Eng.*, vol. 42, no. 1, pp. 21–28, Jan. 1995.

[10] R. Mark and G. Moody, "MIT-BIH Arrhythmia Database Directory," [Online]. Available: http://ecg.mit.edu/dbinfo.html

[11] M. Paoletti, C. Marchesi, "Discovering dangerous patterns in long-term ambulatory ecg recordings using a fast qrs detection algorithm and explorative data analysis," *Computer Methods and Programs in Biomedicine*, vol.82, pp. 20-30, 2006.

[12] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Trans. Biomed. Eng.*, vol. 32, No. 3, pp. 230–236, Mar. 1985.

[13] PhysioBank, "MIT-BIH Long-Term Database Directory," [Online]. Available: http://www.physionet.org/physiobank/database/ltdb/

[14] PhysioToolkit, "The WFDB Software Package," [Online]. Available: http://www.physionet.org/physiotools/wfdb.shtml

[15] Z. Syed, J. Guttag and C. Stultz, "Clustering and symbolic analysis of cardiovascular signals: discovery and visualization of medically relevant patterns in long-term data using limited prior knowledge", EURASIP Journal of Appl. Signal Process, 2007.

[16] W. J. Tompkins, J. G. Webster, Design of microcomputer-based medical instrumentation. Englewood Clifs, NJ 07632: Prentice Hall Inc. ISBN: 0-13-201244-8. pp. 398-3999, 1981.

[17] J.T. Tou and R.C. Gonzalez, *Pattern Recognition Principles*, London, Addison-Wesley, 1974.

[18] "Recommended practice for testing and reporting performance results of ventricular arrhythmia detection algorithms," Association for the Advancement of Medical Instrumentation, Arlington, VA, 1987.