Published in final edited form as:

Proceedings (IEEE Int Conf Bioinformatics Biomed). 2014 November ; 2014: 224–229. doi:10.1109/

BIBM.2014.6999169

Identifying Hypertrophic Cardiomyopathy Patients by Classifying Individual Heartbeats from 12-lead ECG Signals

Quazi Abidur Rahman¹, Larisa G. Tereshchenko^{2,3}, Matthew Kongkatong², Theodore Abraham², M. Roselle Abraham², and Hagit Shatkay^{1,4}

Quazi Abidur Rahman: quazi@cs.queensu.ca; Larisa G. Tereshchenko: tereshch@ohsu.edu; Matthew Kongkatong: matthew.kongkatong@gmail.com; Theodore Abraham: tabraha3@jhmi.edu; M. Roselle Abraham: mabraha3@jhmi.edu; Hagit Shatkay: shatkay@cis.udel.edu

¹Computational Biology and Machine Learning Lab, School of Computing, Queen's University, Kingston, ON, Canada

²Heart and Vascular Institute, Johns Hopkins University, Baltimore, MD, USA

³Knight Cardiovascular Institute, Oregon Health & Science University, OR, USA

⁴Dept. of Computer and Information Sciences & Center for Bioinformatics and Computational Biology, University of Delaware, Newark, DE, USA

Abstract

Test based on electrocardiograms (ECG) that record the heart electrical activity can help in early detection of patients with hypertrophic cardiomyopathy (HCM) where the heart muscle is partially thickened and blood flow is (potentially fatally) obstructed. This paper presents a cardiovascular-patient classifier we developed to identify HCM patients using standard 10-seconds, 12-lead ECG signals. Patients are classified as having HCM if the majority of the heartbeats are recognized as HCM. Thus, the classifier's underlying task is to recognize individual heartbeats segmented from 12-lead ECG signals as HCM beats, where heartbeats from non-HCM cardiovascular patients are used as controls. We extracted 504 morphological and temporal features - both commonly used and newly-developed ones - from ECG signals for heartbeat classification. To assess classification performance, we trained and tested a random forest classifier and a support vector machine classifier using 5-fold cross validation. The patient-classification precision and F-measure of both classifiers are close to 0.85. Recall (sensitivity) and specificity are approximately 0.90. We also conducted feature selection experiments by gradually removing the least informative features; the results show that a relatively small subset of 304 highly informative features can achieve performance measures comparable to that achieved by using the complete set of features.

Keywords

Hypertrophic Cardiomyopathy; Electrocardiogram; Patient classification; Random forests; Support vector machines; Feature selection

I. Introduction

Hypertrophic cardiomyopathy (HCM) is a genetic cardiovascular disease that may cause sudden cardiac death in young people [1]. In HCM patients, a portion of the heart muscle (myocardium) is thickened (hypertrophied). The most consistent trait of HCM is the

thickening of the muscle at the lower left chamber of the heart (left ventricle). An imaging method, two-dimensional echocardiography, is often used to identify left ventricular hypertrophy (LVH). However, this method is not reliable when the thickening of the muscle of the left ventricle is not clearly detectable. Moreover, early prediction of the disease for patients without any thickening of the muscle is not possible through echocardiography [2]. Therefore, the analysis of electrocardiogram (ECG) signals in patients with a family history of HCM and no clear muscle thickening in the left ventricle has high diagnostic value for early detection and prediction of HCM. We have also shown in a recent study [3] that the standard procedure of conducting ECG tests should not be ignored in mass pre-participation screening of young athletes.

Classifiers that automatically identify cardiovascular disease in patients may help reduce both cost and time of the pre-screening process. Historically, the main focus of ECG-classification research has been on identifying life-threatening arrhythmia in cardiovascular patients. Traditional machine learning methods such as artificial neural networks [4], support vector machines [5], random forests [6], and linear discriminants [7] have been used for reliable detection of life-threatening arrhythmia. Morphological (e.g., amplitude values) and temporal (e.g., length of various intervals) features have been primarily extracted for heartbeat-classification.

As mentioned earlier, left ventricular hypertrophy is the most consistent trait among HCM patients. Several criteria have been proposed to detect cardiovascular patients with left ventricular hypertrophy (LVH) from ECG signals. These criteria are usually derived from amplitude values of ECG waveforms. Many studies have been conducted to validate these LVH-detection criteria. While these studies have generally achieved high specificity (approximately 100%), sensitivity has been low (approximately 50%) across different studies [8]. Besides this criterion-based approach, multiple linear regression [11] and rule-based methods [12] have been used to detect cardiovascular patients with LVH, where morphological features were extracted from ECG signals. Potter et al. [9] have tested a set of criteria specifically proposed for detecting HCM [10] on a small group of 56 HCM patients and 56 healthy control subjects. The reported sensitivity and specificity from this study was approximately 90%. However, we are not aware of any previous work that employs machine learning methods for identifying HCM patients from ECG signals.

In this study, we aim to develop a classifier that can distinguish between ECG signals from HCM patients and those from non-HCM controls. Such a classifier will facilitate automated detection of HCM from ECG signals. However, we note that the classifier is not expected to replace extensive diagnosis by an expert cardiologist. Rather, it is intended as an initial screening method that will hopefully detect patients that may have HCM. The automatically detected patients will be referred for further extensive cardiovascular tests and be examined by expert cardiologists.

In order to develop a patient classifier for automated detection of HCM, we have segmented ECG signals into individual heartbeats, extracted features from each heartbeat and then classified these heartbeats by applying machine learning methods. We assigned a patient to the HCM class if the number of heartbeats classified as HCM is equal to or greater than the

number of heartbeats classified as control. For our classification experiments, we have extracted features that have been previously used, as well as some new morphological features from ECG signals. We have applied random forests and support vector machines classifiers to distinguish between heartbeats from HCM and those from non-HCM patients. Using 5-fold cross validation for training and testing, we achieve high performance levels as measured in terms of *precision*, *recall (sensitivity)*, *specificity* and *F-measure*. We also reduce, through feature selection, the number of features required to achieve the same performance level as that obtained by using the complete set of features.

The rest of the paper is organized as follows: Section II describes the ECG dataset used for classification experiments, obtained from HCM patients and from control subjects. In Section III, we discuss feature extraction, classification and feature selection methods and related tools. All classification results are presented in Section IV. Finally, we discuss and analyze the results and present directions for future work in Section V.

II. Data

The ECG dataset used in this study comprises standard 10-second, 12-lead ECG signals from two groups of cardiovascular patients. The first group consists of 221 hypertrophic cardiomyopathy (HCM) patients. Each HCM patient has one or more ECG recordings in the dataset. The total number of ECG signals in the HCM patients' dataset is 754. In the second group there are 541 subjects, all of which were diagnosed with ischemic or non-ischemic cardiomyopathy, and implantable cardioverter defibrillator (ICD) were inserted in their hearts for primary prevention of sudden cardiac death. As none of the ICD patients was diagnosed with HCM, their ECG data is used as the control in the experiments described here. While there may be cases in which a set of healthy controls would be preferable (e.g., pre-screening for HCM among young athletes), we have chosen the ICD patients' ECG dataset as the control because most of the patients referred for ECG tests in a hospital do not usually have a normal cardiac diagnosis; accordingly distinguishing HCM patients from other cardiovascular patients is a realistic, essential task. That said, we expect the methods used in this study to be applicable in other scenarios of distinguishing HCM patients from another group. Each patient in our control dataset has exactly one ECG recording, resulting in a total of 541 ECG signals the control set.

We segmented each ECG signal into individual heartbeats using the freely available ECGPUWAVE tool [13]. A heartbeat is a single cycle in which the heart's chambers relax and contract to pump blood, where each heartbeat comprises multiple waveforms. The ECG waves are created by the electrical signal that passes through the heart chambers (atria and ventricles). Fig. 1 shows a typical heartbeat and its waves: P, Q, R, S, T and U. It also shows inter-wave segments and intervals. While identifying each heartbeat, ECGPUWAVE detects the onset and offset points of the P-wave and the QRS-complex. It also identifies the offset point of the T-wave and the peak of the QRS-complex.

The segmentation of ECG signals was conducted on signals from each of the 12 leads. We then identified the heartbeats that are simultaneously detected on all 12-leads. Each of these

heartbeats was classified using machine learning methods as described in Section III.B. The summary of the dataset is presented in Table I.

III. Methods and Tools

After segmenting the 12-lead ECG signals into individual heartbeats, we extracted features from each heartbeat and represented it as a feature vector for classification. We also applied feature selection to identify highly informative features, and repeated the classification experiments using the selected features. We compared the results obtained from the different classification experiments and assessed the statistical significance of the observed differences. Finally, we identified HCM patients, by classifying each subject based on his/her respective number of heartbeats classified as HCM. The methods and tools used are discussed next.

A. Feature extraction

As described in Section II, we utilized the ECGPUWAVE tool to detect individual waveforms from heartbeats of HCM and ICD patients. We utilized the onset and offset points of various waveforms detected by the tool for extracting temporal and morphological features from each heartbeat. The peak of the QRS-complex was used to measure the RR-intervals between consecutive heartbeats. The morphological features extracted from the QRS complex and the temporal features have been used in the literature [7], [14] for heartbeat classification in a different context, namely, automatic detection of arrhythmia in cardiovascular patients. In the current study, we add morphological features of the P- and T-wave that have not been used before. The complete list of features is shown in Table II. To represent each heartbeat, we extract all 42 features from each of the 12 leads, resulting in a total of 504 features.

B. Classification of heartbeats and detection of HCM patients

As a first step to automatically detect HCM patients from 12-lead ECG signals, we developed a classifier whose task was to assign each instance (heartbeat) into one of two possible classes: HCM or control. Heartbeats from ICD patients are the controls in this study. We applied two standard classification methods: random forests [15] and support vector machine (SVM) [16]. We used the standard classification package in WEKA [17] for random forests and the libsvm package [18] in WEKA for SVM. The random forests algorithm was implemented with 500 trees; the number of features selected at random at each node of a tree was set to $2\sqrt{n}$, where n is the total number of features. Our SVM experiments used Gaussian radial basis function as the kernel.

In the classification experiment, we represented each heartbeat as a 504-dimensional vector of features where 42 features were extracted from each of 12-leads as described in Section III.A. We used the stratified 5-fold cross-validation procedure for training and testing. Although we are classifying individual heartbeats, the final goal of this study is to classify patients into two groups: HCM and ICD. Hence, we partitioned both HCM patients and ICD patients into 5 equal sized groups. Heartbeats from one group of HCM patients and from one group of control patients were included in the test set and the other four groups were used

for training. We repeated the process 5 times such that each heartbeat from an HCM patient or from a control subject is tested exactly once.

After classifying all heartbeats from a subject, we classified that subject as a HCM patient based on the number of heartbeats classified as HCM. If the number of the heartbeats classified as HCM is equal to or higher than that of heartbeats that have been classified as control, the subject is classified as a HCM patient.

To evaluate the performance of both the heartbeat and the patient classification, we have used several standard measures, namely, *precision*, *recall (sensitivity)*, and *specificity*. These measures are defined below, where true positives (*TP*) and true negatives (*TN*) are correctly classified HCM and control heartbeats (or patients), respectively; False positives (*FP*) denote control heartbeats (or patients) that are misclassified as HCM; HCM heartbeats (or patients) incorrectly classified as control are false negatives (*FN*);

$$\begin{aligned} \text{Precision} &= \frac{TP}{TP+FP}, \\ \text{Recall (Sensitivity)} &= \frac{TP}{TP+FN}, \\ \text{Specificity} &= \frac{TN}{TN+FP}. \end{aligned}$$

In addition to these three measures, we also calculate the *F-measure* which is the harmonic mean of precision and recall, defined as:

$$F - \text{measure} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}.$$

We compared the performance measures obtained by random forests and SVM, where the paired t-test [19] was used to assess the statistical significance of the differences along each performance measure.

C. Feature selection

We initially used all 504 features to classify heartbeats as HCM or control beats. Building classifiers from a large feature set can possibly lead to overfitting; moreover, including features that carry only negligible information about the heartbeat-class may incur unnecessary extra training time. To address these issues, we performed feature selection to reduce the number of features. To select features that have high predictive capability, we utilized the well-known *Information Gain* criterion [20] that measures how much information is gained about the heartbeat-class when the value of the feature is known. This criterion is primarily defined for nominal attributes. As all features in our study are numerical, they need to be discretized for calculating information gain. We calculated the information gain using the feature selection package in WEKA, which discretizes numerical attributes following Fayyad and Irani's algorithm [21].

After calculating information gains, we removed the 20 least-informative features and conducted the classification experiment again using both random forests and SVM. We

repeated this procedure by gradually removing 20 features at a time until we observed decline in performance.

IV. Results and Discussion

As explained in Section III.B, the first step in our experiment toward identifying HCM patients was to classify individual heartbeats such that each heartbeat is assigned to one of the two classes: HCM or control. We applied random forests and support vector machine classifiers using the complete set of 504 features for heartbeat classification. Table III shows the results from the 5-fold cross validation experiments. Precision (0.94) and F-measure (0.91) are the same for both classifiers. The small differences in recall and specificity are not statistically significant ($p > 0.35$).

To investigate how these performance measures change when the number of features is reduced, we calculated information gain for selecting highly predictive features. The highest information gain was 0.67 and the lowest was 0.001. Fig. 2 shows a histogram of the information gain distribution across features, where the x-axis shows the information gain values and the y-axis shows the number of features associated with each information gain. As values on the x-axis are rounded to 2 decimal points, an information gain of less than 0.01 is shown as zero (the leftmost column on the graph).

As described in Section III.C, we gradually removed the least-informative features, 20 at a time, and repeated the heartbeat classification experiment using both random forests and SVM. The change in performance in terms of all four measures using random forests for classification is shown in Fig. 3. All four performance measures marginally increase and decrease as we continue removing features until the number of features reaches 304. After that, the performance steadily declines as additional features are removed. All four measures for 304 features are exactly the same as those obtained when using the complete set of 504 features.

We have also plotted the performance measures for SVM while removing 20 features at a time, as shown in Fig. 4. The performance remains almost the same between 504 features and 384 features. Then it starts to decline steadily as we remove more features. The performance measures attained using 504 features and 384 features are exactly the same.

The next step in identifying HCM patients was to classify each subject as one of two classes: HCM or non-HCM. If the percentage of heartbeats classified as HCM was 50% or more, the subject was classified as an HCM patient. Table IV shows patient classification results where heartbeats were classified using all 504 features. Random forests and SVM perform almost the same and the marginal difference between performance measures is not statistically significant ($p > 0.85$).

As 304 features for the random forests classifier and 384 features for the SVM classifier performed exactly the same as the complete feature set when classifying individual heartbeats, we used the respective reduced feature sets to identify HCM patients based on the number of heartbeats classified as HCM. Patient classification results are presented in Table V, where heartbeats were represented using 304 features for random forests and 384

features for SVM. Paired t-tests show no performance-difference between SVM and random forests for classifying patients, when the reduced feature-sets are used for heartbeat classification ($p>0.38$).

The classification results described above show that we were able to achieve high performance level identifying HCM patients from 12-lead ECG data by classifying individual heartbeats using a set of 504 features. They also demonstrate that the reduced feature-sets, obtained by gradually removing the least informative features, performed equally well. The statistical tests applied show that the difference in performance obtained by random forests and by support vector machines is not statistically significant.

V. Conclusion

We have classified individual heartbeats from standard 10-second, 12-lead ECG signals to identify hypertrophic cardiomyopathy (HCM) patients. We have used ECG signals from HCM patients and from non-HCM controls to train and test heartbeat classifiers by applying random forests and support vector machines. A comprehensive set of 504 features extracted from ECG signals was used for heartbeat representation and classification. A subject was identified as a HCM patient if the majority of heartbeats was classified as HCM. The four performance measures from the patient classification experiment using random forests are: *precision 0.84, recall 0.89, specificity 0.93 and F-measure 0.86*. SVM performed similarly, as confirmed by the paired t-test. We have also used the information-gain criterion for feature selection. For random forests, performance measures using 304 selected features were similar to the measures obtained using the complete set of 504 features. For SVM, this was true for a set of 384 informative features.

This work is the first study of its kind, setting out to automatically identify HCM patients from 12-lead ECG signals by classifying heartbeats using machine-learning methods. We have shown that it is possible to attain high performance using random forests or SVMs. We also showed that the information-gain criterion can be effectively used to choose a reduced set of temporal and morphological features that retain a similar level of performance. While in this study we have classified patients simply based on the percentage of individual heartbeats classified as HCM, in future research we shall focus on analyzing and modeling the sequence of heartbeats using advanced machine learning methods.

References

1. Maron, BJ.; Salberg, L. Hypertrophic Cardiomyopathy: For Patients, Their Families and Interested Physicians. Wiley-Blackwell; 2006.
2. Maron BJ. The electrocardiogram as a diagnostic tool for Hypertrophic Cardiomyopathy : Revisited. Ann Noninvasive Electrocardiol. 2001; 6(4):277–279. [PubMed: 11686906]
3. Rahman, QA.; Kanagalingam, S.; Pinheiro, A.; Abraham, T.; Shatkay, H. What We Found on Our Way to Building a Classifier : A Critical Analysis of the AHA Screening Questionnaire. In: Imamura, K.; Usui, S.; Shirao, T.; Kasamatsu, T.; Schwabe, L.; Zhong, N., editors. BHI 2013, LNAI. Vol. 8211. Heidelberg: Springer; 2013. p. 225-236.
4. Yu S, Chou K. Integration of independent component analysis and neural networks for ECG beat classification. Expert Syst Appl. May; 2008 34(4):2841–2846.

5. Melgani F, Bazi Y. Classification of electrocardiogram signals with support vector machines and particle swarm optimization. *IEEE Trans Inf Technol Biomed. Sep; 2008 12(5):667–77.* [PubMed: 18779082]
6. Emanet, N. ECG beat classification by using discrete wavelet transform and Random Forest algorithm. *Fifth International Conference on Soft Computing, Computing with Words and Perceptions in System Analysis, Decision and Control; 2009.*
7. De Chazal P, Reilly RB. patient-adapting heartbeat classifier using ECG morphology and heartbeat interval features. *IEEE Trans Biomed Eng. 2006; 53(12):2535–2543.* [PubMed: 17153211]
8. Schillaci G, Battista F, Pucci G. A review of the role of electrocardiography in the diagnosis of left ventricular hypertrophy in hypertension. *J Electrocardiol. 2012; 45(6):617–23.* [PubMed: 23022303]
9. Potter SLP, Holmqvist F, Platonov PG, Steding K, Arheden H, Pahlm O, Starc V, McKenna WJ, Schlegel TT. Detection of hypertrophic cardiomyopathy is improved when using advanced rather than strictly conventional 12-lead electrocardiogram. *J Electrocardiol. Jan; 2010 43(6):713–8.* [PubMed: 21040828]
10. Corrado D, McKenna WJ. Appropriate interpretation of the athlete's electrocardiogram saves lives as well as money. *Eur Heart J. Aug; 2007 28(16):1920–2.* [PubMed: 17623679]
11. Warner RA, Ariel Y, Gasperina MD, Okin PM. Improved electrocardiographic detection of left ventricular hypertrophy. *J Electrocardiol. Jan.2002 35 Suppl:111–5.* [PubMed: 12539107]
12. Kaiser W, Faber TS, Findeis M. Automatic learning of rules. A practical example of using artificial intelligence to improve computer-based detection of myocardial infarction and left ventricular hypertrophy in the 12-lead ECG. *J Electrocardiol. Jan.1996 29 Suppl:17–20.* [PubMed: 9238372]
13. Laguna, P.; Jane, R.; Bogatell, E.; Anglada, D. ECGPUWAVE. [Online]. Available: <http://www.physionet.org/physiotools/ecgpuwave/>
14. 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(7):1196–1206.* [PubMed: 15248536]
15. Breiman L. Random forests. *Mach Learn. 2001; 45(1):5–32.*
16. Cortes C, Vapnik V. Support-vector networks. *Mach Learn. 1995; 297:273–297.*
17. Hall M, Frank E, Holmes G, Pfahringer B, Reutemann P, Witten IH. The WEKA data mining software. *ACM SIGKDD Explor Newsl. Nov.2009 11(1):10.*
18. EL-Manzalawy, Y.; Honavar, V. WLSVM: Integrating LibSVM into Weka Environment. 2005. [Online]. Available: <http://www.cs.iastate.edu/~yasser/wlsvm.html>
19. Dietterich TG. Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms. *Neural Comput. Oct; 1998 10(7):1895–1923.* [PubMed: 9744903]
20. Mitchell, TM. *Machine Learning.* McGraw-Hill; 1997.
21. Fayyad, U.; Irani, K. Multi-interval discretization of continuous-valued attributes for classification learning. *Thirteenth International Joint Conference on Artificial Intelligence; 1993. p. 1022-1027.*
22. ECG wave. [Online]. Available: <http://lifeinthefastlane.com/ecg-library/basics/t-wave/>

Fig. 1.
A typical heartbeat comprising P, Q, R, S, T, U waveforms and inter-wave segments and intervals [22].

Fig. 2.
Histogram of the information gain distribution across 504 features.

Fig. 3. Performance measures from heartbeat classification using random forests while gradually removing 20 features at a time.

Fig. 4. Performance measures from heartbeat classification using SVM while gradually removing 20 features at a time.

Table I

Summary of the ECG dataset used in this study. Each HCM patient has one or more ECG signals, whereas each of the controls has only one signal in the dataset.

Type of patient	Number of patients	Total number of ECG recordings	Total number of Heartbeats
HCM	221	754	6488
ICD (Control)	541	541	4442

Table II

Complete list of the 42 features extracted from EACH OF THE 12-lead ECG signals for classifying heartbeats. (the total number of features is $42 \times 12 = 504$)

Group	Feature	Definition	Number of features
Temporal (length of intervals)	Pre-RR interval	The interval between the current heartbeat and the previous heartbeat	6
	Post-RR interval	The interval between the current heartbeat and the following heartbeat	
	Average RR-interval	The mean of the RR intervals of a recording and the it is used as the same for all the heartbeats in a recording	
	P-wave duration	The interval between the P-wave onset and offset	
	QRS interval	The interval between the QRS onset and offset	
	T-wave duration	The interval between QRS-offset and T-wave offset	
Morphological (amplitude values)	QRS morphology	10 uniformly sampled amplitude values between the QRS onset and the QRS offset.	36
		Maximum and minimum of original sampled amplitude values in the QRS complex.	
	P and T wave morphology	10 uniformly sampled amplitude values between the wave onset and the wave offset.	
		The maximum and the minimum of the original sampled amplitude values in the P and T wave.	

Table III

Heartbeat classification results using all 504 features.

Classifier	Precision	Recall (Sensitivity)	Specificity	F-measure
RF (all features)	0.94	0.87	0.92	0.91
SVM (all features)	0.94	0.88	0.91	0.91

Table IV

Patient classification results where all 504 features are used for heartbeat classification.

Classifier	Precision	Recall (Sensitivity)	Specificity	F-measure
RF (all features)	0.84	0.89	0.93	0.86
SVM (all features)	0.83	0.90	0.92	0.87

Table V

Patient classification results where reduced sets of 304 (RF) and 384 (SVM) features are used for heartbeat classification.

Classifier	Precision	Recall (Sensitivity)	Specificity	F-measure
RF (304 features)	0.85	0.89	0.93	0.87
SVM (384 features)	0.82	0.89	0.92	0.85