Low-Complexity Detection of Atrial Fibrillation in Continuous Long-Term Monitoring

Andrius Petrėnas[a,*], Vaidotas Marozas[a], Leif Sörnmo[b]

[a]*Biomedical Engineering Institute, Kaunas University of Technology, Kaunas, Lithuania*
[b]*Department of Biomedical Engineering and Center of Integrative Electrocardiology, Lund University, Lund, Sweden*

Abstract

This study describes an atrial fibrillation (AF) detector whose structure is well-adapted both to detection of subclinical AF episodes and to implementation in a battery-powered device for use in continuous long-term monitoring applications. A key aspect for achieving these two properties is the use of an 8-beat sliding window, which thus is much shorter than the 128-beat window used in most existing AF detectors. The building blocks of the proposed detector include ectopic beat filtering, bigeminal suppression, characterization of RR interval irregularity, and signal fusion. With one design parameter, the performance can be tuned to put more emphasis on avoiding false alarms due to non-AF arrhythmias or more emphasis on detecting brief AF episodes. Despite its very simple structure, the results show that the detector performs better on the MIT–BIH Atrial Fibrillation database than do existing detectors, with high sensitivity and specificity (97.1% and 98.3%, respectively). The detector can be implemented with just a few arithmetical operations and does not require a large memory buffer thanks to the short window.

Keywords:
Atrial fibrillation, RR-based detection, low-complexity detection, continuous ambulatory monitoring

1. Introduction

In recent years, atrial fibrillation (AF) monitoring has been performed in connection with cryptogenic ischemic stroke [1], interventional ablation procedures [2], and pharmacological treatment [3], providing information for evaluating the efficacy of different treatment strategies. Ultimately, continuous long-term monitoring, lasting from several weeks to months, should be performed so that all episodes of paroxysmal AF are detected, including the very brief ones, and, therefore, it is essential that detection techniques are developed subject to the constraint that the implementation can be made extremely energy efficient.

It has been shown that the success rate of catheter ablation is highly overestimated when determined from 24-h Holter recordings [2]. This issue can be addressed by considerably extending the monitoring period so that the likelihood of detecting AF episodes increases. However, existing techniques for continuous long-term monitoring reduce the patient's quality of life and often lead to premature termination of the data acquisition [4]. While patient comfort can be improved by shrinking the size of the monitoring device, using a smaller battery, large battery capacity is nonetheless needed in this type of monitoring since battery replacement or recharging should be avoided.

Many different approaches to AF detection have been proposed in recent years: the majority of the detectors take the RR interval series as the starting point [5, 6, 7, 8], whereas a few detectors also involve information on P wave and f wave morphology [9, 10, 11]. However, the latter type of detector is associated with dramatically higher computational demands than are the RR-based detectors, making them unsuitable for implementation in a low power device. Somewhat surprisingly, it has been demonstrated that increased complexity through inclusion of morphologic analysis does not necessarily translate to performance better than that achieved by an RR-based AF detector [9, 10].

Crucial issues to be solved in RR-based AF detection are 1. improved detection rate of brief ("occult") episodes of paroxysmal AF (<30 s) and 2. reduction of the false alarm rate due to ectopic beats and other irregular types of rhythm. While almost all AF detectors require at least a 30-s episode for correct detection, recent studies have suggested that episodes as short as 5 s are of interest to detect since they are associated with increased risk of stroke [12, 1, 13]. Therefore, reliable detection of brief episodes is essential so that an accurate evaluation of AF burden can be produced when studying thromboembolic risk.

In this paper, we propose a novel RR-based AF detector with a low complexity structure that can be implemented with few arithmetical operations. The detector involves blocks for preprocessing, bigeminal suppression, characterization of RR irregularity, signal fusion, and threshold detection; both off- and online versions of the detector are presented (Sec. 2). Using annotated ECG databases as well as simulated signals, both described in Sec. 3, detector performance is investigated in Sec. 4 as a function of different design parameters as well as a function of AF episode length. The resulting performance is compared to that achieved by other AF detectors in the literature.

*Correspondence to: Biomedical Engineering Institute, Kaunas University of Technology, K. Barsausko st. 59, LT-51423 Kaunas, Lithuania. Tel.: +370 37 30 05 28. E-mail address: andrius.petrenas@ktu.lt (A. Petrėnas).
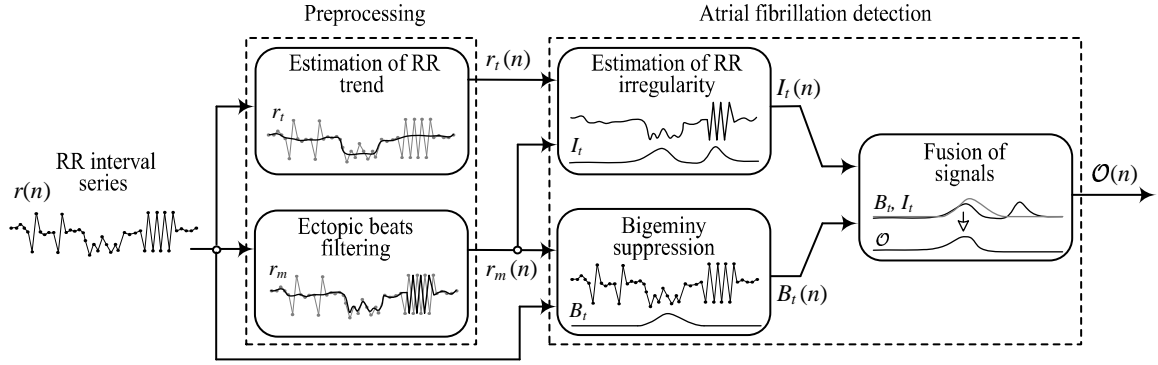
Figure 1: Block diagram of the proposed AF detector.

2. Methods

The proposed detector is based on the observation that AF episodes have increased RR irregularity and usually are associated with increased heart rate. A block diagram of the detector is shown in Fig. 1, where each of the procesing blocks is described in the following.

2.1. Preprocessing

The occurrence of sporadic ectopic beats can be problematic when short time RR series are analyzed and, therefore, it is desirable to introduce techniques which can handle such situations. Preliminary tests showed that the simple 3-point median filter is useful for reducing the influence of ectopic beats in the RR series. This filter is defined by

$$r_m(n) = \text{median}\{r(n-1), r(n), r(n+1)\}, \qquad (1)$$

where $r(n)$ denotes the length of the n:th RR interval (whose unit is in seconds). Median filtering is also useful for rejecting outlier RR intervals due to, e.g., missed QRS complexes. Higher-order median filters were found to be less useful since they smooth AF episodes to such a degree that episodes with low RR irregularity remain undetected.

Since heart rate usually increases during AF episodes, an estimate of the mean RR interval has to be determined, and employed as a feature in the AF detector, cf. [7, 14]. Here, the traditional ensemble averager is replaced by the exponential averager to better track the "trend" in the RR interval series. The exponential averager is defined by [15]

$$r_t(n) = r_t(n-1) + \alpha(r(n) - r_t(n-1)), \qquad (2)$$

where α $(0 < \alpha < 1)$ determines the degree of smoothing, i.e., the lowpass cut-off frequency. Since the exponential averager in (3) has nonlinear phase, forward–backward filtering is performed to achieve linear (null) phase.

2.2. RR interval irregularity

In a sliding detection window of length N, located at time n, the number of all pairwise RR interval combinations differing more than γ seconds is determined, and normalized with its maximum value $N(N-1)/2$, i.e.,

$$M(n) = \frac{2}{N(N-1)} \sum_{j=0}^{N-1} \sum_{k=j+1}^{N} H(|r(n-j) - r(n-k)| - \gamma), \quad (3)$$

where $H(\cdot)$ is the Heaviside step function and $0 \leqslant M(n) \leqslant 1$. It is noted that $M(n)$ is partially based on the same principle as is sample entropy estimation [7].

The primary feature of RR irregularity is provided by the ratio between a smoothed version of $M(n)$ and the RR interval trend $r_t(n)$ in (2),

$$I_t(n) = \frac{M_t(n)}{r_t(n)}, \qquad (4)$$

where $M_t(n)$ is obtained by exponential averaging of $M(n)$. The division by $r_t(n)$ is motivated by the wish to emphasize RR irregularity at higher heart rates. It should be noted that $I_t(n)$ is close to 0 for regular rhythms since the difference between pairs of RR intervals is usually smaller than a properly chosen γ, whereas $I_t(n)$ approaches 1 during AF, see the examples in Fig. 2.

2.3. Bigeminy suppression

It is well-known that bigeminy can be incorrectly interpreted as AF when the detection is RR-based [14]. In order to address this issue, we introduce another measure of RR irregularity, complementary to $I_t(n)$ and with the property of being essentially indifferent to the presence of bigeminy. The measure is defined by

$$B(n) = \left(\frac{\sum\limits_{j=0}^{N-1} r_m(n-j)}{\sum\limits_{j=0}^{N-1} r(n-j)} - 1 \right)^2, \qquad (5)$$

where N is an even-valued integer. Similar to $r_t(n)$ and $M_t(n)$, $B_t(n)$ results from exponential averaging of $B(n)$. For bigeminy as well as for regular rhythms, the ratio in (5) is approximately 1 since $r_m(n)$ and $r(n)$ are similar, and thus $B_t(n)$ is approximately 0, see Fig. 2(b). On the other hand, for AF the variation
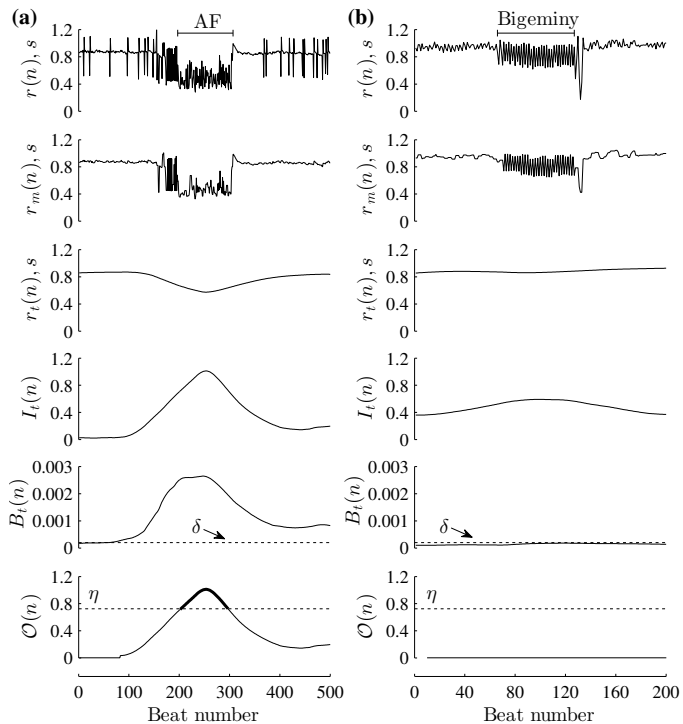
Figure 2: The output of each building block of the proposed detector is illustrated for (a) an AF episode surrounded by ectopic beats and (b) bigeminy. Atrial fibrillation is detected whenever $O(n)$ exceeds the threshold η, then displayed by a thicker line.

in $r_m(n)$ is lower than that in $r(n)$ because of the median filtering, and thus $B_t(n)$ will increase so that it indicates irregularity. The squaring operation in (5) improves the differentiation of AF from other rhythms.

2.4. Signal fusion and detection

Simple signal fusion is employed to produce the decision function $O(n)$: $O(n)$ is identical to $B_t(n)$, unless $B_t(n)$ exceeds a fixed threshold δ when, instead, it becomes identical to $I_t(n)$, i.e.,

$$O(n) = \begin{cases} I_t(n), & B_t(n) \geqslant \delta \\ B_t(n), & B_t(n) < \delta. \end{cases} \quad (6)$$

Figure 2 illustrates $O(n)$ as well as $r_m(n), r_t(n), I_t(n)$, and $B_t(n)$ for an AF episode surrounded by ectopic beats and sinus rhythm with an episode of bigeminy. It is obvious from Fig. 2(a) that median filtering removes the RR intervals related to ectopic beats so that the AF episode can be correctly detected and false alarms avoided. Another feature of the detector is illustrated in Fig. 2(b) where $O(n)$ is shown to be unaffected by the occurrence of bigeminy and, consequently, not misclassified as AF. Although $I_t(n)$ increases during bigeminy, $O(n)$ remains small since $B_t(n)$ remains below δ, cf. (6). Similarly, $B_t(n)$ is close to 0 during sinus rhythm, and therefore $O(n)$ is also close to 0. Atrial fibrillation is detected whenever $O(n)$ exceeds the fixed threshold η.

The number of mathematical operations needed to implement the algorithm is presented in Table 1, showing that the detector requires very few multiplications/divisions for processing

a single RR interval. Of the 8 multiplications, 6 are required for implementation of the forward–backward exponential averager due to multiplication with α, If needed, the α-related multiplications can be approximated by additions and a shift, see Discussion.

Table 1: The number of arithmetic operations required per RR interval.

Multiplications	Divisions	Additions/subtractions
8	2	45

2.5. Online atrial fibrillation detection

Since forward–backward filtering requires that the time-reversed signal is processed, this type of filtering is best suited for offline processing. On the other hand, for online processing, forward–backward filtering with the exponential averager in (2) is replaced by forward filtering and a second-order exponential averager, defined by

$$r_t(n) = \alpha^2 r(n) + 2(1 - \alpha)r_t(n - 1) - (1 - \alpha)^2 r_t(n - 2). \quad (7)$$

To a minor extent, detector performance depends on the phase response of the selected exponential averager. Depending on the choice of α, the exponential averager in (7) produces a group delay that has to be taken into account. Accordingly, α is set to the same value for the respective exponential averagers associated with $r_t(n)$, $I_t(n)$, and $B_t(n)$ to ensure identical group delay.

In the following, the AF detector is referred to as either *offline* or *online*, depending on whether filtering is performed according to (2) or (7). These two versions have the same computational complexity.

3. Datasets and performance evaluation

3.1. Clinical signals

The algorithm was developed on the Long Term Atrial Fibrillation (LTAF) database [16, 17], composed of 84 ECG recordings from patients with paroxysmal or persistent AF, most recordings with a 24-h duration. The entire database consists of nearly 9 million beats of which 59% occur during AF.

The MIT–BIH Atrial Fibrillation database [18, 17] and the MIT–BIH Normal Sinus Rhythm (NSR) database [17] were used for performance evaluation. The MIT–BIH AF database includes 25 AF recordings of approximately 10-h duration, and contains in total more than 1 million beats, of which 43% occur during AF. Two subsets of the AF database were also analyzed in order to facilitate the comparison with published results, namely, by excluding records 00735 and 03665 since these only contain RR interval information [19, 9], and by excluding records 04936 and 05091 since these contain incorrect annotations [5, 20]. The resulting two subsets are labeled AFDB$_1$ and AFDB$_2$, respectively.

The MIT–BIH NSR database contains 18 ECG recordings of approximately 24-h duration, with a total of almost 2 million

beats. Since no significant arrhythmias are present, it is well-suited for evaluation of detector specificity. To facilitate the comparison, the performance of the proposed detector has been evaluated on the commonly used combinations of the MIT–BIH AF and the MIT–BIH NSR databases.
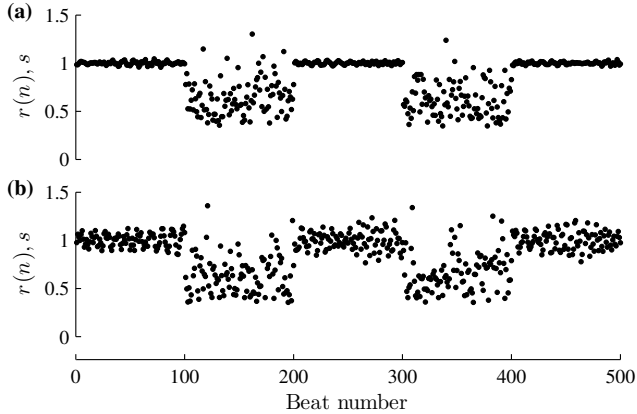
Figure 3: Two examples of paroxysmal AF, each containing two AF episodes with 100 beats; note that the RR interval series are simulated. The mean and standard deviation of the heart rate during sinus rhythm is (a) 60 ± 1 beats per minute (bpm) and (b) 60 ± 5 bpm.

3.2. Test signals

The emerging interest in the detection of brief episodes of paroxysmal AF [1, 12] motivated us to investigate performance as a function of episode duration. Due to the lack of annotated databases with brief episodes, the recently proposed paroxysmal AF simulator in [11] has been applied. RR intervals during sinus rhythm are generated by the model in [21], whereas an atrioventricular node model is used to generate RR intervals during AF [22]. When combining these two models, different RR interval series can be generated with variable episode duration, mean heart rate, and heart rate variability as illustrated in Fig. 3.

Detector performance was tested using different RR series with alternating episodes of SR and AF. The entire database with simulated signals consists of four datasets, each dataset defined by the mean and standard deviation of the heart rate during sinus rhythm, namely, 60 ± 1, 60 ± 5, 100 ± 1, and 100 ± 5 bpm. Each dataset is, in turn, divided into nine subsets with fixed AF episode lengths, ranging from 20 to 180 beats in steps of 20 beats. Thus, the database is composed of 36 subsets, each consisting of 5000 RR intervals for which 100 realizations were computed.

3.3. Performance measures

The performance was investigated in terms of the area under the curve (A) of the receiver operating characteristic (ROC), sensitivity (Se) and specificity (Sp). Sensitivity is defined by the number of correctly detected AF beats divided by the total number of AF beats, whereas specificity is defined by the number of correctly detected non-AF beats divided by the total number of non-AF beats. All other types of rhythm, including atrial flutter, were labeled as non-AF.

Figure 4: The influence of different parameters on detection performance in terms of A: (a) α and γ for $N = 8$ and $\delta = 2 \cdot 10^{-4}$, (b) α and N for $\gamma = 0.03$ and $\delta = 2 \cdot 10^{-4}$, and (c) α and δ for $\gamma = 0.03$ and $N = 8$. These results are based on the LTAF database.

4. Results

4.1. Parameter settings

Figure 4 displays A as a function of the exponential averaging parameter α for the LTAF database. For all investigated values of γ, N, and δ, the results show that A improves as α decreases. While these results suggest that a small α should be chosen, such a choice also means that the risk of missing brief episodes of paroxysmal AF increases. Therefore, it is important to complement the results in Fig. 4 with others that pinpoint detection performance as a function of episode length.

Using simulated signals, Fig. 5 demonstrates that the detection of brief episodes improves as α increases. For example, for $\alpha = 0.1$, an area under the curve of $A = 0.92$ is obtained for 20-beat episodes, whereas, for $\alpha = 0.02$, the same value of A is obtained for 60-beat episodes. Hence, it can be concluded from the results in Figs. 4 and 5 that the choice of α should be a trade-off between avoiding false alarms due to non-AF arrhythmias (calling for a small α) and detecting brief AF episodes (calling for a large α).

Another important result conveyed by Fig. 5 is that better performance is obtained in situations with low heart rate and low variability (i.e., 60 ± 1 bpm) than in situations with high heart rate and high variability (i.e., 100 ± 5 bpm).

Based on the results in Figs. 4 and 5, the following parameter values were chosen as a trade-off between the above-mentioned performance aspects: $\gamma = 0.03$ s, $N = 8$, $\delta = 2 \cdot 10^{-4}$, and $\alpha = 0.02$. Unless otherwise stated, these values are used in the following.

4.2. Detection threshold

Figure 6 displays the distribution of the detector output $O(n)$ for AF and non-AF beats in the LTAF database. Based on the properties of these two distributions, the detection threshold η was chosen as that particular value where sensitivity and specificity are identical, i.e., $\eta = 0.725$.

4.3. Briefest episode length

The briefest possible AF episode that can be detected was determined by means of simulated RR intervals with one, single AF episode. Starting with an episode length of five beats,
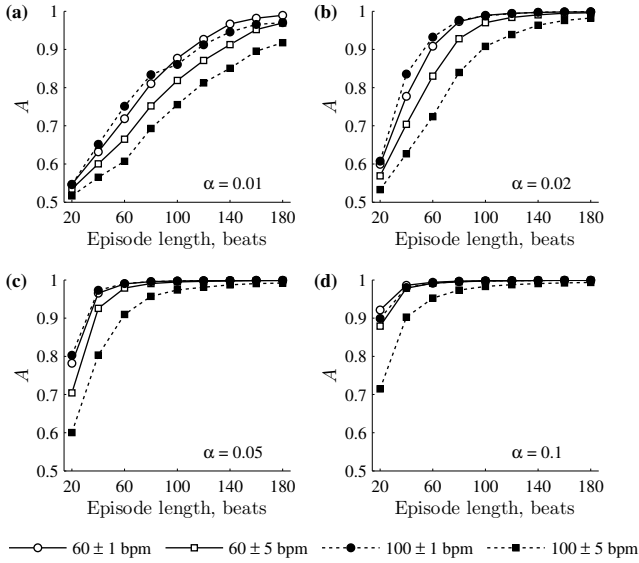
Figure 5: Detection performance A as a function of AF episode length for (a) $\alpha = 0.01$, (b) $\alpha = 0.02$, (c) $\alpha = 0.05$, and (d) $\alpha = 0.1$. These results are based on simulated signals which do not contain any arrhythmias except AF.
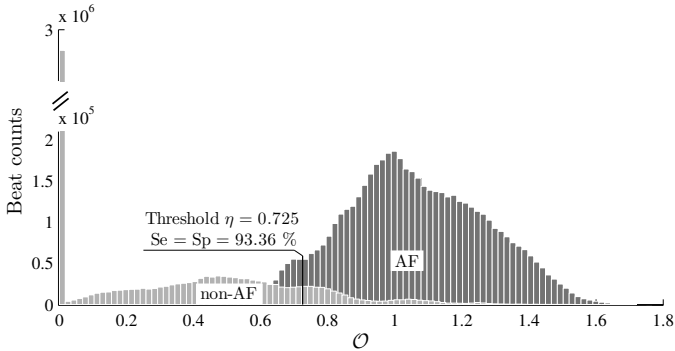

Figure 6: Distribution of the output $O(n)$ during AF and non-fibrillatory rhythms (non-AF). The large bin close to 0 for non-AF beats is due to the bigeminy suppression block. These results are based on the LTAF database.

the length was, in this particular test, incremented by one beat at a time until the episode was detected. The episode length is presented in Table 2 for different values of α—a parameter which is particularly influential on detection performance—together with the corresponding detection delay. For $\alpha = 0.02$, the shortest episode detected contained 60 beats, whereas for $\alpha = 0.05$, episodes as short as 15 beats could be detected

4.4. Detection on ECG databases

Detection performance was studied for various combinations of the MIT–BIH AF and MIT–BIH NSR databases, see Table 3. Using the offline detector with $\alpha = 0.02$, the sensitivity/specificity were found to be 97.1/98.3% on the MIT–BIH AF database, whereas the online version performed marginally worse with 96.9/98.2%. When omitting the two records with incorrect annotations, i.e., when analyzing $AFDB_2$, the sensitivity increased from 97.1% to 98.0% for the offline version. When evaluating performance on the MIT–BIH NSR database,

Table 2: The briefest possible episode detected for different values of α and the corresponding detection delay. In this particular test, an episode was considered as detected whenever the annotation and the detector output overlapped with at least 50%.

α	Episode length, beats	Delay, beats
0.005	300	318
0.01	120	158
0.02	60	78
0.05	15	30
0.1	8	14

a specificity of 98.6% was achieved for both the off- and online detectors (sensitivity was not evaluated since no AF episodes were present).

Tuning the detector to finding briefer episodes, e.g., by using $\alpha = 0.05$, Table 3 shows that such tuning comes at the expense of slightly reduced performance since the sensitivity/specificity drop from 97.1/98.3% to 96.7/97.9% on AFDB for the offline detector.

4.5. Off- and online detection

The slight difference in performance between the off- and online versions is due to the different filters used for trend estimation, defined by (2) and (7). The influence of a nonlinear phase on the RR interval trend $r_t(n)$ is illustrated in Fig. 7. Although the phase distortion is negligible when α is large, i.e., 0.1, a slight exponential reaction during rhythm transitions can be noted for $\alpha = 0.02$ or smaller, leading to slower reaction when heart rate changes occur and higher phase disturbances.


Figure 7: RR interval trend estimation using forward–backward and unidirectional filtering techniques. Note that the delay of the unidirectional filter has been compensated for.

5. Discussion

A recent study has shown that AF is often overlooked after interventional therapies when the standard strategy for treatment evaluation is used, i.e., at least two 24-h Holter recordings [23]. This finding suggest that continuous AF monitor-

Table 3: Sensitivity and specificity of the proposed detector, evaluated for various combinations of the MIT–BIH Atrial Fibrillation (AFDB) and MIT–BIH Normal Sinus Rhythm (NSRDB) databases. Records 00735 and 03665 are excluded in $AFDB_1$, and records 04936 and 05091 in $AFDB_2$.

| Database | $\alpha = 0.02$ | | | | $\alpha = 0.05$ | | | |
| | offline | | online | | offline | | online | |
	Se, %	Sp, %	Se, %	Sp, %	Se, %	Sp, %	Se, %	Sp, %
AFDB	97.1	98.3	96.9	98.2	96.7	97.9	96.5	97.9
$AFDB_1$	97.1	98.1	96.8	98.0	96.6	97.7	96.5	97.8
$AFDB_2$	98.0	98.2	97.7	98.1	97.5	97.8	97.3	97.8
AFDB & NSRDB	97.1	98.5	96.9	98.5	96.7	98.4	96.5	98.4
$AFDB_1$ & NSRDB	96.8	98.2	96.5	98.1	96.4	97.7	96.2	97.8
$AFDB_2$ & NSRDB	97.3	98.2	97.0	98.2	96.8	97.8	96.7	97.8
NSRDB	NA	98.6	NA	98.6	NA	98.6	NA	98.6

Table 4: Detector performance evaluated on the MIT–BIH Atrial Fibrillation database (AFDB).

Method	Year	Database	Se, %	Sp, %
Proposed detector	2014	AFDB	97.1	98.3
Zhou et al. [8]	2014	AFDB	96.9	98.3
Carvalho et al. [10]	2012	$AFDB_1$	93.8	96.1
Huang et al. [6]	2011	AFDB	96.1	98.1
Lake et al. [7]	2011	AFDB	91	94
Lian et al. [25]	2011	AFDB	95.8	96.4
Dash et al. [5]	2009	$AFDB_2$	94.4	95.1
Tateno & Glass [26]	2001	AFDB	94.4	97.2

ing performed during much longer time spans should be recommended so that the success rate is not overestimated. However, current technology does not allow continuous monitoring since most AF detectors require a substantial amount of computations which render efficient implementations difficult, especially for detectors which involve morphologic ECG features. A first step towards efficient implementation was recently taken in [24] where the RR-based detector in [5] was implemented in application specific integrated circuit (ASIC). The results suggest that the energy required for long-term operation, i.e., for several years, is well within the battery capacity of an existing implantable device.

The present detector is compared to the best-performing detectors in the literature with respect to sensitivity and specificity, using the MIT–BIH AF database, see Table 4. Like in most other studies, performance is presented with one decimal. It is evident that the detector by Zhou et al. [8] performs almost as good as the present detector, although the difference in performance increases slightly when the results are reported with two decimals: sensitivity/specificity are 97.12/98.28% and 96.89/98.25%, respectively.

When evaluating detector performance, it is important, for the reasons mentioned earlier, to also consider the ability to detect brief AF episodes. Of the detectors in Table 4 which employ a window length of 128 beats, i.e., [26, 5, 6, 25], it was only Lian et al. [25] who also reported on performance for shorter windows. Comparing their results obtained for a length of 32 (i.e., the shortest window studied) with the results of the present detector obtained for $\alpha = 0.05$, the sensitivity is 94.4% vs. 96.7% and the specificity 92.6% vs. 97.9%, and thus the

present detector offers a considerable improvement in performance.

Atrial fibrillatory episodes should have a duration of at least 30 s to be considered clinically significant [27], implying that detectors with a 128-beat window tend to miss brief clinical episodes. With the growing interest in detection of occult AF episodes [28, 1, 12, 13, 29], with durations from 5 to 30 s, there are even stronger reasons to develop AF detectors which can operate with window lengths much shorter than 128. Such a goal may not be easily achieved, however, when the detection process involves the estimation of probabilities, e.g., used for computing the sample entropy, since a shorter window implies increased statistical uncertainty.

Thanks to the flexibility introduced by α, the present detector can be tuned to detect brief episodes, although, just as with any detector, improved detection of brief episodes comes at the expense of lower specificity. Since the detector was developed on the LTAF database, mostly containing very long AF episodes, it is not surprising that better performance was obtained for a small α ($\leqslant 0.01$), cf. Fig. 4. However, since the detection delay becomes unacceptably large for such small values, $\alpha = 0.02$ was used as a suitable trade-off. For the online detector, a smaller α leads to increased phase distortion and, therefore, it may be motivated to consider more sophisticated approaches to trend estimation [30, 31].

A specificity of 98.6% was achieved on the MIT–BIH Normal Sinus Rhythm database, see Table 3, indicating that the detector produces few false alarms in the presence of respiratory sinus arrhythmia. This result is slightly better than those reported in [6] and [8] where the specificity was found to be 98.2% and 98.3%, respectively.

When more complex arrhythmias are encountered, reduced performance is expected. This reduction can estimated by analyzing the MIT–BIH Arrhythmia database [18, 17] which contains a variety of more complex arrhythmias. The resulting sensitivity/specificity were found to be 97.8/86.4%, indicating that most AF episodes could be reliably detected, while certain arrhythmias are misdetected as AF. Zhou et al. [8], being one of the very few authors who have reported on the performance on this database, achieved sensitivity/specificity of 97.3/90.8%, where specificity is notably better than that of the present detector. This difference in performance may be explained by the much longer detection window used by Zhou et al. which pro-
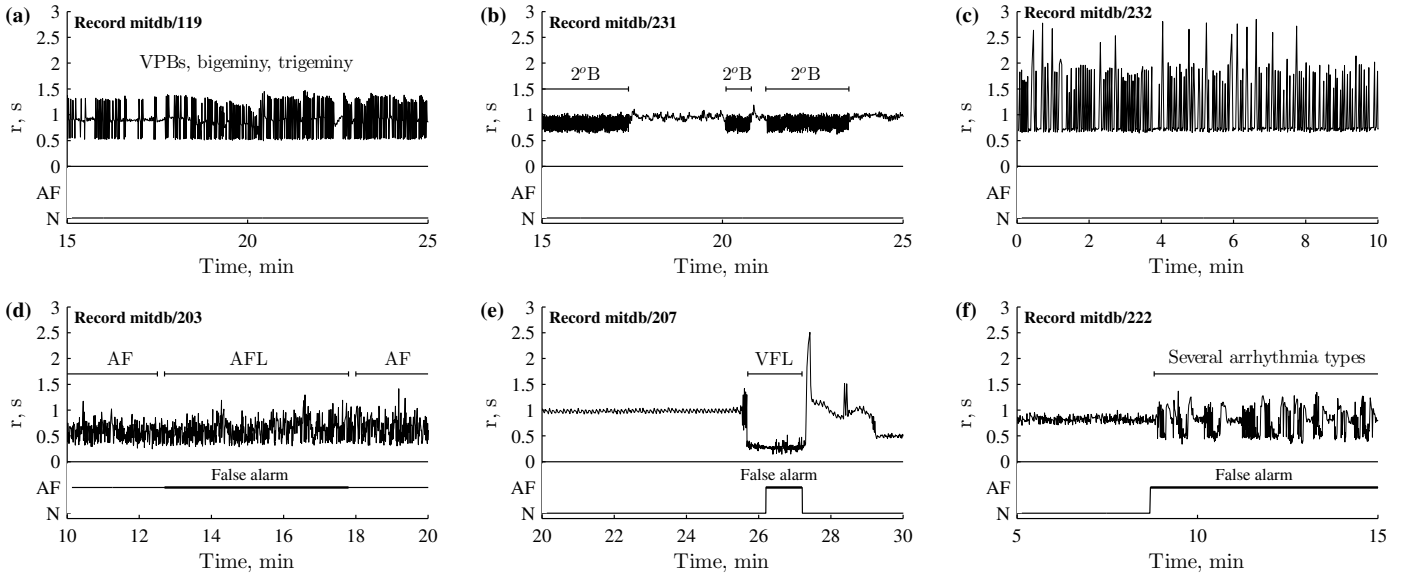
Figure 8: Examples of detector performance for various arrhythmia episodes of the MIT–BIH Arrhythmia database. The top row displays cases free of false alarms during (a) multiple ventricular premature beats (VPBs), (b) 2nd degree atrioventricular block (2°B) and (c) sinus bradycardia. The bottom row displays cases with false alarms due to (d) atrial flutter (AFL) with irregular ventricular rhythm, (e) ventricular flutter (VFL), and (f) an episode of combined arrhythmias including AF, AFL, atrial bigeminy, supraventricular tachycardia, atrioventricular junctional rhythm, and atrial premature beats.

duces better specificity when long AF episodes are encountered (as is the case in this database).

The examples in Fig. 8 shed additional light on the pros and cons of AF detection when performed in the presence of various arrhythmias, none of them being AF. In the first three examples, the detector demonstrates excellent performance when encountering frequent ectopic beats (Fig. 8(a)), episodes of 2nd degree atrioventricular block (Fig. 8(b)), and sinus bradycardia (Fig. 8(c)). On the other hand, the performance degrades when atrial flutter (Fig. 8(d)) or ventricular flutter (Fig. 8(e)) are encountered since both these types of flutter are difficult to distinguish from AF when confining the analysis to RR intervals. Figure 8(f) displays an episode of a complex supraventricular arrhythmia which causes the detector to produce a false alarm.

It is likely that these false alarms in Fig. 8(d)–(f) can be avoided by the introduction of a two-stage detection scheme. With its high sensitivity, the present detector can be employed in a first stage to find possible AF episodes. In a second stage, a more advanced AF detector is employed which involves morphologic information, like the detector described in [11], to determine in further detail whether the detected episodes are correct or not. Since morphologic analysis is computationally much more costly, the second-stage detector should preferably be implemented in a server or a smartphone.

While implementational aspects of the detector are outside the scope of this study, e.g., whether ASIC or a field programmable gate array (FPGA) should be used, a few observations on detector complexity can nonetheless be made. The present detector requires a window of only 8 RR intervals which, as already pointed out, is highly desirable from the viewpoint of detecting brief episodes. Moreover, such a short window implies very modest memory requirements which is desirable from the viewpoint of energy consumption. With its extremely simple structure, the detector requires just a few arithmetical operations (Table 1) and no table lookups and, consequently, the implementation can be made very battery-conserving. It is noted that further complexity reduction can be achieved by replacing the multiplications in (2), (3), and (7) by additions and shifts which are less costly, reducing the number of multiplications from 8 to 2. For example, the choice $\alpha = 0.02$ can be closely approximated with $5/256 \approx 0.01953$ which is implemented by 5 additions and an 8-bit shift.

Zhou et al. [8] reported on the computation time needed to analyze different public databases, but did not provide any information on the required number of arithmetic operations per RR interval. Since that detector makes use of high-order filters, in addition to the buffer required for the 128-beat window, much more memory is required than for the present detector.

The low complexity of the present detector is partly achieved by avoiding rather involved steps for handling of ectopic beats such as those in [5, 10]. Instead, simple filtering and flagging techniques, i.e., (1) and (5), are employed for the purpose of reducing the number of false alarms due to bigeminy—an approach which seems to perform well according to the results in Table 4.

A limitation of the present study is that the proposed detector was not evaluated on real ECG data with occult paroxysmal AF since no such database has yet been annotated. Instead, an approach involving test signals has been pursued which still can provide valuable insight on detection performance.

7

6. Conclusions

The present study shows that the proposed AF detector, despite its extreme simplicity, offers better performance than do the detectors described in the literature. An important feature of the detector is its use of a short window, only 8 beats, facilitating the detection of brief, subclinical AF episodes. The detector is particularly well-suited for implementation in a battery-powered device, e.g., an external or implantable event recorder, thanks to the very few arithmetical operations required for each RR interval.

Conflict of interest statement

None declared.

References

[1] A. C. Flint, N. M. Banki, X. Ren, V. A. Rao, A. S. Go, Detection of paroxysmal atrial fibrillation by 30-day event monitoring in cryptogenic ischemic stroke: The stroke and monitoring for PAF in real time (SMART) registry, Stroke 43 (10) (2012) 2788–2790.

[2] T. Hanke, E. I. Charitos, U. Stierle, A. Karluss, E. Kraatz, B. Graf, A. Hagemann, M. Misfeld, H. H. Sievers, Twenty-four–hour Holter monitor follow-up does not provide accurate heart rhythm status after surgical atrial fibrillation ablation therapy: Up to 12 months experience with a novel permanently implantable heart rhythm monitor device, Circulation 120 (11 suppl 1) (2009) S177–S184.

[3] P. Zimetbaum, Antiarrhythmic drug therapy for atrial fibrillation, Circulation 125 (2) (2012) 381–389.

[4] HRS/EHRA/ECAS Expert Consensus Statement on Catheter and Surgical Ablation of Atrial Fibrillation:, Recommendations for personnel, policy, procedures and follow-up: A report of the heart rhythm society (hrs) task force on catheter and surgical ablation of atrial fibrillation, Europace 9 (6) (2007) 335–379.

[5] S. Dash, K. Chon, S. Lu, E. Raeder, Automatic real time detection of atrial fibrillation, Ann. Biomed. Eng. 37 (2009) 1701–1709.

[6] C. Huang, S. Ye, H. Chen, D. Li, F. He, Y. Tu, A novel method for detection of the transition between atrial fibrillation and sinus rhythm, IEEE Trans. Biomed. Eng. 58 (4) (2011) 1113–1119.

[7] D. E. Lake, J. R. Moorman, Accurate estimation of entropy in very short physiological time series: The problem of atrial fibrillation detection in implanted ventricular devices, Am. J. Physiol. (Heart Circ. Physiol.) 300 (1) (2011) H319–H325.

[8] X. Zhou, H. Ding, B. Ung, E. Pickwell-MacPherson, Y. Zhang, Automatic online detection of atrial fibrillation based on symbolic dynamics and shannon entropy, BioMedical Engineering OnLine 13 (18).

[9] S. Babaeizadeh, R. E. Gregg, E. D. Helfenbein, J. M. Lindauer, S. H. Zhou, Improvements in atrial fibrillation detection for real-time monitoring, J. Electrocardiol. 42 (6) (2009) 522–526.

[10] P. Carvalho, J. Henriques, R. Couceiro, M. Harris, M. Antunes, J. Habetha, Model-based atrial fibrillation detection, in: A. Gacek, W. Pedrycz (Eds.), ECG Signal Processing, Classification and Interpretation, Springer London, 2012, pp. 99–133.

[11] A. Petrėnas, L. Sörnmo, A. Lukoševičius, V. Marozas, Detection of occult paroxysmal atrial fibrillation, Med. Biol. Eng. & Comput. (2014) (in press).

[12] R. C. S. Seet, P. A. Friedman, A. A. Rabinstein, Prolonged rhythm monitoring for the detection of occult paroxysmal atrial fibrillation in ischemic stroke of unknown cause, Circulation 124 (2011) 477–486.

[13] A. A. Rabinstein, J. E. Fugate, J. Mandrekar, J. D. Burns, R. C. Seet, S. A. Dupont, T. J. Kauffman, S. J. Asirvatham, P. A. Friedman, Paroxysmal atrial fibrillation in cryptogenic stroke: A case control study, J. Stroke Cerebrovascular Dis. 22 (8) (2013) 1405 – 1411.

[14] P. Langley, M. Dewhurst, L. D. Marco, P. Adams, F. Dewhurst, J. Mwita, R. Walker, A. Murray, Accuracy of algorithms for detection of atrial fibrillation from short duration beat interval recordings, Med. Eng. & Physics 34 (10) (2012) 1441–1447.

[15] L. Sörnmo, P. Laguna, Bioelectrical Signal Processing in Cardiac and Neurological Applications, Elsevier (Academic Press), Amsterdam, 2005.

[16] S. Petrutiu, A. V. Sahakian, S. Swiryn, Abrupt changes in fibrillatory wave characteristics at the termination of paroxysmal atrial fibrillation in humans, Europace 9 (2007) 466–470.

[17] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C. K. Peng, H. E. Stanley, PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals, Circulation 101 (2000) E215–220.

[18] G. Moody, R. Mark, A new method for detecting atrial fibrillation using R-R intervals, in: Proc. Comput. Cardiol., Vol. 10, 1983, pp. 227–230.

[19] R. Couceiro, P. Carvalho, J. Henriques, M. Antunes, M. Harris, J. Habertha, Detection of atrial fibrillation using model-based ECG analysis, in: Proc. Int. Conf. Pattern Recognition (ICPR), Vol. 19, 2008, pp. 1–5.

[20] J. Lee, B. Reyes, D. McManus, O. Mathias, K. Chon, Atrial fibrillation detection using an iPhone 4S, IEEE Trans. Biomed. Eng. 60 (1) (2013) 203–206.

[21] P. McSharry, G. Clifford, L. Tarassenko, L. Smith, A dynamical model for generating synthetic electrocardiogram signals, IEEE Trans. Biomed. Eng. 50 (3) (2003) 289–294.

[22] V. D. A. Corino, F. Sandberg, L. T. Mainardi, L. Sörnmo, An atrioventricular node model for analysis of the ventricular response during atrial fibrillation, IEEE Trans. Biomed. Eng. 58 (12) (2011) 3386–3395.

[23] E. I. Charitos, U. Stierle, P. D. Ziegler, M. Baldewig, D. R.Robinson, H. Sievers, T. Hanke, A comprehensive evaluation of rhythm monitoring strategies for the detection of atrial fibrillation recurrence: Insights from 647 continuously monitored patients and implications for monitoring after therapeutic interventions, Circulation 126 (7) (2012) 806–814.

[24] O. Andersson, K. H. Chon, L. Sörnmo, J. Neves Rodrigues, A 290 mV sub-V-T ASIC for real-time atrial fibrillation detection, IEEE Trans. Biomed. Circ. Sys. (2014) (in press).

[25] J. Lian, L. Wang, D. Muessig, A simple method to detect atrial fibrillation using RR intervals, Am. J. Cardiol. 107 (2011) 1494–1497.

[26] K. Tateno, L. Glass, Automatic detection of atrial fibrillation using the coefficient of variation and density histograms of RR and deltaRR intervals, Med. Biol. Eng. & Comput. 39 (2001) 664–671.

[27] The Task Force for the Management of Atrial Fibrillation of the European Society of Cardiology (ESC), Guidelines for the management of atrial fibrillation, Eur. Heart J. 31 (19) (2010) 2369–2429.

[28] A. H. Tayal, M. Tian, K. M. Kelly, S. C. Jones, D. G. Wright, D. Singh, J. Jarouse, J. Brillman, S. Murali, R. Gupta, Atrial fibrillation detected by mobile cardiac outpatient telemetry in cryptogenic TIA or stroke, Neurology 71 (2008) 1696–1701.

[29] A. Kishore, A. Vail, A. Majid, J. Dawson, K. R. Lees, P. J. Tyrrell, C. J. Smith, Detection of atrial fibrillation after ischemic stroke or transient ischemic attack: A systematic review and meta-analysis, Stroke 45 (2) (2014) 520–526.

[30] M. Bianchi, M. Boyle, D. Hollingsworth, A comparison of methods for trend estimation, Applied Economics Letters 6 (2) (1999) 103–109.

[31] S. Kim, K. Koh, S. Boyd, D. Gorinevsky, ℓ_1 trend filtering, SIAM Review 51 (2) (2009) 339–360.