State of the Art

New Developments in the Automatic Analysis of the Surface ECG: The Case of Atrial Fibrillation

Franco Chiarugi

Biomedical Informatics Laboratory, Institute of Computer Science, Foundation for Research and Technology – Hellas, Heraklion, Crete, Greece

Key words:
Digital electrocardiography, classification of atrial fibrillation.

Manuscript received:
June 2, 2008;
Accepted:
June 17, 2008.

Address:
Franco Chiarugi

Biomedical Informatics Laboratory
ICS-FORTH
N. Plastira 100
Vassilika Vouton
70013 Heraklion,
Crete, Greece
e-mail:
chiarugi@ics.forth.gr

This article discusses current approaches to the analysis of digital electrocardiographic data, with special reference to the detection and classification of episodes of atrial fibrillation (AF). AF is the most common cardiac arrhythmia encountered in clinical practice. It is a supraventricular tachyarrhythmia characterised by uncoordinated atrial activation, with resulting deterioration of atrial mechanical functionality.[1,2] AF accounts for approximately one third of hospitalisations for cardiac rhythm disturbance, affecting 4.5 million people in the EU and up to 2.3 million people in the USA.[3] However, the public health burden of AF could increase dramatically over the next 50 years[3] in parallel with the ageing population, since the prevalence and the incidence of AF are known to increase with age.[4]

The guidelines drawn up by the American College of Cardiology, American Heart Association, and the European Society of Cardiology[1,2] include a classification scheme for AF that is unrelated to an immediate, reversible cause (which would mean that the AF is not likely to recur or to have prognostic significance after the disappearance of the acute aetiology). The scheme specifies four categories: first detected episode (provides no knowledge of previous events and uncertainty about duration); paroxysmal (recurrent, but can terminate spontaneously in less than 7 days, and commonly less than 48 hours); persistent (recurrent, lasting more than 7 days, and typically requires cardioversion to terminate); and permanent (fails to terminate after cardioversion, or relapses less than 24 hours after termination, or no cardioversion attempted). Note that when a patient has had 2 or more episodes after the clinical detection of the first episode of AF, the arrhythmia is considered to be recurrent. Thus, the guidelines focus on the classification and treatment of the arrhythmia based on its duration (paroxysmal, persistent and permanent) rather than on its aetiology.[5]

AF is associated with substantial mortality and morbidity due to thromboembolism, heart failure, and impaired cognitive function. In particular, persistent or permanent AF entails a serious risk of thromboembolism, as a result of thrombus formation within the atria, which can cause stroke or other thromboembolic events. Paroxysmal AF accounts for 35 to 66% of total cases[6,7] and may deteriorate into persistent or permanent AF. Considering that the ideal therapeutic goal for AF is to achieve and maintain sinus rhythm,[8] the use of aggressive procedural techniques should be avoided in the case of paroxysmal episodes. Thus, the discrimination between paroxysmal and persistent or permanent AF, and the prediction of paroxysmal AF termination, can be invaluable in order to avoid useless therapeutic interventions, reduce the associated clinical costs, minimise the risks for the patient and improve the patient's quality of life.

Digital electrocardiography

Digital electrocardiography offers new possibilities for the interpretation of the ECG. Advanced computerised analysis, using sophisticated algorithms for the separation and classification of atrial and ventricular rhythms, allows the individual components of the ECG signal to be evaluated in isolation.

The basic requirement for any accurate signal processing of ECGs is the availability of digital ECGs. Since 1957 electrocardiography has inevitably been influenced by the development of modern computers. Indeed, computer interpretation of the ECG was one of the first applications of computers in health care. The first attempt to automate ECG analysis using a digital computer was made as early as 1957[9] and other attempts soon followed.[10] In the late seventies and eighties computer processing of ECGs increased rapidly[11] and shifted from university to industry and to commercial ECG devices. By 1990 more than 100 million ECGs per annum were being interpreted by a computer program.[12] In parallel with the evolution of computer technology, ECG devices evolved from analogue recorders to devices with the ability to store ECGs locally in digital format. Twenty years ago, transmission of ECGs was mainly performed using analogue techniques, but during the last fifteen years almost all newly developed ECG devices with medium or high capabilities offer features for digital recording and interpretation and the possibility to transfer a digitally stored ECG to a computer system.

Digital ECGs are susceptible to analysis by a variety of computerised algorithms. Although many cardiologists may still question the value of computerised ECG interpretation, it is widely recognised that computer algorithms are capable of performing many other analytical tasks—such as analysis of heart rate variability, late potentials, and QT dispersion—which are beyond the capacity of a human operator.[13]

One relevant tool for the understanding of the pathophysiological mechanisms of AF is the analysis and interpretation of digital atrial electrograms (AEG), which are recordings obtained (invasively) from the atrial surface.[14] To study AF signals properly, similar techniques can be applied to the AEG and the ECG, taking into account that atrial activity is more evident in the former than in the latter; on the other hand, multi-channel AEG recordings are rarely available, in contrast to ECGs. If satisfactory results can be obtained from the surface ECG, then this approach is surely preferable, considering that the AEG has the disadvantage (from the patient's point of view) of being an invasive examination.

Suitable surface ECGs for the analysis of atrial activity can be obtained in digital form from resting ECG examinations (short-term ECGs, typically 10 s) or Holter ECG recordings (long-term ECGs, typically 24 hours). The difficulty of catching the arrhythmia during its manifestation favours the use of excerpts selected from Holter ECG recordings that include AF episodes. However, short-term ECGs could also be used provided they are recorded during an AF episode.

Detection of atrial fibrillation in the surface ECG

In AF the normal atrial depolarisation represented by the P-wave on the sinus rhythm ECG degenerates into a rapid, chaotic pattern of fibrillation waves (f-waves) that vary in size, shape and timing.[2] The QRS complex in AF often has no morphological changes, but the ventricular response is usually irregular,[15] which results in RR intervals with highly variable duration.

AF may occur in isolation or in association with other arrhythmias, most commonly atrial flutter (AFL) or atrial tachycardia (AT). AFL is a more organised arrhythmia than AF and is characterised by a saw-tooth pattern of regular atrial activation called flutter waves (F-waves) on the ECG.[1,2] The surface ECG thus contains all the relevant parameters necessary for the identification of an AF episode. However, although it is easy for a trained cardiologist to recognise AF on the ECG, the classification of an AF episode (as paroxysmal, persistent, or permanent) based on the surface ECG adds a new level of difficulty.

Perhaps this is one of the areas where the computer can provide significant assistance. It would be interesting to verify if the relevant information for the classification of an AF episode is present in the surface ECG. But what are the clinical parameters (if indeed they exist) that can be evaluated from the surface ECG and used for AF classification and prognosis?

It has not been completely explained how an episode of AF is maintained or is terminated. The duration of paroxysmal AF episodes changes with the patient and the episode, and the moment in time when paroxysmal AF self-terminates is not known.[16] The most widely accepted theory to explain AF is based on the continuous propagation of multiple wavelets (called re-entries) wandering throughout the atria and dependent on the refractory period, mass and conduction velocity.[17] Several studies have demonstrated a reduction in the number of re-entries just before AF termination, with the production of simpler wave fronts where f-waves evolve to P-waves. According to these studies, the atrial activity slightly evolves to a more organised pattern prior to AF termination.[18,19]

From the signal processing point of view the need to separate atrial from ventricular activity in the AEG or ECG seems evident. The separation of atrial and ventricular components is hampered by the fact that the two signals are often superimposed in each AEG/ECG channel. Consequently, the design of algorithms for the cancellation of the ventricular activity is not an easy task and a perfect cancellation cannot be performed, especially in the case of signals with extrasystoles, artefacts, baseline wandering and high frequency noise. Once the atrial activity is separated from the ventricular, several features can be extracted from the two signals using time and/or frequency domain analysis in order to find suitable indexes for the classification of AF episodes. In the analysis that follows, the atrial component consists mainly of an oscillatory wave whose characterisation will be made in the frequency domain. Besides spectral parameters of the atrial component, measures extracted from the ventricular activity, such as the heart rate, will be considered.

Techniques for separating atrial and ventricular rhythm

The fibrillation process might be better analysed on samples from intervals without ventricular activity, i.e. the T-Q intervals. However, such a straightforward approach does not fully take advantage of the features of the fibrillation process; samples during the QRST interval can equally be used for analysing the atrial activity, provided that the ventricular activity has first been properly removed. The availability of additional samples reflecting atrial activity helps not only to improve the accuracy of the atrial feature evaluation but also to avoid the problem of vanishing T-Q intervals at high heart rates.[20]

QRST cancellation through cross-channel adaptive filtering is a possible approach. Using this technique, the lead with the highest atrial component power (which can be identified using various techniques) is used as the desired input of an adaptive filter using a recursive least squares (RLS) algorithm. The filter has as its reference input the signal with the lowest atrial activity. The high QRST power (compared to the f-wave power) drives the adaptation of the filter coefficients, providing at the error output a signal where the QRST is cancelled and at the other output a signal where the ventricular activity is enhanced.

Average beat subtraction is the most widespread technique for atrial activity extraction and is based on the fact that AF is uncoupled from the ventricular activity. This method is based on a classification of the detected QRSTs and the subtraction from the original signal of the average beat (QRST) of the class to which each beat belongs. Usually, the classification is performed using template matching techniques that are applied to the entire QRST complex, or to the QRS complex and the T-wave separately. In this respect, the use of an adaptive template in conjunction with the correct spatiotemporal alignment of every QRS complex has proven to perform very well, even though it requires the availability of multi-channel recordings.[20-22]

Another recently proposed approach to atrial signal extraction exploits the property that atrial and ventricular activities arise from different bioelectrical sources, applying the so-called blind source separation algorithm.[23] This is a technique for the estimation of original sources from the observation of their merged signals only, without any *a priori* information about the sources—hence the term blind. These different sources can be separated when multi-lead ECG recordings are available, by taking advantage of their different statistical properties. Blind source separation algorithms are usually based on principal component analysis (PCA) or independent component analysis (ICA); these are the two main separation methods and both have been also applied to AF analysis.[20]

Annotated databases

The availability of public annotated databases is of paramount importance to the scientific community in order to allow an easy test of new algorithms and an easy comparison with the outcomes of similar studies or different approaches tested on the same database.

The PhysioNet/Computers in Cardiology 2004 challenge was focused on this question: "Is it possible to predict if (or when) an episode of atrial fibrillation will end spontaneously?" In order to answer these questions the participants in this challenge had to develop computerised algorithms trained and evaluated on a public annotated database provided on the PhysioNet website,[24,25] composed of annotated learning sets and two different test sets.

The study of AF analysis in our laboratory was started with a view to participating in the PhysioNet/

Computers in Cardiology 2004 challenge. In the following sections the approach we implemented, which won 3 out of 4 prizes in the challenge, is described in detail, together with some further subsequent developments made in some parts of the algorithm.

The ECG records of the AF Termination Challenge Database,[25] made available on the PhysioNet site[24] as learning and testing sets, were used for the classification of the AF episodes. This database contains real data (surface ECGs) for a total of 80 records. Each record, extracted from a two-lead, 128 Hz Holter ECG recording, is 1 minute in length. Learning set N is composed of 10 records of non-terminating AF (N-type AF), defined as AF that was not observed to have terminated for the duration of the long-term recording, at least an hour following the record. Learning sets S and T are each composed of 10 records of paroxysmal AF, but in learning set S AF terminated one minute after the end of the record (S-type AF), while in learning set T AF terminated immediately after the end of the record (T-type AF).

For the discrimination between terminating and non-terminating AF test set A was made available, composed of 30 records, of which about one-half contain N-type AF and the remainder T-type AF. In addition, test set B was provided, composed of 20 records, 10 with S-type AF and 10 with T-type AF.

In addition to the above, the records of the MIT-BIH Arrhythmia Database[26] were used for further significant improvements in the QRS detection and classification algorithms that were utilised in the approach we adopted for the separation of ventricular and atrial activities. The use of a different database annotated in terms of QRSs and not in terms of AF episodes was necessary to allow a reliable evaluation of the performance of such algorithms.

The MIT-BIH Arrhythmia Database is composed of real data (surface ECGs) for a total of 48 records. The records are about half-hour excerpts of two-channel ambulatory ECG recordings obtained from 47 subjects studied by the BIH Arrhythmia Laboratory between 1975 and 1979. The recordings are digitised at 360 Hz with 11-bit resolution over a 10 mV range. Each beat of each record has a reference annotation identifying the QRS position and type.

Ventricular rhythm (QRS detection and related techniques)

QRS detection is the first fundamental step in the evaluation of the ventricular rhythm. QRS detection can be performed[27] in the time domain (derivative-based algorithms or algorithms based on digital filters), in a transformed domain (wavelet-based QRS detection or Hilbert transform-based QRS detection) or using other techniques (e.g. neural networks).

The approach we selected was in the time domain with derivative-based algorithms.[28] In detail, the process consists of the following steps: filtering (0.2-40 Hz) of both ECG signals and interpolation to 1024 Hz; QRS detection by an adaptive threshold on the QRS-enhanced signal (QeS) obtained by summing of absolute filtered derivatives of each channel. The derivative filter was selected as a band-pass filter with a derivative behaviour in the band of interest.

The adaptive threshold was initially evaluated as 40% of the average among the QRS-enhanced signal peaks in 2 s periods, rejecting cases that fell outside the 98% percentiles. At each detection, the threshold was re-evaluated taking into account the value of the current detected peak. In addition, a "closed-eye" period of 200 ms after the detected QRS was set up in order to avoid false detections unreasonably close to the previous QRS. Also, after each detection, the threshold was artificially increased (150% of its regular value) and reduced with a linear trend, again in order to avoid detection of high T-waves as QRS.

The QRS detector was only qualitatively evaluated on the AF Termination Challenge Database files, mainly because the database was not annotated for QRS detection but only for AF episode classification and the main focus was the classification of the AF episode. However, the results were considered satisfactory and a quantitative evaluation was postponed and performed on a more suitable database (MIT-BIH Arrhythmia Database with QRS and rhythm annotations). During the quantitative evaluation of the QRS detection algorithm some improvements were made and the final algorithm achieved an excellent performance on the MIT-BIH Arrhythmia Database.

The main modifications were the insertion of stricter band-pass filtering (5-15 Hz) in the pre-filtering phase of the signal, in order to select more accurately the typical frequencies contained in the QRS complexes.[29] The algorithm for the threshold re-evaluation after each QRS detection was reviewed. In the modified version the average QeS peak (QeSap) is continuously updated after each QRS detection using the QeS peak (QeSp) detected in the current QRS with the following formula (n is the progressive number of the detected QRS):

$$QeSp(n) = \min(QeSp(n), 1.5*QeSap(n-1))$$
$$QeSap(n) = 0.97*QeSap(n-1) + 0.03*QeSp(n)$$

The beginning of a QRS is detected when the QeS exceeds the threshold (0.4*QeSap(n-1)), while the end of the QRS is revealed when the QeS drops to the threshold and remains down for a sufficient number of consecutive samples.

The original signals (ch. 1 and ch. 2), QeS and the detection threshold in an excerpt from a record of the MIT-BIH Arrhythmia Database are shown in Figure 1.

A final improvement was made by introducing an estimation of the noise around each detected QRS in each channel. In fact, when the noise in one or both leads is considerably high, the performance of the detector is significantly reduced, but it was observed that, when the noise was present in only one channel, the exclusion of this noisy channel from the QeS improved the QRS detection results. Of course this technique can be applied only in the case of multi-channel ECGs.

The noise level in each lead for each detected QRS was estimated using the signal in the following T-P interval. Similarly, the QRS power and the ratio between the T-P interval average power and the QRS average power were evaluated. If this ratio was below 0.1, the complex was considered not to be noisy; if it was between 0.1 and 0.2, it was considered to be low in noise, and, if it exceeded 0.2, then it was considered noisy. Thus, based on such ratio, a three-level noise index (NI) was obtained and for any interval of an ECG, a noise score (NS) was estimated by averaging the NI of the QRSs detected in that interval. Intervals containing noise (noisy intervals) were then evaluated. The appearance of a number of consecutive noisy QRSs determines the beginning of a noisy interval, which ends once a few consecutive non-noisy QRSs appear. For each noisy interval, the QRS detection was also performed using only channel 1 and only channel 2, re-evaluating the NS in both cases. Finally a criterion for the selection of the best channel (1, 2 or 1+2) for the QRS detection was implemented.

The performance was calculated based on the confusion matrix, from which the true negatives (TN, correct evaluation of the absence of a QRS) cannot be derived, but only the true positive (TP), false positive (FP) and false negative (FN) are available. Thus, the available indexes were only the sensitivity (TP/ (TP+FN)) and the positive predictive value (PPV= TP/(TP+FP)). The performance of the implemented QRS detector was very high (sensitivity 99.76%, PPV 99.81%) and compared favourably to other published studies evaluated using the same database. The algorithm proposed by Hamilton[30] achieved a sensitivity of 99.76% and PPV 99.80%, which are substantially the same as the algorithm described here, but the reported statistical results were obtained from a subset of the total an-

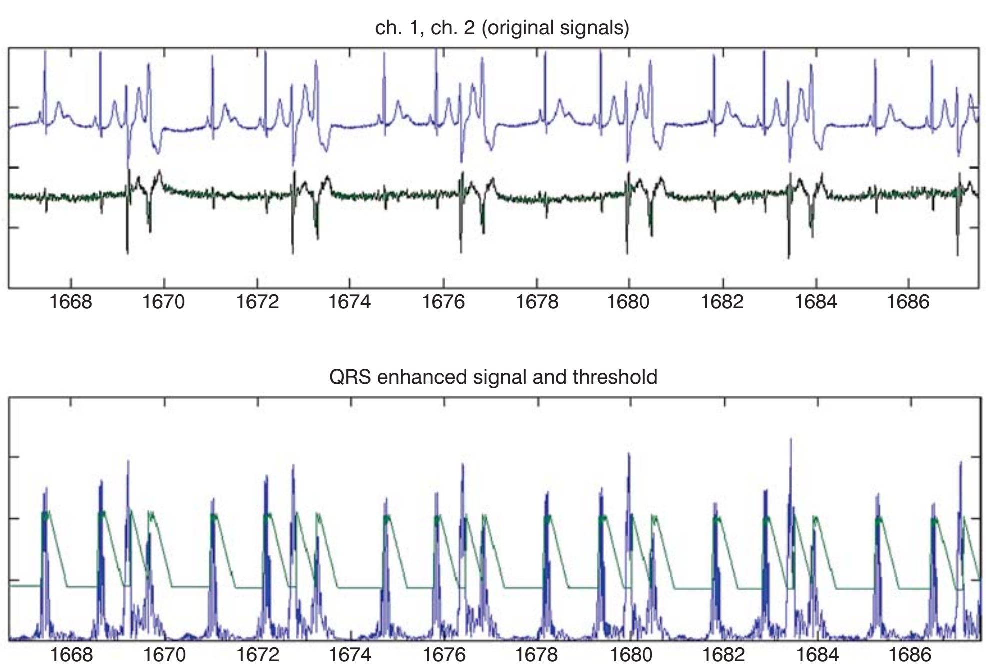

Figure 1. Original signals (ch. 1 and ch. 2), QeS and detection threshold in an excerpt from a record of the MIT-BIH Arrhythmia Database in the presence of multiform premature ventricular complexes (on the abscissa, time is expressed in seconds).

notated beats (about 91,000 annotated QRSs). Similarly, the best results obtained by Christov[31] had sensitivity 99.74% and PPV 99.65%, a little lower than the ones obtained with the algorithm described here.

The sequence of the detected QRSs (in the absence of false positives and false negatives) identifies the ventricular rhythm. From the ventricular rhythm the following parameters were considered for testing as regards their significance in the classification of AF episodes: RRmean (average of the RR intervals over the entire record), HRmean (average of the heart rate over the entire record) and RRstd (standard deviation of the RR intervals over the entire record).

Atrial rhythm (P- or f-/F-wave detection and related techniques)

In the analysis of an AF episode, the characterisation of atrial activity may best be carried out using spectral analysis, from which the significant parameters can be derived. In fact, detection of single f-/F-waves can be of interest in the case of atrial rhythm classification and in the analysis of transitions from sinus rhythm to AF episode and *vice versa*,[32] but that was far beyond the goal of the PhysioNet Challenge.

In a more generic approach to the full analysis of the atrial rhythm for a generic surface ECG, the detection of P-waves is a crucial task in detecting and discriminating between supraventricular and ventricular arrhythmias.[33] In contrast to the ventricular activity, the detection of atrial activity is a complex and unsolved problem, especially in arrhythmic situations. P-waves are difficult to detect because of their low signal-to-noise ratio, their high inter-patient variability and their possible invisibility on the ECG signal in the case of atrioventricular (AV) dissociations, where QRS, T- and P-wave occurrences do not follow their normal sequence and the P-wave may be hidden by a higher wave (QRS or T).[34]

The lack of annotated beat-to-beat P-wave databases has hampered a quantitative evaluation of several techniques found in the literature. From an analysis of some significant methods in the recent literature, they are mainly based on these approaches:[35] a) localised search area methods; b) source separation methods; c) P-wave detection exploiting the knowledge of the current PP rhythm.

The localised search area methods consist in searching for the P-wave in a localised area outside the QRST (generally before the QRS). Once the area is identified, the P-wave can be detected directly using different signal processing techniques such as template matching,[32] digital filtering,[36,37] wavelets,[38,39] wavelet transform and neural network,[40] or a combination of hidden Markov model and wavelet.[41] The localised search area approach gives satisfactory results for ECGs where the QRS is easy to identify and the atrial activity is synchronised with the ventricular (e.g. sinus rhythm). However, relying only on the QRS locations for the estimation of the P-waves, it cannot properly deal with situations of AV dissociation.[35]

The source separation approaches involve viewing P-wave detection as a source separation problem. They first separate the atrial activity from the ventricular activity, by either QRST cancellation or direct separation. This produces an atrial activity signal where the QRST complexes (noise) have been removed and consequently the P-waves (signal) have an increased signal-to-noise ratio and can be more easily detected. These approaches are expected to perform better in the presence of AV dissociations. The most common approach is the subtraction of a QRST template from the original signal.[34,42] Other methods have been proposed using neural networks and blind source separation. After cancellation, atrial activity may be analysed by a classification phase[33,34] or by spectral analysis.[43] Although these approaches separate the ventricular activity from the atrial activity, they can remain sensitive to the presence of abnormal beats that with simple techniques are not properly cancelled. Moreover, the QRS occurrence is not employed in detecting the P-wave, but it could be utilised in situations other than AV dissociations in order to support P-wave activity detection.[35]

P-wave detection exploiting the knowledge of the current PP rhythm is an integrated approach, which uses localised search areas driven by the current PP rhythm estimation.[35] It has shown a notable performance, not only in sinus rhythm but also in different types of arrhythmias.

Analysis of atrial rhythm in AF

In the context of the PhysioNet Challenge, the problem was not related to the detection of AF episodes but to the analysis of a known AF episode; there was no need to search for P-waves or to detect each single f-/F-wave. In fact, the atrial rhythm did not need to be identified: it was already known we were in the presence of an AF episode and we simply had to identify and extract some significant features from the atrial signal.

As a first approach, the cross-channel adaptive filtering technique for the separation of ventricular and atrial activities was implemented using a 2nd degree polynomial adaptive filter with forgetting factor lambda 0.996. This technique performed quite well except in cases of ECG signals with extrasystoles or artefacts, where the QRST cancellation left some residuals in the QRST areas that might alter the subsequent signal processing.[28]

In an attempt to obtain a more accurate cancellation, a different approach was selected and two main considerations indicated the choice of an average beat subtraction (ABS) approach: a) in operative conditions (especially using devices for ECG monitoring outside the healthcare premises) it is very common to acquire ECGs with a very limited number of ECG channels; b) a precondition for the required solution was the possibility of easy down-scaling to single channel ECGs (typical of wearable devices).

The selected ABS approach was based on QRST cancellation operated by QRS detection, QRST clustering with class centroid estimation, and subtraction of the estimated class centroids from the signal after establishing a proper beat alignment. An example of the atrial activity obtained using the above processing is shown in Figure 2.

In this way, the drawback of the cross-channel adaptive filter cancellation introducing artefacts on abnormal beat occurrences was overcome. However, baseline movements, high QT interval variability strictly correlated with the high RR variability typical of AF episodes, and early occurrence of ventricular depolarisation, causing superimposition of QRST templates during the subtraction, might still cause imperfect QRST cancellation, producing artefacts in the residual signals. Moreover, the cancellation of a QRST belonging to a morphology class with only one element resulted in zeroing the QRST interval and in cancelling the atrial component as well.

The ABS technique implemented provides an atrial signal for each original channel. The availability of two residuals allows a further step of processing where this multi-channel atrial activity is exploited for the creation of an enhanced atrial activity signal (EAAS), where the significant parameters can be better evaluated. In fact, the EAAS is obtained as the residual with higher atrial activity, or in case of significant cross-correlation, as the sum or the difference of the two residuals.

A frequency domain analysis of the EAAS seemed more appropriate, considering the low amplitudes and the irregularities in the f-waves, which can easily jeopardise a time domain analysis. The parameters with probable significance evaluated from the atrial activity were as follows: frequency of the power spectral density peak (fPmaxb), otherwise called dominant atrial frequency (DAF); Pdmaxb (max value of the power spectral density); wb3dB (3 dB bandwidth of the peak); Pb3dB (power in the 3 dB band around the peak, the area under the curve in the power spectrum diagram).

Beat classification/clustering (morphological)

The adoption of ABS approaches requires beat classification or clustering in order to subtract from each beat the representative beat of its class. This classification is morphological because it is based only on the QRST complex morphology and not on ventricular rhythm considerations.

In the approach we adopted, QRST morphology clustering of both channels was performed using L1 (sum of absolute differences) distance. Horizontal wiggling was used, since the reference point provided by the QRS detection was not by itself sufficiently reliable for beat alignment. In addition, up and down vertical shifting was applied to compensate for resid-

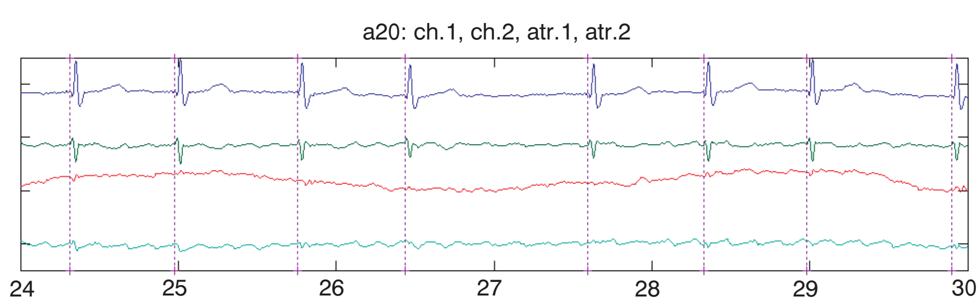

Figure 2. From top to bottom: channel 1 (original signal), channel 2 (original signal), extracted atrial activity of channel 1, extracted atrial activity of channel 2, in an excerpt from a record of the AF Termination Challenge Database. Vertical scales are different for the 4 signals.

ual baseline wandering. From each QRST complex in the ECG signal the corresponding QRST template was subtracted after its alignment using the horizontal and vertical references found in the previous step.

QRST morphology classification was carried out separately for each single lead (Figure 3), since the main purpose was the cancellation of the ventricular activity from the surface signals. Both leads should be used simultaneously for a classification related to physiological investigations. Because of the noise (mainly baseline wandering), the different overlap of the f-waves on the T-wave in different beats, the beat-to-beat variability, and the errors in the beat alignment, the QRST templates identified were usually a few more than the morphology templates actually present in the ECG signal.

The process of QRST cancellation produced two residual signals (one for each ECG channel) where the ventricular activity was cancelled. Figure 4 shows for the second channel, from the top, the ECG signal, the artificial signal composed of the QRST templates properly aligned with the original ECG, and the residual signal (atrial activity). The method was robust with respect to artefacts and was able to perform cancelling of extrasystoles effectively. The assessment of the performance was mainly made qualitatively because a quantitative evaluation was not possible, given that the AF Termination Challenge Database was not annotated in terms of beat type. Subsequently, further research was made in an attempt to refine and assess the method quantitatively using the annotated MIT-BIH Arrhythmia Database, which is annotated for a classification related to physiological investigation and has the same number of classes in both leads.

The algorithm was improved using a different pre-filtering in order to reduce the baseline wandering that could significantly alter the signal morphology. The two-lead ECG signal was first sub-sampled at 1/5 of the original frequency in order to pass it through a chain of median filters of 200 ms and 700 ms.[44] The first median filter has the goal of eliminating the QRSs while the second eliminates the T-waves. At the end of the filtering, a signal representing the course of the baseline was obtained. This signal, low-pass (1 Hz) filtered and re-sampled at the original frequency, was then subtracted from the original signal, giving the original two leads without the baseline wandering. Then, the signal was band-pass filtered in order to obtain a further reduction of the residual noise.

The classification problem of the pre-filtered signal is a typical problem of unsupervised clustering, because 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 a possible classification of all beats was performed, while in the second phase the created classes were re-estimated and, if necessary, redefined.

In order to begin the first phase of the grouping

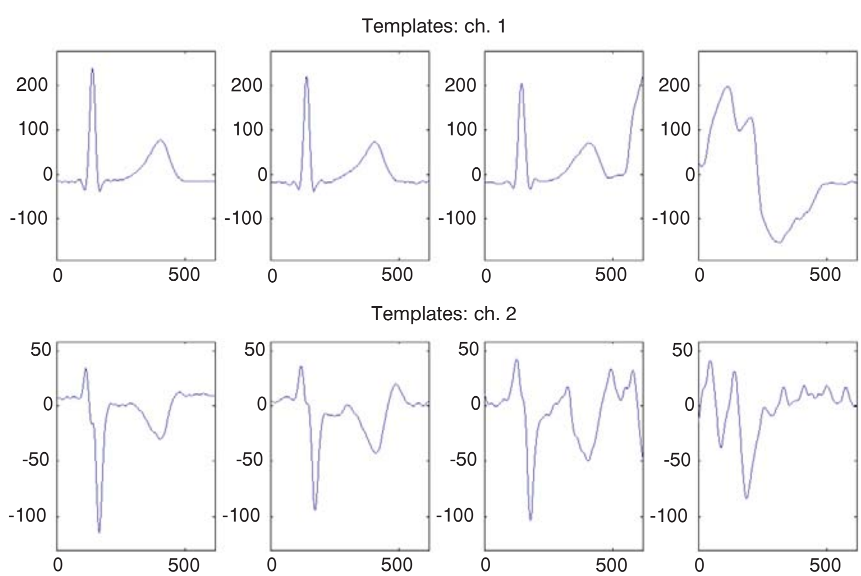

Figure 3. An example of QRST clusters from the AF Termination Challenge Database.

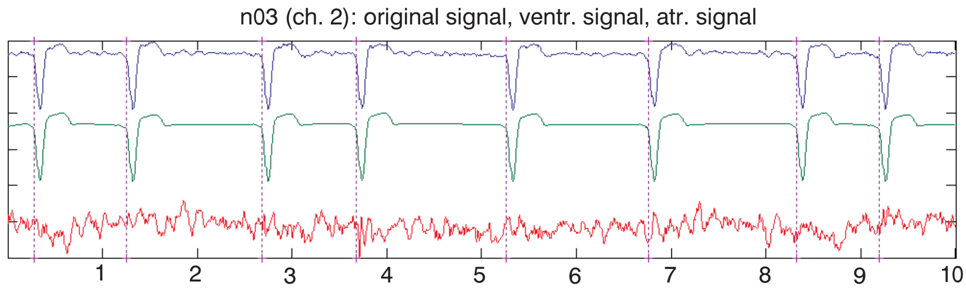

Figure 4. An excerpt from channel 2 in a record of the AF Termination Challenge Database with, from top to bottom, the original signal, and the extracted ventricular and atrial signals (with different vertical scales).

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.

The following algorithm was applied to each record of the database. First a rough centroid (template) is calculated for all beats of the signal, one for each lead, with duration approximately 0.6 seconds. 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 a distance vector. 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.

Figure 5 shows a comparison between the centroid of the dominant QRST complexes after this ini-

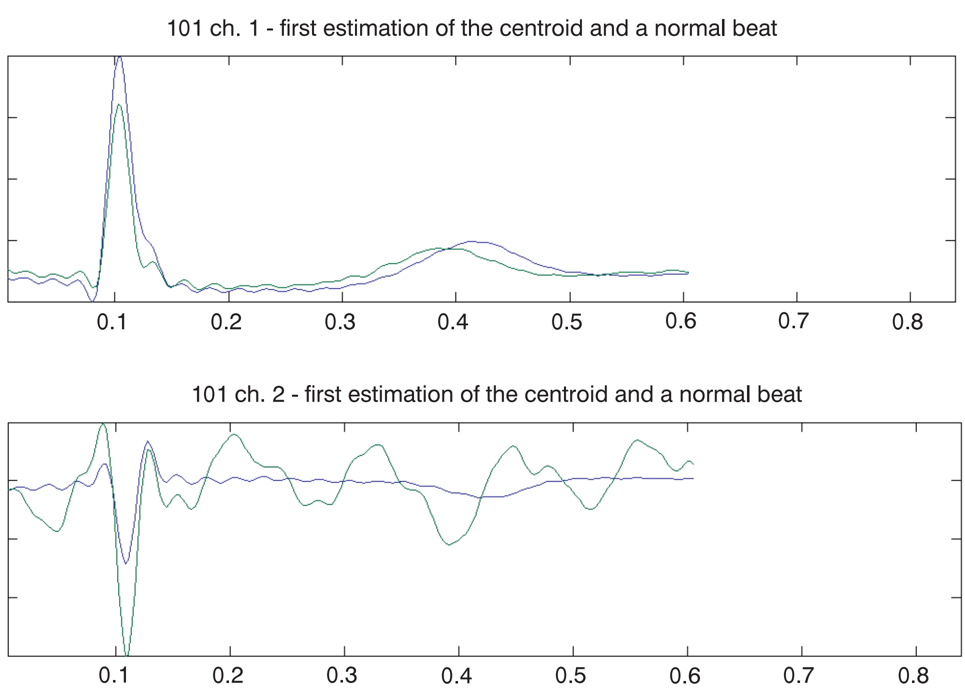

Figure 5. First estimation of the centroids (blue) and a normal beat (green) in a record of the MIT-BIH Arrhythmia Database.

tialisation phase and a generic dominant complex, for both channels.

Once this centroid has been obtained, the 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. Along with L1, the L2 distance is also calculated and these two features, combined with the centroid-to-beat correlation coefficient, are the main attributes 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 or not.

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 of the remaining beats 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 can be made (from the residual beats) and, thus, a final group with all residual unclassified beats is additionally formed.

In the second phase, all groups but the one with the dominant beats of the signal 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 degree of L1 similarity with the dominant group, and their status as dominant or non-dominant beats is re-evaluated. The centroids of each reformed group are compared to the centroid of the dominant group using the same features of the first phase. The beats of the groups whose comparison criteria are satisfied are put into the dominant ones. All other groups remain non-dominant. The validation criteria are very similar to the ones used in the first phase, but the thresholds are even stricter.

The last group, with the remaining unclassified beats, is 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.

This clustering algorithm, though, cannot be reliable without taking into consideration noise issues. In fact, noise can significantly change the morphology of a QRST complex, making its proper classification difficult. For this reason, before processing the two-channel signal to obtain a beat classification, the noise level is estimated for each detected QRS, adopting the same algorithm used for the noise estimation in the QRS detection. Based on this estimation for every complex, the criteria for the clustering algorithm of the first phase and for the process of the last group with the remaining beats, during the second phase, are slackened.

The performances of such a clustering algorithm, in relation only to the dominant class, were evaluated on all the records of the MIT-BIH Arrhythmia Database. Very satisfactory results were obtained for dominant class discrimination: sensitivity 98.84%, specificity 95.12%, positive predictive value 99.45% and negative predictive value 90.22%.

Spectral analysis of the atrial rhythm

For the feature extraction from the atrial activity, the non-stationary signal and the occurrence of residual artefacts led to the choice of the short time Fourier transform (STFT) applied to the EAAS. A Gaussian window of 3.33 s was selected in an attempt to obtain a satisfactory compromise between limiting the insertion of the residual time artefacts in too many spectra and obtaining a satisfactory frequency resolution. The time frequency representation obtained shows a DAF together with components short in time and wide in frequency, related to abrupt residual artefacts. Figure 6 shows the time frequency representations obtained for the learning set cases n05 (non-terminating AF) and t06 (terminating AF).

A mean power spectrum was obtained by discarding, for each frequency, 10% percentiles and averaging the STFT squared magnitudes in time, in a further attempt to minimise the effect of noise and artefacts in the evaluation of the atrial activity parameters. Figure 7 shows the resulting spectra for cases n05 of learning set N and t06 of learning set T.

Results

The PhysioNet/Computers in Cardiology Challenge 2004 made available only the classification of the

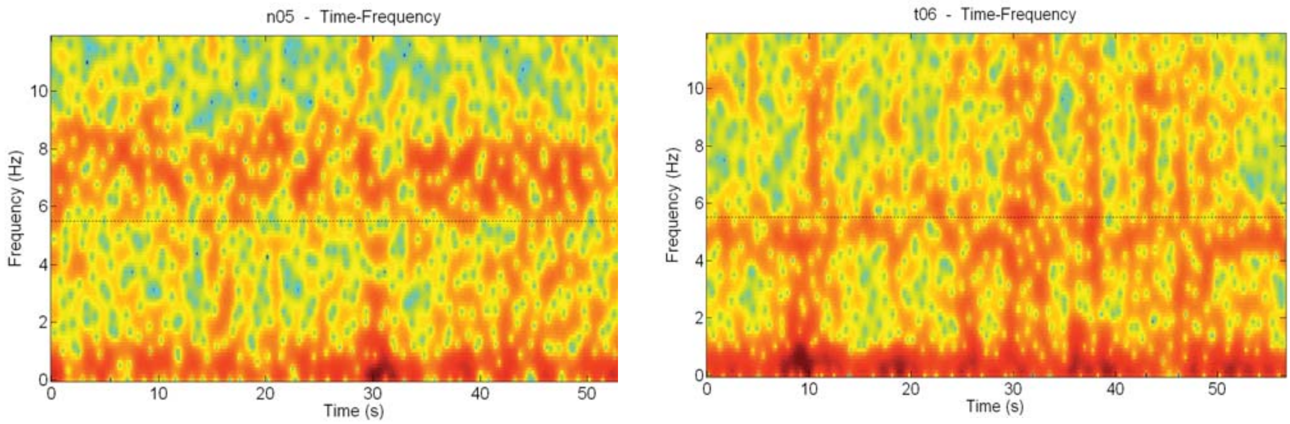

Figure 6. The time-frequency representation for record n05 (excluding the very low frequencies the dominant atrial frequency, DAF, can be perceived in the interval 6-8 Hz) and record t06 (excluding the very low frequencies the DAF can be perceived in the interval 4-5 Hz) of the AF Termination Challenge Database.

learning sets for the setup of the classifier. The classification of other 30 (test set A) and 20 (test set B) cases was not known and those cases were reserved for scoring the performance of each competitor. At the end of PhysioNet/Computers in Cardiology Challenge 2004, the true classification of the test sets became available and more statistical tests could be applied in order to get more information about the discriminating power of the features and to build a better classifier.

As previously mentioned, for an accurate statistical analysis, parameters other than the frequency of the power spectral density peak (DAF or fPmaxb) were extracted from the mean power spectrum of the AF component: Pdmaxb (max value of the power spectral density), wb3dB (3 dB bandwidth of the peak), Pb3dB (power in the 3 dB band around the peak, the area under the curve in the power spectrum diagram). In addition, the following ventricular activity parameters were considered: RRmean (average RR interval), HRmean (average heart rate) and RRstd (standard deviation of the RR intervals).

Discrimination between N- and T-types

The frequency of the power spectral density peak (DAF or fPmaxb) revealed great power in discriminating N-type and T-type records on first analysis. On the N/T learning set, the dominant frequency of the atrial component had a discriminant power much higher than the other features. The second most powerful parameter was the HRmean and the third was RR mean, highly correlated with HRmean. On the 50 cases of the entire N/T dataset (including the test set A, once its classification became publicly available) ANOVA univariate statistics provided similar results to those obtained on the N and T learning sets, with the exception of the Pdmaxb and wb3dB, which were

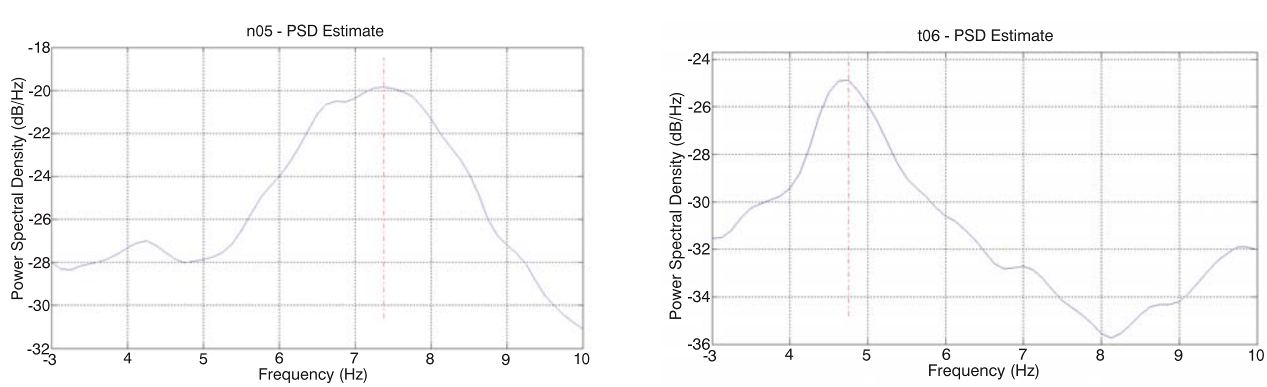

Figure 7. The average power spectrum with the indication of the peak for records n05 and t06 of the AF Termination Challenge Database.

not previously significant but now achieved a moderate discriminating power. The feature fPmaxb had the highest discriminating power; its statistical significance was more than 10 times greater than the maximum of the other features (HRmean).

The use of only the fPmaxb feature resulted in 46/50 (92%) correct classifications, confirming the highest discriminating power of fPmaxb with respect to the other parameters. The discrimination between non-terminating (N-type cases) and those terminating immediately after the end of registration (T-type cases) was effective using the fPmaxb and HRmean features.

Discrimination between N- and S-types

The discrimination between N-type and S-type cases could provide some information about the possibility of distinguishing cases of medium- or long-term terminating AF cases from non-terminating cases. The overall N-S dataset consists of 46 cases, including the records of the N and S learning sets and the N-type and S-type cases extracted from test sets A and B (once their classification became publicly available).

Although the performance in discriminating N/S-type was slightly lower than that obtained in N/T-type discrimination, the robustness of the parameter fPmaxb and the small decrease in performance obtained using this feature (91% correct classifications as best result) with respect to that obtained in the task of N/T-type discrimination (92%) leads us to be optimistic about the possibility of distinguishing AF which will terminate in the medium or long term from the non-terminating form.

Discussion and conclusions

This article has focused on a detailed discussion of a specific approach to the detection and classification of AF on the surface ECG. It is worth stressing the value of public, annotated ECG databases in the evaluation and development of this and other approaches. Of course, the use of a standardised data set limits the conclusions that may be drawn from the application of a given algorithm; on the other hand, it allows the direct comparison of different approaches to a specific problem, without the need to speculate about possible confounding factors. It can also provide a target for developers, who in the PhysioNet/Computers in Cardiology Challenge 2004 were challenged to achieve an acceptable level of compliance between the results of their computer algorithms and the consensus estimate of a panel of human referees.[45]

Methods of information processing should never be completely standardised, as this would tend to stifle the creative thought processes that can sometimes find solutions to the most intractable problems of analysis. Nevertheless, the existence of fixed points of reference, with consensually established definitions and rules of use,[46] offers a common proving ground for the testing and improvement of diverse analytical methods and can thus enhance the developmental process by highlighting the pros and cons of different approaches.

In the PhysioNet/Computers in Cardiology Challenge 2004, our best result (27/30 on a blind test set) was obtained by cancelling the QRST from the ECG by cross-channel adaptive filtering and classifying the cases by a linear classifier (based on fPmaxb and average heart rate) whose coefficients were chosen manually on the basis of the results obtained from the learning set.[47]

Subsequently, many improvements were made to each step of the first approach. A more accurate ventricular activity cancellation was performed with morphological classification, interpolation and wiggling for class assignment and cancellation, through class centroid subtraction for each ECG leads.

The implemented method, based on average cluster template subtraction, produced in principle a residual signal with fewer zeroed or artificially reconstructed areas than other methods from the PhysioNet/Computers in Cardiology Challenge 2004, which performed a previous zeroing of premature ventricular complexes and aberrant beats,[48] or a blanking of the QRS complex (without the T-wave) after the averaged QRST subtraction by interpolating the signal over the QRS part of each detected heart beat.[49] However, no significant improvement was obtained with respect to the first attempt and the best parameter subset was still composed of fPmaxb and HRmean. Inclusion of the latter feature improved only slightly on the performance of fPmaxb alone.

Other relevant studies participating in the PhysioNet/Computers in Cardiology Challenge 2004[48,49] obtained results comparable with the method described here. Summarising, several of the most successful approaches began with the subtraction of the QRST complexes from the ECG signals, obtaining a residual signal dominated by atrial activity that was then analysed in the frequency domain. In these studies, evidence collected from the analysis of the learning set supported the hypothesis that atrial activity slows and regularises before the self-termination of AF,[50] as has also been confirmed by other researchers.[18,19] From a physiological point of view, these results are consistent with the existing AF theories, such as bioelectric remodelling, since the atrial refractory cycle (closely related to the inverse of DAF) tends to decrease with the maintenance of the arrhythmia.[48,51,52]

This work confirms that success rates greater than 90% can easily be achieved in the N/T-type discrimination. Furthermore, we may be hopeful that the good values obtained in the N/S-type discrimination task (89-91%) can be used in the identification of medium- or long-term terminating AF episodes.

Clinical implications

Even though a huge database containing non-terminating AF cases and AF that spontaneously terminated after different time epochs (ten minutes, half an hour, hours, ten hours) would be necessary for a more robust assessment of the discrimination capabilities of the secondary features (apart from DAF) and a more accurate investigation of the possibility of long-term prediction, the results discussed above are very relevant. The diagnostic method implemented for the discrimination between terminating and non-terminating AF, using non-invasive techniques and specifically surface ECG signals, seems to have the potential for future use in clinical applications, providing a clear benefit for patients, doctors and health systems, since it can help in avoiding useless therapeutic intervention (cardioversion) and minimising the risks for the patient. Indeed, the decision as to whether and how to treat AF, with medication or electrical cardioversion, is a common and sometimes crucial dilemma in clinical practice. Since AF terminates spontaneously in more than 30% of cases of AF seen by physicians in offices and hospitals,[53] the physician should try to avoid unnecessary therapy, with its attendant risks, while at the same time looking out for the patient's best interest. If an accurate and reliable way could be found to predict the self-termination of AF even a short time in advance, the benefit would be great in terms of patient comfort and safety, and cost to health care providers. Of course, an approach like the one presented here needs considerably more testing, ideally in randomised, controlled studies, before being applied in routine clinical practice.

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 author would like to thank M. Varanini and his co-workers at the Institute of Clinical Physiology of the Italian Research Council in Pisa, Italy and Prof. G. Vrouchos of ICU/CCU Department of the Venizeleio-Pananeio Hospital in Heraklion, Crete, Greece for their valuable collaboration in the achievement of some significant outcomes described in this paper.

Furthermore, the author would like to thank his colleagues at the Biomedical Informatics Laboratory of ICS-FORTH for their invaluable support and co-operation.

References

1. Fuster V, Rydén LE, Asinger RW, et al. ACC/AHA/ESC guidelines for the management of patients with atrial fibrillation. A report of the American College of Cardiology/American Heart Association Task Force on Practice Guidelines and the European Society of Cardiology Committee for Practice Guidelines and Policy Conferences (Committee to develop guidelines for the management of patients with atrial fibrillation) developed in collaboration with the North American Society of Pacing and Electrophysiology. Eur Heart J. 2001; 22: 1852-1923.
2. Fuster V, Rydén LE, Cannom DS, et al. ACC/AHA/ESC 2006 guidelines for the management of patients with atrial fibrillation: a report of the American College of Cardiology/American Heart Association Task Force on Practice Guidelines and the European Society of Cardiology Committee for Practice Guidelines (Writing Committee to Revise the 2001 Guidelines for the Management of Patients With Atrial Fibrillation). J Am Coll Cardiol. 2006; 48: 854-906.
3. Turpie AGG. New oral anticoagulants in atrial fibrillation. Eur Heart J. 2008; 29: 155-165.
4. Heeringa J, van der Kuip DA, Hofman A, et al. Prevalence, incidence and lifetime risk of atrial fibrillation: the Rotterdam study. Eur Heart J. 2006; 27: 949-953.
5. Vardas PE, Mavrakis HE. Atrial fibrillation: a symptom treated as a disease? Hellenic J Cardiol. 2006; 47: 191-193.
6. Greenlee RT, Vidaillet H. Recent progress in the epidemiology of atrial fibrillation. Curr Opin Cardiol. 2005; 20: 7-14.
7. Levy S, Camm AJ, Saksena S, et al. International consensus on nomenclature and classification of atrial fibrillation: A collaborative project of the Working Group on Arrhythmias and the Working Group of Cardiac Pacing of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology. J Cardiovasc Electrophysiol. 2003; 14: 443-445.
8. Ganotakis ES, Mikhailidis DP, Vardas PE. Atrial fibrillation, inflammation and statins. Hellenic J Cardiol. 2006; 47: 51-53.
9. Pipberger HV, Arms RJ, Stallmann FW. Automatic screen-

ing of normal and abnormal electrocardiograms by means of digital electronic computer. Proc Soc Exp Biol Med. 1961; 106: 130-132.
10. Caceres CA, Steinberg CA, Abraham S, et al. Computer extraction of electrocardiographic parameters. Circulation. 1962; 25: 356-362.
11. Willems JL, Arnaud P, van Bemmel JH, Degani R, Macfarlane PW, Zywietz C. Common standards for quantitative electrocardiography: goals and main results. CSE Working Party. Methods Inf Med. 1990; 29: 263-271.
12. Macfarlane PW. A brief history of computer-assisted electrocardiography. Methods Inf Med. 1990; 29: 272-281.
13. Chiarugi F, Lombardi D, Lees PJ, Chronaki CE, Tsiknakis M, Orphanoudakis SC. Support of daily ECG procedures in a cardiology department via the integration of an existing clinical database and a commercial ECG management system. Ann Noninvasive Electrocardiol. 2002; 7: 263-270.
14. Rieta JJ, Hornero F. Comparative study of methods for ventricular activity cancellation in atrial electrograms of atrial fibrillation. Physiol Meas. 2007; 28: 925-936.
15. Alcaraz R, Rieta JJ. Non-linear organization analysis of paroxysmal atrial fibrillation. Conf Proc IEEE Eng Med Biol Soc. 2007; 2007: 1957-1960.
16. Petrutiu S, Sahakian AV, Swiryn S. Abrupt changes in fibrillatory wave characteristics at the termination of paroxysmal atrial fibrillation in humans. Europace. 2007; 9: 466-470.
17. Alcaraz R, Rieta JJ. Non-linear analysis of the main atrial wave to estimate organization in paroxysmal atrial fibrillation. Proc Computers in Cardiology. 2007; 34: 545-548.
18. Bollmann A, Lombardi F. Electrocardiology of atrial fibrillation. Current knowledge and future challenges. IEEE Eng Med Biol Mag. 2006; 25: 15-23.
19. Giazitzoglou E, Triantafyllis D, Katritsis D. Organised atrial fibrillation simulating atrial flutter. Hellenic J Cardiol. 2006; 47: 123-127.
20. Bollmann A, Husser D, Mainardi L, et al. Analysis of surface electrocardiograms in atrial fibrillation: techniques, research, and clinical applications. Europace. 2006; 8: 911-926.
21. Stridh M, Sörnmo L. Spatiotemporal QRST cancellation techniques for analysis of atrial fibrillation. IEEE Trans Biomed Eng. 2001; 48: 105-111.
22. Langley P, Rieta JJ, Stridh M, et al. Comparison of atrial signal extraction algorithms in 12-lead ECGs with atrial fibrillation. IEEE Trans Biomed Eng. 2006; 53: 343-346.
23. Rieta JJ, Castells F, Sanchez C, et al. Atrial activity extraction for atrial fibrillation analysis using blind source separation. IEEE Trans Biomed Eng. 2004; 51: 1176-1186.
24. PhysioNet: the research resource for complex physiologic signals [Internet]. Available from: http://www.physionet.org/.
25. Spontaneous termination of atrial fibrillation: a challenge from PhysioNet and Computers in Cardiology 2004 [Internet]. Available from: http://www.physionet.org/challenge/2004/.
26. MIT-BIH Arrhythmia Database [Internet]. Available from: http://www.physionet.org/physiobank/database/mitdb/.
27. Kohler BU, Hennig C, Orglmeister R. The principles of software QRS detection. IEEE Eng Med Biol Mag. 2002; 21: 42-57.
28. Chiarugi F, Varanini M, Cantini F, Conforti F, Vrouchos G. Noninvasive ECG as a tool for predicting termination of paroxysmal atrial fibrillation. IEEE Trans Biomed Eng. 2007; 54: 1399-1406.
29. Chiarugi F, Sakkalis V, Emmanouilidou D, Krontiris T, Varanini M, Tollis I. Adaptive threshold QRS detector with best channel selection based on a noise rating system. Proc Computers in Cardiology. 2007; 34: 157-160.
30. Hamilton PS. Open source ECG analysis software documentation. Somerville, MA, USA: E P Limited; 2002.
31. Christov II. Real time electrocardiogram QRS detection using combined adaptive threshold. Biomed Eng Online. 2004; 3: 28.
32. Dotsinsky I. Atrial wave detection algorithm for discovery of some rhythm abnormalities. Physiol Meas. 2007; 28: 595-610.
33. Carrault G, Cordier MO, Quiniou R, et al. Temporal abstraction and inductive logic programming for arrhythmia recognition from electrocardiograms. Artif Intell Med. 2003; 28: 231-263.
34. Hernandez AI, Carrault G, Mora F. Improvement of a P-wave detector by a bivariate classification stage. Trans Inst Meas Control. 2000; 22: 231-242.
35. Portet F. P wave detector with PP rhythm tracking: evaluation in different arrhythmia contexts. Physiol Meas. 2008; 29: 141-155.
36. Goutas A, Ferdi Y, Herbeuval JP, et al. Digital fractional order differentiation-based algorithm for P and T-waves detection and delineation. ITBM-RBM. 2005; 26: 127-132.
37. Laguna P, Jane R, Caminal P. Automatic detection of wave boundaries in multilead ECG signals: validation with the CSE database. Comput Biomed Res. 1994; 27: 45-60.
38. Martinez JP, Almeida R, Olmos S, et al. A wavelet-based ECG delineator: evaluation on standard databases. IEEE Trans Biomed Eng. 2004; 51: 570-581.
39. Sovilj S, Jeras M, Magjarevic R. Real time P-wave detector based on wavelet analysis. Proc 12th IEEE Mediterranean Electrotechnical Conf. (MELECON 2004). 2004; 2004: 403-406.
40. Sternickel K. Automatic pattern recognition in ECG time series. Comput Methods Programs Biomed. 2002; 68: 109-115.
41. Graja S, Boucher JM. Hidden Markov tree model applied to ECG delineation. IEEE Trans Instrum Meas. 2005; 54: 2163-2168.
42. Lemay M, Jacquemet V, Forclaz A, et al. Spatiotemporal QRST cancellation method using separate QRS and T-waves templates. Proc Computers in Cardiology. 2005; 32: 611-614.
43. Stridh M, Sornmo L, Meurling CJ, et al. Sequential characterization of atrial tachyarrhythmias based on ECG time-frequency analysis. IEEE Trans Biomed Eng. 2004; 51: 100-114.
44. 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.
45. Macfarlane PW, Willems JL. The CSE European Working Party. The CSE project: Progress as viewed by the cooperating centres. In: Selvester RH, Geselowitz D, editors. Computerized interpretation of the electrocardiogram. Proc Engineering Foundation; 1983. p. 293.
46. Fayn J, Rubel P, Macfarlane PW. Can the lessons learned from the assessment of automated electrocardiogram analysis in the Common Standards for Quantitative Electrocardiography study benefit measurement of delayed contrast-enhanced magnetic resonance images? J Electrocardiol. 2007; 40: 246-250.
47. Cantini F, Conforti F, Varanini M, Chiarugi F, Vrouchos G. Predicting the end of an atrial fibrillation episode: the PhysioNet challenge. Proc Computers in Cardiology. 2004; 31: 121-124.
48. Petrutiu S, Sahakian SV, Ng J, et al. Analysis of the surface electrocardiogram to predict termination of atrial fibrillation: The 2004 Computers in Cardiology/PhysioNet Challenge. Proc Computers in Cardiology. 2004; 31: 105-108.
49. Hayn D, Edegger K, Scherr D, et al. Automated prediction of

spontaneous termination of atrial fibrillation from electrocardiograms. Proc Computers in Cardiology. 2004; 31: 117-120.

50. Moody GB. Spontaneous termination of atrial fibrillation: a challenge from PhysioNet and Computers in Cardiology 2004. Proc Computers in Cardiology. 2004; 31: 101-104.
51. Mora C, Castells J, Ruiz R, et al. Prediction of spontaneous termination of atrial fibrillation using time frequency analysis of the atrial fibrillatory wave. Proc Computers in Cardiology. 2004; 31: 109-112.
52. Goette A, Honeycutt C, Langberg JJ. Electrical remodeling in atrial fibrillation. Time course and mechanisms. Circulation. 1996; 94: 2968-2974.
53. Goudevenos JA, Vakalis JN, Giogiakas V, et al. An epidemiological study of symptomatic paroxysmal atrial fibrillation in northwest Greece. Europace. 1999; 1: 226-233.