Computers in Biology and Medicine

Heartbeat monitoring from adaptively down-sampled electrocardiogram

Luca Mesin[1]

[1] *Mathematical Biology and Physiology, Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129, Turin, Italy*

Keywords: Holter monitoring; ECG; adaptive sampling; RR rhythm; arrhythmias identification.

Corresponding author:

Luca Mesin, Ph.D.

Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino; Corso Duca degli Abruzzi 24, Torino, 10129 ITALY

Phone: 0039-011-0904085; Fax: 0039-011 5644099; e-mail : **luca.mesin@polito.it**

Abstract

Background and Objective. Heartbeats Holter monitoring is important for the detection of arrhythmias and possible anomalies, which are predictive of cardiovascular risks and infections. Reducing the number of acquired samples is useful to save energy and memory, but a proper down-sampling schedule is needed to record all useful information.

Method. An adaptive algorithm for the non-uniform down-sampling of data is used to reduce the mean sampling frequency of ECG data. The acquired data are processed to extract RR rhythm and to classify the heartbeats among a set of possible types of arrhythmias.

Results. The proposed method is tested in terms of the ability to estimate the heart rate and to classify the heartbeats from the MIT-BIH Arrhythmia data down-sampled below the Nyquist limit. The mean accuracy in identifying the heartbeats was over the 98% and the RMS error in estimating the RR time series was lower than the 1%. Variability, spectral and complexity indexes extracted from RR series were estimated with a mean error that was lower than 10%. Classification accuracy was above the 95%.

Conclusions. An adaptive method to down-sample ECG data is discussed. It can be useful to save energy and to reduce memory occupation, while still preserving important information on the heartbeats.

1. Introduction

The long-term outpatient monitoring is used for the remote surveillance of sensitive people. It allows a rapid intervention when needed, developing an individualized care [1], with positive effects on the management of clinical services and on the quality of life of patients [2]. The heart rate (HR) variability (HRV) is an important predictor of cardiovascular risk [3]. It can be monitored by Holter acquisitions, which can then be processed by a number of mathematical techniques that quantify overall fluctuations, spectral composition and degree of complexity. Using these techniques, HRV was found to be altered, for example, in the case of systemic infection [4]. Moreover, cardiac dysfunctions can affect not only the pacing, but also the beat waveforms, which could be classified to further characterize the heartbeats [5]. The automatic classification could be useful to avoid subjective and time-consuming human-based examinations and to allow for a better management of cardiac disorders in clinics.

Reducing the number of sampled data allows to save energy and memory, making the data-logger less cumbersome and increasing the time span of the monitoring [1][2][6][7]. A method for the adaptive down-sampling of biomedical data was recently proposed [8]. It

allows a real time tuning of the sampling frequency, showing to be more flexible than other adaptive algorithms. Moreover, it outperformed both uniform down-sampling and compressive sensing in terms of the recovery of high amplitude or energetic components (measured by the averaged rectified error, ARE, and the percentage root mean square difference, PRD, respectively, which are the performance indexes usually used in literature to compare different methods [9]). However, some signals contain contributions with low average amplitude or energy that provide most information. For example, the QRS complex of the electrocardiogram (ECG) provides all information to detect the heartbeats, but can have small ARE and PRD. Moreover, details of the ECG waveform can be important when there is the need of classifying different heartbeats. Thus, it is interesting to check the performance of a down-sampling algorithm also in terms of the retrieved information and not only measuring average reconstruction errors. Hence, this paper tests the ability of the adaptive algorithm proposed in [8] to record the essential ECG samples to allow HR monitoring and beat classification.

2. Methods

ECG data

The MIT-BIH Arrhythmia Database was used [10][11][12]. It consists of 48 ECG recordings with two channels sampled at 360 Hz for 30 minutes. Only the first channel was considered (i.e., in most records, a modified limb lead II, obtained by placing the electrodes on the chest). Data contain most of the energy below 45 Hz, so that they were resampled at 100 Hz, after an antialiasing lowpass filter[1]. In order to remove low frequency noise, the data were also highpass filtered[2].

Adaptive ECG down-sampling

The method proposed in [8] was used to reduce further the number of measurements, starting from data resampled at 100 Hz. The adaptive sampling is based on a prediction algorithm (using a multi-layer perceptron, MLP [13]) which predicts the subsequent sample from a set of delayed data and estimates the uncertainty of the prediction (based on 100 estimates obtained considering uniform random data chosen within a range equal to their previously estimated uncertainty [8]). The algorithm requires measuring a new sample if its prediction is

[1] Chebyshev Type II filter, cutoff at 45 Hz with no more than 1 dB of passband ripple and at least 50 dB of attenuation in the stopband, which started at 50 Hz; notice that powerline interference is at 60 Hz. To avoid phase distortion, the filter was used twice, the second round with time reversal, to get zero phase.

[2] Anti-causal (zero-phase) Chebyshev Type II filter, with stopband edge at 0.1 Hz and cutoff at 1 Hz, with no more than 1 dB of passband ripple and at least 20 dB of attenuation in the stopband.

more uncertain than a threshold selected by the user. For each patient, the optimal delayed data were selected processing the first 5 minutes of signal based on the theory of time series embedding [14]. The prediction algorithm was then trained on the first 4 minutes and validated on the following minute of signal, selecting the best topology of the MLP in terms of generalization of the predictions (refer to [8] for details and references). These preliminary offline procedures allow to adapt the algorithm to the specific data to be sampled. The method was then applied on the rest of the ECG (i.e., the last 25 minutes of ECG were used for testing).

Different levels of down-sampling were obtained by varying the uncertainty threshold. The percentage of reduction cannot be imposed precisely a-priori, as the algorithm adapts to the tested data. A preliminary test based on 10 s of data was used to select the thresholds that allowed to reduce the acquired samples of about either the 20% or the 45%. Then, five thresholds were selected linearly distributed between these two extreme values. The down-sampling was also limited not to be pushed above the 80% [8].

Heartbeat estimation

The down-sampled data were used to estimate by cubic interpolation the original ones (already down-sampled at 100 Hz). These reconstructed data were then used to estimate the heartbeats. Some data showed amplitude variations that needed to be equalized in order to get a robust estimation of the heartbeats. Moreover, both positive and negative peaks were found (the first usually corresponding to normal beats and the latter to premature ventricular contractions). The procedure described in Appendix A was found to be sufficiently accurate to detect the main positive and negative peaks, after compensating for amplitude variations.

Data analysis

The locations of the heartbeats are indicated in the database website. They were down-sampled to 100 Hz and used as reference RR series. Moreover, beats corresponding to negligible peaks in the ECG (with amplitudes about 10 times smaller than those of the R peaks) were discarded: usually they were indicated in the annotations as missed or as the result of either artifacts, or pre-excitations, or atrial fibrillation, or atrial flutter, or nodal (A-V junctional) rhythm. The adopted algorithm, even when applied to the original data, could not identify such peaks. Then, the residual beats were compared to those identified automatically from the down-sampled ECG. Each beat was classified as either a true positive, or a false positive or negative. In particular, a beat was considered missed when the minimum delay between the estimated and the correct location of the beat was larger than 120 ms[3].

[3] This large tolerance was accepted because sometimes 2 negative peaks closer than 100 ms were present and the

The statistical distributions of the RR series extracted from the original or the down-sampled data were also compared. The estimated RR series was cleaned by excluding RR intervals which were out of a range, which was adapted to the ECG of the patient by selecting it on the basis of the training data: the allowed minimum delay between subsequent peaks was the 90% of the minimum RR interval in the training set; the maximum delay was the 50% larger than the maximum RR interval in the training set. Then, the following indexes characterizing the HRV were extracted from the RR series [4]: standard deviation (stdRR), root mean square of successive difference (RMSSD), low and high frequency components (LF and HF, defined as the percentage area under the curve of the power spectral density[4], PSD, of the RR series considering the ranges [0.04-0.18] and [0.18-0.4] Hz, respectively) and sample entropy, as a measure of complexity [15]. Furthermore, the root mean square error in estimating the RR time series distribution and its PSD were computed.

Finally, using the algorithm described in Appendix B, the peaks identified from down-sampled data were classified in either of the following 8 classes: normal beat, paced, atrial premature beat, right bundle branch block beat, left bundle branch block beat, premature ventricular contraction and fusion of normal and either ventricular or paced beat.

3. Results

Figure 1 shows a representative example of application to test data from a subject. Notice that by increasing the reduction, some R peaks can be missed (as the estimated waveform is down-sampled also around the QRS complex), introducing an error in the estimation of the RR series.

Figure 2 shows the main indexes considered to characterize the RR time series extracted from the same data as in Figure 1. Both the estimations and the percentage variation with respect to the correct value (i.e., the one computed considering the RR series provided by the MIT-BIH Arrhythmia Database) are shown. Notice that the discrepancy with respect to the correct values increases with the reduction of the number of acquired samples.

Figure 3 shows a summary of the performance of the method when applied to the whole dataset. Scatter plots are considered, showing different performance indexes versus the

algorithm identified either of the 2. In such a case, the beat was not considered as missed. However, the effect of the discrepancy was assessed considering the RR series. On the other hand, when the ECG waveform showed a clear QRS complex, the maximal delay between the estimated and correct location of the peak was in the order of a few samples.

[4] The PSD was computed by the parametric Burg algorithm, using the minimum order for an autoregressive model guaranteeing to get a variance up to 5% larger than that obtained using a model of order 50.

reduction level. The accuracy of R peak identification is shown. Moreover, the indexes estimated from the RR series are displayed as percentage of their correct value. Average errors increase with the reduction level, even if there is a large variation among different data.

Figure 4 shows the cross-correlation between the waveforms related to each heartbeat detected from the down-sampled data and the original (once identified the peaks from the down-sampled data, windows from 250 ms before to 500 ms after them were considered on both the original and the down-sampled data; the maximum of the cross-correlation was considered in each window, as a measure of similarity between the waveforms). Two examples of problematic portions of data are shown in Figure 4A. Notice that sometimes the down-sampled data are not able to represent correctly the R peaks, as already shown in Figure 1C. However, most of the identified heartbeats are associated to waveforms that are highly correlated to the original ones (average median cross-correlation above 99.5%, with a decreasing trend when increasing the reduction level, Figure 4B).

Figure 5 shows a summary of the performance in the classification of different heartbeats. Notice from the specific example shown in Figure 5A (i.e., the one corresponding to the worst accurate classification) that the waveforms can also be quite different within the same class (as in the case of the premature ventricular contractions, which present two quite different recurrent shapes): in this case, searching for the different waveforms (computed by averaging only similar waves) would improve the classification accuracy. As shown in Figure 5B, the classification accuracy is very similar (around 95%) considering different levels of sample reduction (with an expected trend to decrease when the reduction is higher). The algorithm had different performances in identifying different waveforms: Figure 5C shows the percentage of correctly identified waveforms, for each of the considered classes. Furthermore, Table 1 shows the confusion matrix of the classification of different waveforms considering all the test data, with the largest reduction of samples: notice how normal and paced beats are usually correctly identified, whereas some problems are found for example with atrial premature beats and premature ventricular contractions, which are sometimes confounded as normal beats.

Figure 6 shows the distribution of mean processing time for each data considered. The results were grouped rank ordered with respect to the 5 uncertainty thresholds. There is a decreasing trend of the processing time, as the sample reduction becomes larger. The processing time is about the 10% of the duration of the epoch (indicating that a real time application would be possible), considering the specific hardware and software used: a sequential, interpreted implementation of the algorithm in Matlab was run on a single processor of a PC with

Intel(R) Core i7-2630QM, Quad-Core, clock frequency of 2 GHz, 6 GB of RAM and 64 bits operating system.

4. Discussion

The long term monitoring of the heartbeat is extensively used for self-controlling and in patients [6][7][16][17][18]. Decreasing the number of acquired samples is useful to save the energy required for the sampling and to reduce the memory storage [9][19][20]. In the case of remote monitoring in which wireless sensors sample and transmit data to a base station, a further compression of data is useful to reduce the energy spent for the transmission [21]. Moreover, quantization of the difference signal with a low number of bits and coding have been successfully applied in ECG applications [22][23] to reduce the number of transmissions (increasing further the compression ratio, defined as a measure of the reduction of the number of bits needed to represent the original signal).

4.1 Discussion of results and possible applications

The innovative algorithm proposed in [8] was used to resample the signals of the MIT-BIH Arrhythmia Database with different levels of down-sampling (thus the compression ratio is measured only in terms of the reduction of the mean sampling frequency) and its accuracy in recovering the RR series and in classifying the beat waveforms was tested.

The algorithm showed good accuracy in the identification of the heartbeats from down-sampled data. Notice for example that the mean accuracy in identifying the heartbeats was over the 98% and the RMS error in estimating the RR time series was lower than the 1%. The indexes were estimated with a mean error that is lower than about 10%, but with some variation across different data.

The classification of the heartbeats was quite accurate (above 95% in the average, even if the accuracy depends also on the type of waveforms), even considering the simple approach described in Appendix B. The accuracy of the classification was marginally affected by the level of reduction, showing that when a beat is identified, the ECG waveform is usually well estimated. This conclusion is also supported by the high correlation (above 99% in the average) between the original waveforms and those identified from down-sampled data.

Important applications could be possible in the diagnosis of cardio-arrhythmias [24]. For example, an early detection of atrial fibrillation is important to introduce therapies protecting the patient from the effects of the arrhythmias and from its progression, but it is often

prevented by its silent nature (so that about 30% of patients are not aware of a so-called asymptomatic atrial fibrillation). The long-term ECG monitoring can allow to detect a possible problem. The proposed algorithm allows to estimate a reliable RR rhythm from down-sampled data, from which a large HRV could be possibly identified, suggesting that there is the risk that atrial fibrillation occurred during the measurement. A further inspection of the ECG waveforms would then allow to diagnose the arrhythmia. After indicating the normal and arrhythmic waves in a portion of the signal, the automatic classification could be run on the whole data to get a further assessment of the patient.

4.2 Limitations and future works

The MIT-BIH Arrhythmia Database is a good benchmark to test different methods in the same conditions. However, experimental data from Holter monitoring are expected to contain more noise than that included in the considered dataset and to be affected by large motion artefacts. A preliminary test of the algorithm on noisy data was shown in [8], where the method was used to down-sample ECG recorded by a custom-made acquisition prototype [25]. The performances of the method, in terms of reconstruction errors with similar down-sampling ratio, were comparable to those discussed here. However, additional tests are required to assess the reliability of the heartbeat estimation from down-sampled noisy data. An adequate filtering approach (possibly different from the one used for the data considered in this paper) is expected to be needed to reduce noise and slow motion artefacts.

4.3 Considerations on the development of a prototype and on power saving

To test the relevance of the proposed algorithm, a Holter system embedding it should be developed. Different solutions can be investigated, requiring an optimal hardware and software co-design in order to get indeed a real time implementation and a power saving.

1. The algorithm can be run on an embedded microcontroller, consuming a power that should be less than the one saved by discarding samples.
2. The processing can be performed by a base station, wirelessly connected to the system; in this case, the portable system saves the power related to processing, but it spends the one needed for the communication.

Some preliminary considerations on power consumption can be given. First, consider an embedded solution, implemented using a low power microcontroller, e.g., MSP430 by Texas Instruments. To save energy, the system is active only for sampling and processing, otherwise it can be in stand-by (with a consumption that is about 200 times lower than when it is active).

The microcontroller switches on to enable the ADC and then it shuts off. Making the analog-to-digital conversion takes about 3.5 μs and consumes about 2640 μW [26] (with the clock at 8 MHz). When the algorithm decides to discard a measurement, the power consumption is that of the idle state, i.e., 5 μW. Considering a constant sampling rate of 100 Hz, the average consumption of the ADC is about 5.92 μW (about 5 μW due to the idle state and 0.92 μW due to the peak power consumption during the sampling). Assuming to discard about one-half of the samples, the power saved by discarding measurements is around 7.8%.

The energy required to run the algorithm depends on many factors, including the clock frequency and the number of clock cycles, which is related to the processing time. With the implementation discussed in this paper, the processing time for a sample was about 1 ms (Figure 6). Assuming that the power consumption during processing is similar to that required by the ADC, the processing time should be reduced (to be lower than the 3.5 μs required for an AD conversion) for the average power to be less than that saved by discarding a sample. As an alternative, a dual-core solution could be used, using the MSP430 for the sampling and an application-specific instruction-set processor (ASIP, with high computational performances and low power consumption) for the processing [26]. To reduce the computational cost, different optimizations could be included: a compiled routine should be used (instead of the interpreted and sequential code here considered); some parts of the algorithm could be run in parallel (e.g., the hidden neurons of the MLP and the different randomly chosen data processed to estimate the prediction uncertainty); the uncertainty could be computed using a lower number of estimates (e.g., considering only extreme values, but performances should be checked).

Consider now the alternative of devoting the data storage and processing to a base station. Power consumption was measured for a chip including the same microcontroller as before, i.e., the MSP430, and a ZigBee Network Processor, i.e., CC2530ZNP [27] (many alternatives are also available, e.g., the Bluetooth low energy, intended to provide transmission with low power consumption). The power required for waking up and switching off the CC2530, preparing the data and transmitting them is much larger than that saved by avoiding a measurement (about 7 mA·ms for conversion of 2 samples, 90 mA·ms for preparation and 100 mA·ms for transmission [27]). Thus, transmitting the data to save the power needed for the processing is not a viable solution with such a hardware. However, there could be applications in which data transmission to a base station is needed (e.g., to avoid the storage of data by the portable system and/or to allow the real-time monitoring of patients from a base

station). In those cases, the down-sampling would allow to decrease also the number of transmissions, providing a great power saving.

4.4 Final considerations

The tests conducted support the use of the adaptive down-sampling approach proposed in [8] in Holter monitoring of ECG. The mean number of acquisitions could be pushed as low as about 60 per second (far below Nyquist limit), still getting an accurate RR rhythm (with the possibility of extracting valuable information on its variability, spectral content and complexity). Moreover, reliable beats waveform representation and classification can be obtained, allowing to detect automatically possible arrhythmias or other cardiac dysfunctions. Further work is needed to test the method in the field. Potential applications in portable devices are expected, guaranteeing low power consumption (needed for long term continuous monitoring) and good accuracy (important for clinical applications or alerting systems).

5. Conclusions

This paper discusses the application to ECG of a real time algorithm that schedules a non-uniform sampling adapted to the measured data, which reduces the number of acquired samples below the Nyquist limit. The method recorded enough information to allow to estimate accurately the RR series and to classify the heartbeats among different types of arrhythmias. The technique can be applied for the individual self-assessment of healthy people or the continuous monitoring of sensitive patients, saving energy and memory, thus helping to make the data-logger less cumbersome and to increase the time span of the recording.

Appendix A – Algorithm for the estimation of the heartbeat

The estimation of the heartbeats was based on a pre-processing, needed to equalize the amplitude of the signals, followed by a classical method for peaks identification (i.e., the Pan-Tompkins algorithm [28]). Specifically, the following steps were considered for the pre-processing (which were applied to already bandpass filtered data, down-sampled to 100 Hz, as indicated in the text).

1. The data were nonlinearly processed with the function

$$y(t) = \frac{4}{\pi}\arctan\left(\frac{x(t)}{3\sigma_x}\right) \quad (1)$$

where $x(t)$ is the ECG signal and σ_x is its standard deviation. This processing allows to reduce outliers (e.g., spikes with amplitude larger than $3\sigma_x$).

2. The data were written as the difference between the positive and negative part

$$y(t) = y_+(t) - y_-(t), \quad \text{where} \quad y_+(t) = \max(y(t),0),\ y_-(t) = -\min(y(t),0) \tag{2}$$

3. The main peaks of $y_\pm(t)$ were identified by searching for the local maxima above a threshold (10% of the maximum).
4. Such peaks were then interpolated to get a signal at 100 Hz, which was then lowpass filtered (with cutoff 0.5 Hz), obtaining smooth estimations of the amplitudes of the positive and negative peaks, indicated by $A_\pm(t)$, respectively.
5. The equalized ECG was finally obtained by summing the positive and negative parts divided by the peak amplitudes

$$x_{new}(t) = \frac{y_+(t)}{A_+(t)} - \frac{y_-(t)}{A_-(t)} \tag{3}$$

The equalized ECG, $x_{new}(t)$, had positive and negative parts, each including spikes with similar amplitudes. The heartbeats were detected by processing the equalized ECG by the Pan-Tompkins algorithm [28], which is based on a derivative filter followed by a quadrature operation (that emphasizes the QRS complex), a moving average (on 15 samples) and a threshold above which the peaks are found (the standard deviation of the signal was used as threshold). The algorithm extracted both the R peaks and the most prominent negative peaks. When a negative peak was found, the algorithm searched for a positive peak in the portion of data 150 ms before: if it was not found, the negative peak was kept. Once obtained the samples corresponding to the peaks, in order to improve the resolution of their location, the signal was locally interpolated around each peak with a quadratic function, whose extreme was considered as updated estimate.

Appendix B – Algorithm for the classification of the heartbeats

Eight different classes of heartbeats were considered: normal beat, paced, atrial premature beat, right bundle branch block beat, left bundle branch block beat, premature ventricular contraction and fusion of normal and either ventricular or paced beat. Beats belonging to the different classes were identified in the training set (based on the original data resampled at 100 Hz and using the annotations provided in the dataset) and their average waveforms, RMS amplitude and mean delay from the previous heartbeat were computed. From each waveform,

the following indexes were computed: correlation C with the mean waveform and percentage error with respect to the RMS amplitude E_A and mean delay E_D

$$C = \max\left[\frac{\int w_i(t)s_k(t)dt}{\int s_k^2(t)dt}\right] \qquad E_A = \frac{|\|w_i(t)\|_2 - \|s_k(t)\|_2|}{\|w_i(t)\|_2} \qquad E_D = \frac{|RR_i^w - RR_k^s|}{RR_k^s} \tag{4}$$

where $w_i(t)$ is the waveform corresponding to the i[th] class out of the 8 listed above, $s_k(t)$ is the k[th] beat to be classified (considering a window of 750 ms around the estimated peak), RR_i^w is the mean RR interval of the i[th] class and RR_k^s is the last RR interval of the considered peak $s_k(t)$. The information provided by these 3 indexes was processed on the training set selecting the directions of maximal Fisher discrimination. The first two directions were then used on the test data to perform the classification: the maximal a-posteriori likelihood, obtained by weighting the two Fisher directions, was chosen

$$class = \underset{i}{\arg\max}\left[|\lambda_1| P_i \frac{e^{-\frac{[\vec{d}_1\cdot(\vec{x}-\vec{\mu}_i)]^2}{2\sigma_{1i}^2}}}{\sigma_{1i}} + |\lambda_2| P_i \frac{e^{-\frac{[\vec{d}_2\cdot(\vec{x}-\vec{\mu}_i)]^2}{2\sigma_{2i}^2}}}{\sigma_{2i}}\right] \tag{5}$$

where the vectors $\vec{d}_k$ (k=1, 2) are the directions of maximal Fisher discrimination, the scalar values λ_k are the eigenvalues associated to such directions, $\vec{\mu}_i$ is the mean value of the indexes in (4) for the i[th] class, σ_{ki} is the standard deviation of the indexes of the i[th] class projected along the k[th] direction and $\vec{x}$ is the vector of indexes associated to the waveform to be classified.

Acknowledgements.

The author thanks Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, for the economic support of this study.

Conflict of interest

The author states no conflict of interest.

Other declarations

There was no source of funding for this research (apart from Department internal resources).

References

[1] M. Patel, J. Wang, Applications, challenges, and prospective in emerging body area networking technologies, IEEE Wirel Commun, 17 (2010) 80–88.

[2] S. Park, S. Jayaraman, Enhancing the quality of life through wearable technology. IEEE Eng in Med and Biol. 22 (2003) 41–48.

[3] H. Tsuji, M.G. Larson, F.J. Jr Venditti, E.S. Manders, J.C. Evans, C.L. Feldman, D. Levy. Impact of reduced heart rate variability on risk for cardiac events. The Framingham Heart Study. Circulation, 94 (1996) 2850-2855.

[4] S. Ahmad, T. Ramsay, L. Huebsch, S. Flanagan, S. McDiarmid, I. Batkin, L. McIntyre, S.R. Sundaresan, D.E. Maziak, F.M. Shamji, P. Hebert, D. Fergusson, A. Tinmouth, A.J. Seely. Continuous multi-parameter heart rate variability analysis heralds onset of sepsis in adults. PLoS One. 14 (2009) 4(8):e6642.

[5] C. Ye, B.V. Kumar, M.T. Coimbra. Heartbeat classification using morphological and dynamic features of ECG signals. IEEE Trans Biomed Eng. 59 (2012) 2930-2941.

[6] B. Latrè, B. Braem, I. Moerman, C. Blondia, P. Demeester, A survey on wireless body area networks, Wirel Netw. 17 (2011) 1–18.

[7] F. Pinciroli, C. Pagliari, Understanding the evolving role of the Personal Health Record, Comput Biol Med. 59 (2015) 160-163.

[8] L. Mesin, A neural algorithm for the non-uniform and adaptive sampling of biomedical data, Comput Biol Med. 71 (2016) 223-230.

[9] D. Craven, B. McGinley, L. Kilmartin, M. Glavin, E. Jones. Compressed Sensing for Bioelectric Signals: A Review. IEEE J Biomed Health Inform. 19 (2015) 529-40.

[10] G.B. Moody, R.G. Mark. The impact of the MIT-BIH Arrhythmia Database. IEEE Eng in Med and Biol. 20 (2001) 45-50.

[11] A.L. Goldberger, L.A.N. 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-e220

[12] G.B. Moody, R.G. Mark, A.L. Goldberger. PhysioNet: a Web-based resource for the study of physiologic signals. IEEE Eng in Med and Biol 20 (2001) 70-75.

[13] S. Haykin, Neural Networks: A Comprehensive Foundation, Prentice Hall, 1999.

[14] L. Mesin, Introduction to Biomedical Signal Processing, ISBN 9788892322721, 2017.

[15] J.S. Richman, J.R. Moorman. Physiological time-series analysis using approximate entropy and sample entropy. Am J Physiol Heart Circ Physiol. 278 (2000) H2039-49.

[16] P.S. Pandian, K. Mohanavelu, K.P. Safeer, T.M. Kotresh, D.T. Shakunthala, P. Gopal, V.C. Padaki, Smart Vest: wearable multi-parameter remote physiological monitoring system. Med Eng Phys. 30 (2008) 466-77.

[17] F. Buttussi, L. Chittaro, MOPET: a context-aware and user-adaptive wearable system for fitness training. Artif Intell Med. 42 (2008) 153-163.

[18] M. Chan, D. Estève, J.Y. Fourniols, C. Escriba, E. Campo, Smart wearable systems: current status and future challenges, Artif Intell Med. 56 (2012) 137-156.

[19] A.M. Dixon, E.G. Allstot, D. Gangopadhyay, D.J. Allstot, Compressed sensing system considerations for ECG and EMG wireless biosensors, IEEE Trans Biomed Circuits Syst. 6 (2012) 156-166.

[20] L. Mesin, S. Aram, E. Pasero, A neural data-driven algorithm for smart sampling in wireless sensor networks, EURASIP J Wirel Commun and Netw, 23 (2014) 1-8.

[21] F. Dabiri, H. Noshadi, M. Sarrafzadeh, Behavioural reconfigurable and adaptive data reduction in body sensor networks, Int J Auton. and Adapt Comm Systems, 6 (2013) 207-224.

[22] L. Polania, R. Carrillo, M. Blanco-Velasco, K. Barner, Exploiting prior knowledge in compressed sensing wireless ECG systems, IEEE J.Biomed.Health Inf. 19 (2015) 508–519.

[23] D. Craven, B. McGinley,L. Kilmartin, M. Glavin, E. Jones, Energy-efficient Compressed Sensing for ambulatory ECG monitoring, Comput Biol Med. 71(2016) 1–13.

[24] A.J. Camm, G. Corbucci, L. Padeletti. Usefulness of Continuous Electrocardiographic Monitoring for Atrial Fibrillation, Am J Cardiol. 110(2012) 270–276.

[25] L. Mesin, A. Múnera, E. Pasero, A low cost ECG biometry system based on an ensemble of Support Vector Machine classifiers, in Smart Innovation Systems and Technologies, (Editors: S. Bassis, A. Esposito, C. F. Morabito, E. Pasero) Springer, (2016) 425-433.

[26] F. Bouwens, J. Huisken, H. De Groot, M. Bennebroek, A. Abbo, O. Santana, J. van Meerbergen, A. Fraboulet. A dual-core system solution for wearable health monitors. In Proceedings of the 21st Edition of the Great Lakes Symposium on VLSI, GLSVLSI '11, New York, NY, USA, (2011) 379–382.

[27] http://www.ti.com/lit/an/swra381/swra381.pdf

[28] J. Pan, W.J. Tompkins. A real-time QRS detection algorithm, IEEE Trans. Biomed Eng. 32 (2007) 230-236.

Table caption

Confusion matrix of the classification of the different waveforms (N - normal beat; P - paced; A - atrial premature beat; R - right bundle branch block beat; L - left bundle branch block beat; V - premature ventricular contraction; F - fusion of normal and ventricular beat; f - fusion of normal and paced beat), considering the largest reduction of each of the file of the MIT-BIH Arrhythmia Database.

Figure Captions

Figure 1. Example of processing of a specific signal (first derivation of data 119 of the MIT-BIH Arrhythmia Database, already down-sampled to 100 Hz, i.e., close to the Nyquist limit). A) Portion of data, with superimposed the estimation of the amplitude of positive and negative peaks used for equalization (the type of the waveforms is also indicated). B) Equalized data with indication of the R peaks. C) Down-sampled data with minimum and maximum reduction (some peaks are missed using the maximal reduction). D) Original RR time series and error in estimating it using the maximally reduced data. E) Probability density function (PDF) of the RR computed from the original and the maximally down-sampled data. F) Power spectral density (PSD) of the RR (after removing the mean) computed from the original and the maximally down-sampled data (PSD obtained by the Burg parametric method). G) Performances of R peaks identification in terms of true positive (TP), false positive (FP) and false negative (FN) considering a mistake if the delay between the correct and the estimated spikes was more than 120 ms.

Figure 2. Indexes extracted from the RR time series extracted from original and down-sampled data from a specific signal (same as for Figure 1). A) Standard deviation, B) root mean square of successive difference, C) sample entropy, D) low frequency contribution, E) high frequency contribution, F) percentage error in estimating the power spectral density (PSD) and the probability density function (PDF) of the RR.

Figure 3. Accuracy and percentage error of indexes extracted from the whole MIT-BIH Arrhythmia Database in relation with the percentage of down-sampling.

Figure 4. A) Examples of portions of original and down-sampled data, with the indication of the cross-correlation between the waveforms (identified from the down-sampled data). B) Distribution of correlations between original and down-sampled data (different percentiles are

shown) for each subject in the MIT-BIH Arrhythmia Database, in the case of maximum reduction level. Moreover, on the bottom, the average median cross-correlation is shown as a function of the reduction level.

Figure 5. Accuracy in classifying automatically the heartbeats. A) Examples of different heartbeats found in a specific file. On the left, all waveforms of a specific type are aligned (number of waveforms given in parenthesis) and averaged (lighter trace); on the right, two portions of the signal are shown indicating the different waveforms. B) Accuracy of the classification for each file (maximum reduction level is considered, but similar results are obtained for all levels) and average accuracy as a function of the reduction level. C) Percentage of correctly identified waveforms, comparing the estimates with the classification given by experts and provided in the MIT-BIH Arrhythmia Database (N - normal beat; P - paced; A - atrial premature beat; R - right bundle branch block beat; L - left bundle branch block beat; V - premature ventricular contraction; F - fusion of normal and ventricular beat; f - fusion of normal and paced beat).

Figure 6. Mean time taken to process 1 s of data. All MIT-BIH Arrhythmia Database was considered. The time to process each test file was measured (details on the implementation and on the PC used to run the algorithm are given in the text). Median, quartiles and range of processing time are shown (outliers indicated individually).

Table 1

		Estimated class							
		N	**P**	**A**	**R**	**L**	**V**	**F**	**f**
Correct classification	**N**	59998	13	256	8	0	925	312	170
	P	0	5836	0	0	0	0	0	85
	A	490	0	1564	42	0	152	5	0
	R	25	0	20	5833	0	1	0	0
	L	0	0	0	2	6781	57	0	0
	V	866	49	59	29	12	4719	99	12
	F	146	0	0	5	0	28	436	0
	f	96	51	0	0	0	1	0	655

Confusion matrix of the classification of the different waveforms (N - normal beat; P - paced; A - atrial premature beat; R - right bundle branch block beat; L - left bundle branch block beat; V - premature ventricular contraction; F - fusion of normal and ventricular beat; f - fusion of normal and paced beat), considering the largest reduction of each of the test data (25 minutes of each of the file of the MIT-BIH Arrhythmia Database).

Figure 1

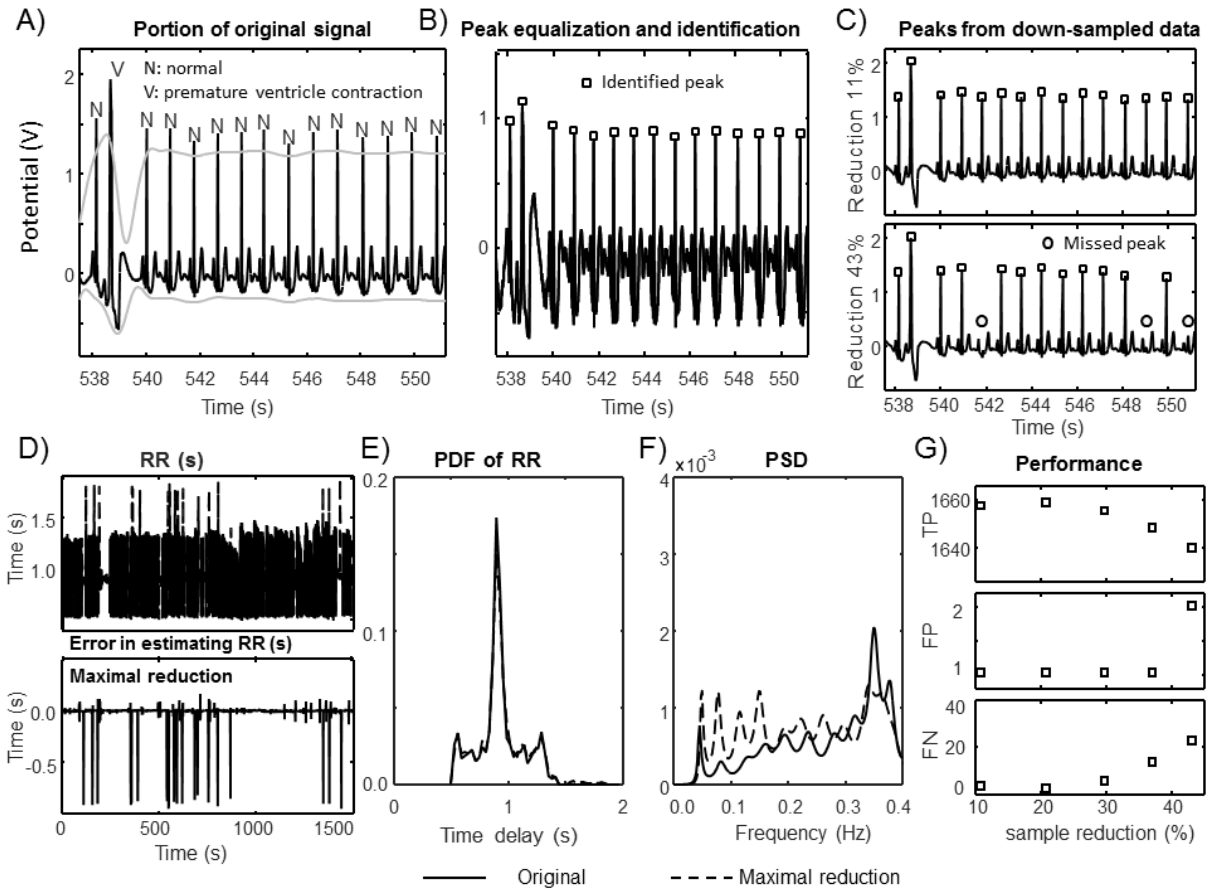

Example of processing of a specific signal (first derivation of data 119 of the MIT-BIH Arrhythmia Database, already down-sampled to 100 Hz, i.e., close to the Nyquist limit). A) Portion of data, with superimposed the estimation of the amplitude of positive and negative peaks used for equalization (the type of the waveforms is also indicated). B) Equalized data with indication of the R peaks. C) Down-sampled data with minimum and maximum reduction (some peaks are missed using the maximal reduction). D) Original RR time series and error in estimating it using the maximally reduced data. E) Probability density function (PDF) of the RR computed from the original and the maximally down-sampled data. F) Power spectral density (PSD) of the RR (after removing the mean) computed from the original and the maximally down-sampled data (PSD obtained by the Burg parametric method). G) Performances of R peaks identification in terms of true positive (TP), false positive (FP) and false negative (FN) considering a mistake if the delay between the correct and the estimated spikes was more than 120 ms.

Figure 2

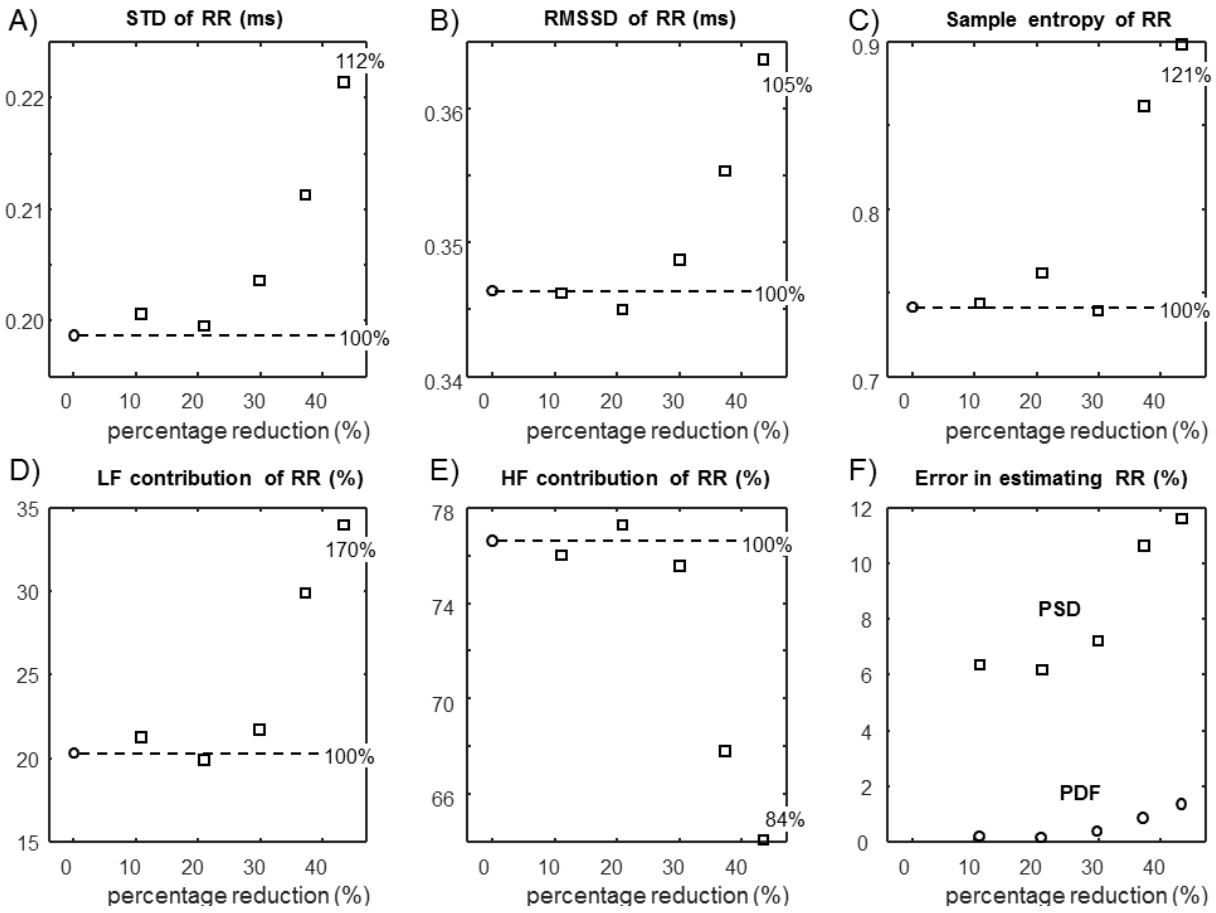

Indexes extracted from the RR time series extracted from original and down-sampled data from a specific signal (same as for Figure 1). A) Standard deviation, B) root mean square of successive difference, C) sample entropy, D) low frequency contribution, E) high frequency contribution, F) percentage error in estimating the power spectral density (PSD) and the probability density function (PDF) of the RR.

Figure 3

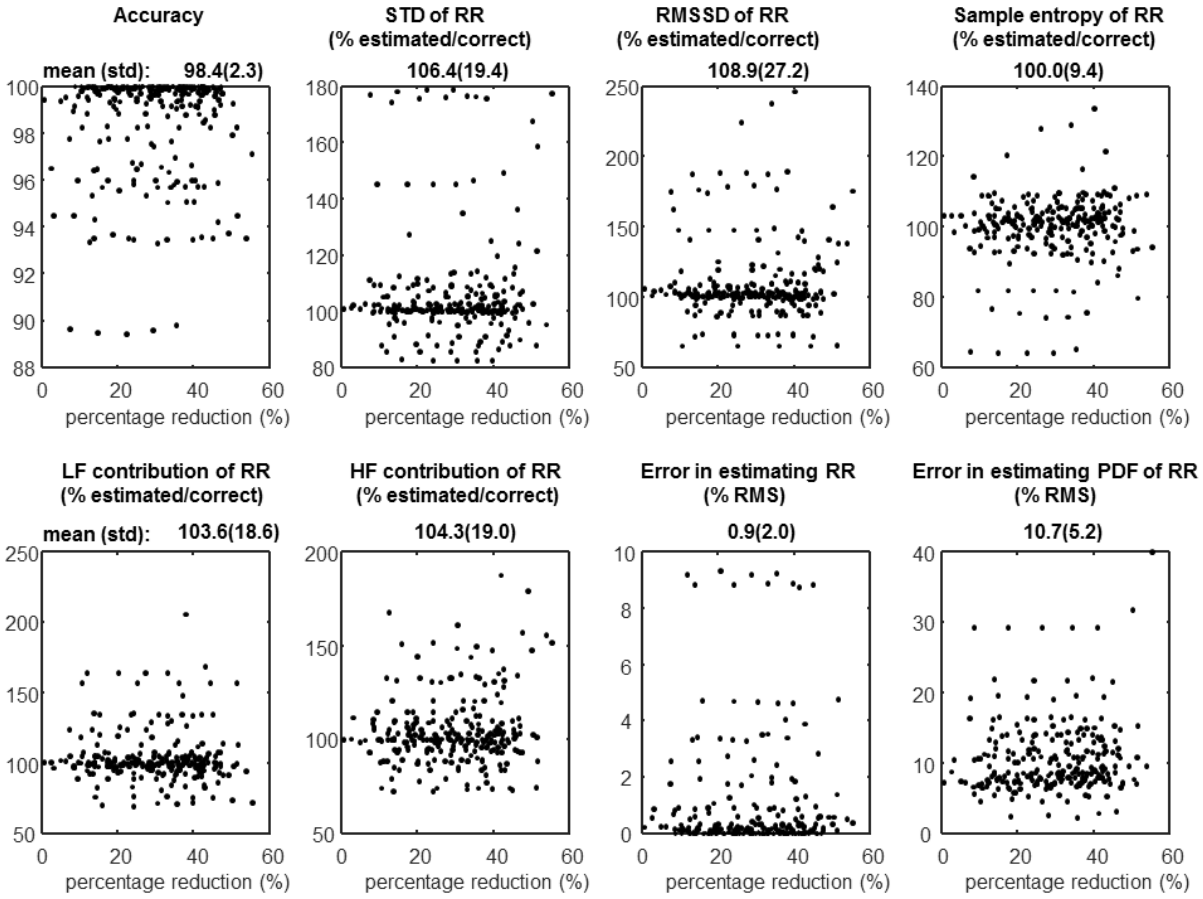

Accuracy and percentage error of indexes extracted from the whole MIT-BIH Arrhythmia Database in relation with the percentage of down-sampling.

Figure 4

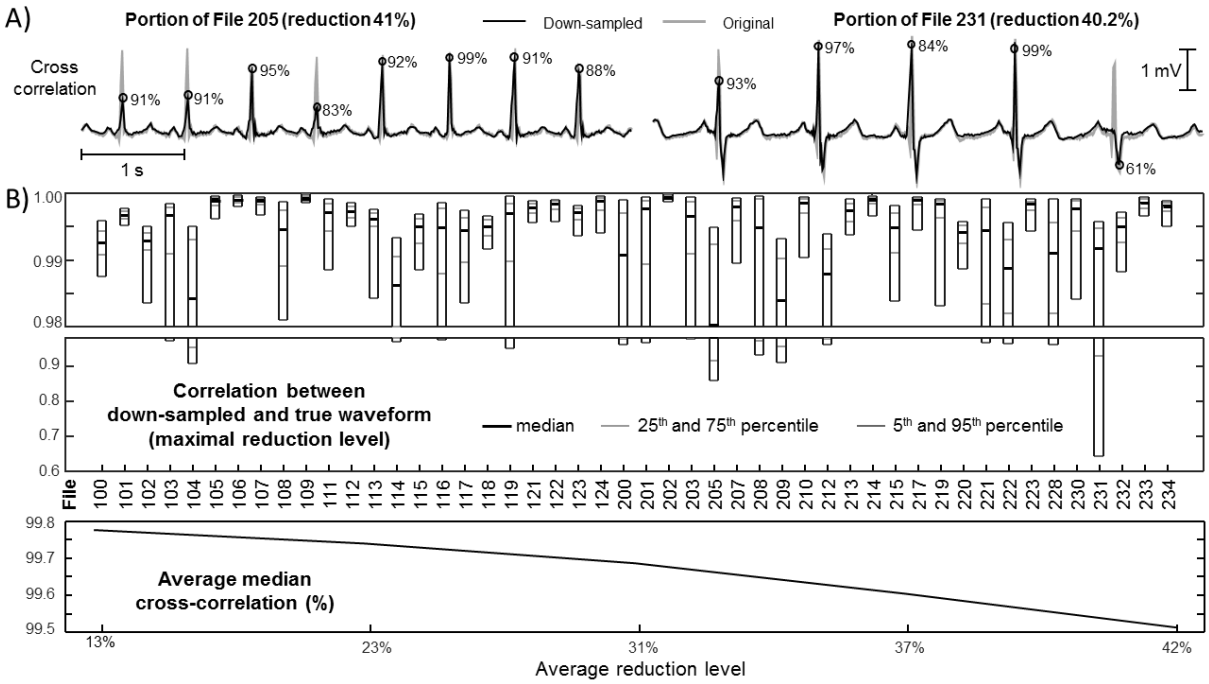

A) Examples of portions of original and down-sampled data, with the indication of the cross-correlation between the waveforms (identified from the down-sampled data). B) Distribution of correlations between original and down-sampled data (different percentiles are shown) for each subject in the MIT-BIH Arrhythmia Database, in the case of maximum reduction level. Moreover, on the bottom, the average median cross-correlation is shown as a function of the reduction level.

Figure 5

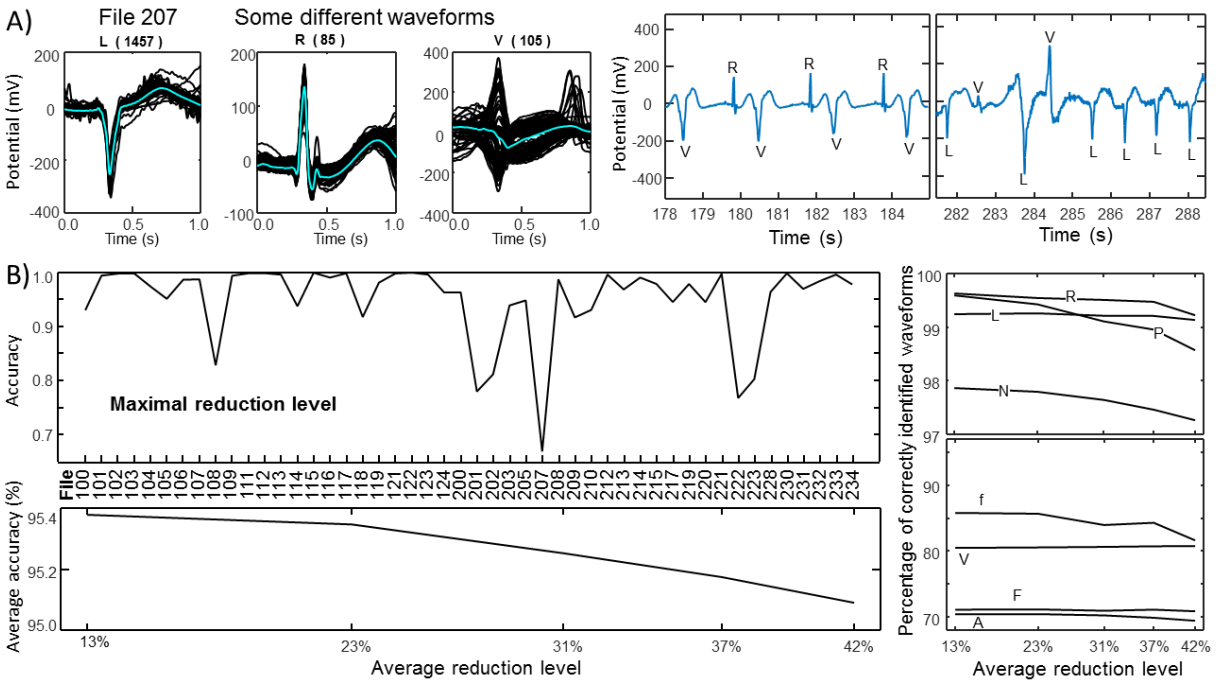

Accuracy in classifying automatically the heartbeats. A) Examples of different heartbeats found in a specific file. On the left, all waveforms of a specific type are aligned (number of waveforms given in parenthesis) and averaged (lighter trace); on the right, two portions of the signal are shown indicating the different waveforms. B) Accuracy of the classification for each file (maximum reduction level is considered, but similar results are obtained for all levels) and average accuracy as a function of the reduction level. C) Percentage of correctly identified waveforms, comparing the estimates with the classification given by experts and provided in the MIT-BIH Arrhythmia Database (N - normal beat; P - paced; A - atrial premature beat; R - right bundle branch block beat; L - left bundle branch block beat; V - premature ventricular contraction; F - fusion of normal and ventricular beat; f - fusion of normal and paced beat).

Figure 6

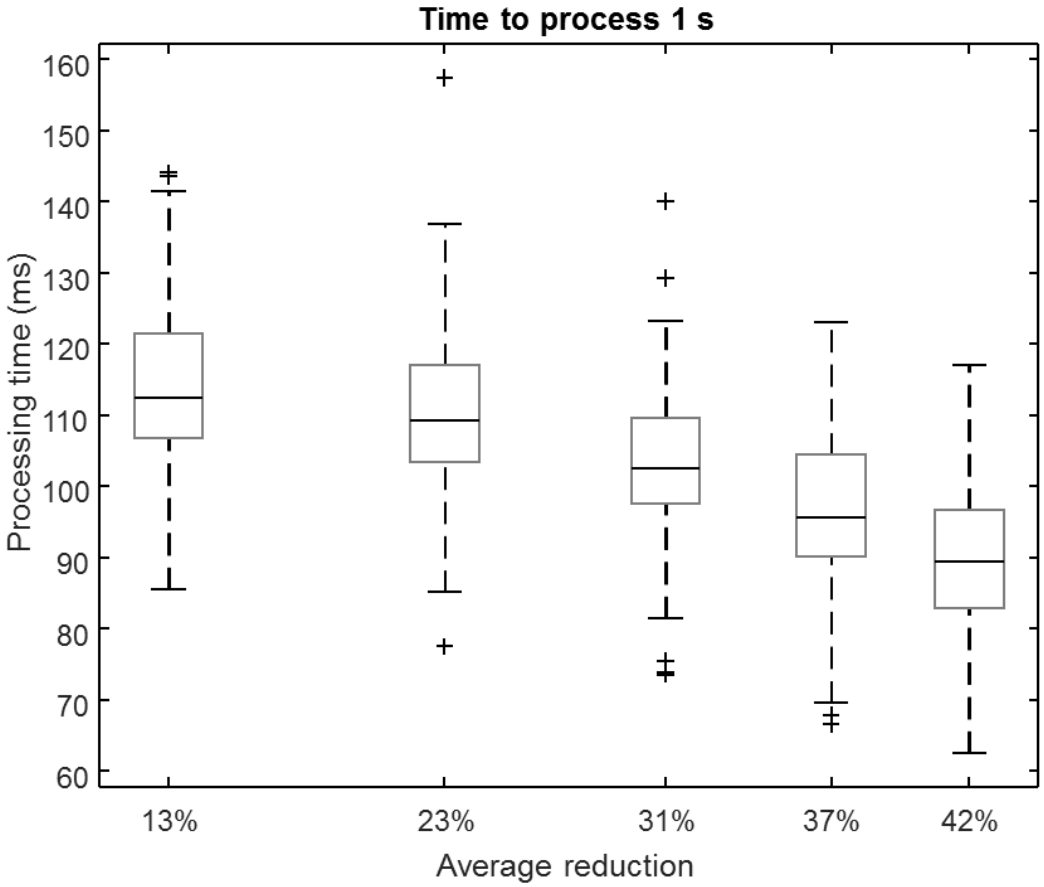

Mean time taken to process 1 s of data. All MIT-BIH Arrhythmia Database was considered. The time to process each test file was measured (details on the implementation and on the PC used to run the algorithm are given in the text). Median, quartiles and range of processing time are shown (outliers indicated individually).