Measurement 223 (2023) 113757

Contents lists available at ScienceDirect

Measurement

journal homepage: www.elsevier.com/locate/measurement

Accurate ECG monitoring by Gaussian feature streaming

Alessandra Galli, Giada Giorgi, Claudio Narduzzi *

Department of Information Engineering, University of Padua, via G. Gradenigo, 6/b, Padova, I-35131, Italy

ARTICLE INFO

Dataset link: http://physionet.org

Keywords:
Compression
Electrocardiogram
Fixed segmentation
Gaussian dictionary
Morphological feature set
Wearable monitoring

ABSTRACT

Wearable cardiac monitors can usefully contribute to early detection of potential cardiovascular pathologies. However ECG trace data streaming over wireless links creates some significant challenges, due to the amount of data to be transmitted. We employ a signal analysis approach based on a Gaussian dictionary to model ECG traces in a compressed way. The algorithm operates on fixed-length segments, and achieves effective compression for wireless data transmission, associating just 6 bytes to each Gaussian feature. At the same time it enables accurate reconstruction of ECG traces from the reduced data set. We tested our method on a set of 46 ECG recordings taken from the Physionet MIT-BIH Arrhythmia Database, obtaining 90% data compression rates, while percent relative deviation of reconstructed traces is always below 5%.

1. Introduction

Recent years have witnessed an explosive growth in the availability of wearable cardiac monitoring devices, from fitness trackers to medical grade monitors analysing multiple-lead electrocardiogram (ECG) signals [1]. On account of their reduced impact on patient quality of life and mobility, these devices can contribute to heart profile evaluation and enhance diagnostics. In particular, continuous heart monitoring plays an important role in early detection of potential cardiovascular pathologies, that are a major cause of death worldwide.

In a mobile context, data acquisition and processing devices are tasked with delivering healthcare measurement information to a data collection and analysis centre, most likely a cloud-based one [2]. For instance, Bluetooth low-energy (BLE) is often employed for short-range data transmission from sensing units to a smartphone, the latter then providing the link to a cloud-based application [3]. The combination provides good throughput, while range is dependent on wireless network coverage and transmission endurance can be limited by energy consumption. Other solutions are based instead on 5G communication technology, which offers higher bandwidth and low latency, with the advantage of direct connection to the mobile network [4].

Unobtrusive long-term cardiac monitoring requires *guaranteed throughput* to reliably transfer an accurate continuous flow of measured data [5], hence endurance and transmission range need to be emphasized. It is also useful to remember, as a reference, that one ECG channel sampled at 500 Hz with 12-bit resolution produces 6000 bits per second. This yields over 500 Mb per lead in a single day [6], creating some significant challenges to mobile applications streaming data over wireless links.

Several papers have presented effective data compression algorithms for wireless applications [7,8]. The focus is on the use of simple, low complexity algorithms to reduce the amount of transmitted data and prolong sensor battery life, while minimizing information loss [9]. However, compression is often treated as an independent problem, so that signal reconstruction is required before ECG-specific feature analysis [10].

Signals recorded during largely unrestricted daily life activities can be affected by acquisition noise, motion and electrode artifacts [11], placing greater emphasis on de-noising. In this regard, compressive sensing (CS) exploits to advantage the underlying sparse-signal assumptions [12,13]. Besides achieving useful compression ratios, CS allows detection of significant ECG features directly from compressed data, dispensing with the need to reconstruct waveforms first [6,14]. On the contrary, traditional Holter recorders *preserve* information about whole traces, enabling trained clinicians to carry out detailed analyses of waveforms where the need arises.

The signal analysis approach we present in this paper relies on a compact morphological feature set that enables accurate reconstruction of ECG traces, while providing an effective compression scheme for wireless data transmission. The focus is on reliable, continuous determination of model parameters describing an unbroken sequence of signal segments. We concentrate on showing the ability to generate accurate feature streams, regardless of the specific subject state, as a way to build confidence in the use of features. For this purpose we adopt a signal model based on Gaussian kernels, exploiting the fact that a moderate number of suitable components suffices to accurately represent each segment. This allows to meet different challenges at the same time:

* Corresponding author.
E-mail address: claudio.narduzzi@unipd.it (C. Narduzzi).

https://doi.org/10.1016/j.measurement.2023.113757
Received 6 April 2023; Received in revised form 10 October 2023; Accepted 23 October 2023
Available online 27 October 2023

1. achieve greater robustness against noise and artifacts in recorded traces;
2. support accurate and reliable trace interpretation for analysis and diagnosis;
3. provide compact representations of trace data for archiving, transmission and efficient analysis.

The paper expands on the signal analysis approach presented in [15], hence representing its extended discussion.

The quality of data for ECG analysis, reconstruction and interpretation is discussed in the final Section, for which purpose we refer to the well-known set of ECG traces contained in the MIT-BIH Arrhythmia Database, hosted at https://physionet.org [16]. These traces are extensively employed for research and represent a thoroughly studied set of measured data, carefully annotated with beat labels and revised over the years, providing a reliable reference benchmark that includes normal heart beats as well as anomalies of different kinds.

2. Gaussian modelling of ECG segments

A Gaussian kernel $g_\sigma(t)$ is defined as:

$$g_\sigma(t) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{t^2}{2\sigma^2}} \tag{1}$$

where σ is the *dispersion* (shape) parameter. On account of the kernel good localization property, a sum-of-Gaussians is an adaptable and versatile model, that suits very well the features of ECG traces. The use of Gaussian mixture models was proposed in [17] for the generation of realistic synthetic ECG traces and in [18] for ECG compression and classification. Several authors have since considered Gaussian models in ECG-related works and adapted them to different purposes.

Cardiac cycles correspond to sequences of elementary waves representing different stages of activity (P wave, QRS complex formed by Q, R and S waves, and T wave), that can be modelled by I Gaussian kernels in the form:

$$x(nT_s) = \sum_{i=1}^{I} a_i g_{\sigma_i}(nT_s - \tau_i), \tag{2}$$

where a_i is the magnitude of the ith Gaussian component and τ_i its *location*, or time position relative to a given reference point.

Each elementary wave in a cardiac cycle could be modelled by just one Gaussian kernel, but the addition of a second kernel for P and T waves is suggested in the literature to help describe asymmetries, yielding a total of *seven* Gaussian kernels, at least, to model a single heart beat [18]. In practice the location and number of wave complexes can vary within ECG physiological bounds, and may also be affected by some pathology.

An ECG trace is a sequence of individual cardiac cycles. If (2) is employed as a model for single heart beats, its application with ECG traces must rely on preliminary partitioning into shorter lengths containing a single heart beat each. Single-beat segmentation is then critically dependent on detection of peaks associated with the R waves, which requires a reliable R-peak detection algorithm running on the monitoring device. Since cardiac cycles can vary in length as a consequence of heart rate variability, ECG trace partitioning usually results in variable-length segments.

In our approach we adopt instead *fixed* segmentation, that allows to dispense with R-peak detection [19]. ECG traces are partitioned into equal-length segments, represented by a sample vector $\underline{\mathbf{x}} = [x(n_1T_s), \dots, x(n_2T_s)]^T$, where T_s is the sampling interval, $n_1T_s < 0 < n_2T_s$ with $n_2 - n_1 + 1 = N$, and superscript 'T' denotes transposition. Segment length is equal to NT_s *regardless* of the number of heart beats within. The use of a Gaussian dictionary was shown to provide promising results in the analysis of ECG traces.

Although fixed segmentation allows an easier implementation from a firmware perspective, two kinds of problems are introduced, due to the generally random position where segment edges fall within an ECG trace:

- the number of cardiac cycles within a segment is variable, depending on segment length and heart rate. An average rate may be estimated for individual subjects in a steady condition, but even this can change with subject activity;
- segment edges might fall in critical positions over significant parts of the trace, such as a QRS complex. In this case model estimation accuracy can be hampered by significant edge effects.

The first issue is addressed by considering (2) as a generic model, applied to a full fixed segment, and setting a generic upper bound I_{max} on the total number of kernels within. The second problem is harder and is in fact critical to the success of the proposed approach. It requires careful implementation of segment overlapping, and will be the focus of discussion in Section 4.

3. Dictionary-based analysis

For any given segment, parameter estimates for model (2) can be obtained by solving the least-squares problem:

$$\min_{a_i,\sigma_i,\tau_i} \sum_{n_1}^{n_2} \left[x(nT_s) - \sum_{i=1}^{I} a_i g_{\sigma_i}(nT_s - \tau_i) \right]^2. \tag{3}$$

Unfortunately, this problem is non-linear in the continuous parameters τ_i and σ_i, and is often associated with computationally-intensive algorithms, ill-suited to a wearable context [20,21].

One step towards reducing the complexity of the problem is a switch of dispersion and location parameters from continuous-valued to discrete-valued, thus creating a finite two-dimensional parameter grid. For dispersion values, the grid was designed so that when the time location of an elementary wave is matched by a dictionary element, their relative amplitude deviation never exceeds $\pm 5\%$ and is usually much less. This resulted in a non-uniform grid, increasing by 10% steps within a range determined by physiological bounds and by bandwidth limitations of the data acquisition system [22]. Sample positions are a natural choice for the set of location parameter values, the range of variation will be discussed later.

It is important to remark that by these criteria a pre-defined standard grid independent from subjective features can be defined, with the benefit of enhancing repeatability and comparability among analysed ECG traces. This leads almost naturally to approaching the problem by dictionary-based analysis. The *dictionary* can be seen as a predefined pool of Gaussian kernels, each characterized by a different combination of values for parameters τ and σ. Each dictionary element (*atom*) can be described as a column vector of N sampled values of a Gaussian kernel:

$$\underline{\mathbf{g}}_{(\tau_h,\sigma_j)} = \left[g_{\sigma_j}(n_1T_s - \tau_h), \dots, g_{\sigma_j}(n_2T_s - \tau_h) \right]^T. \tag{4}$$

where τ_h is a specific time position in a set of H allowable values and σ_j is a specific dispersion taken from a set of K possible choices.

The generic dictionary matrix $\mathbf{D}$ has the form:

$$\mathbf{D} = \left[\mathbf{D}_{\sigma_1} \; \mathbf{D}_{\sigma_2} \; \dots \; \mathbf{D}_{\sigma_J} \right] \tag{5}$$

where each $N \times H$ matrix block contains H column vectors $\underline{\mathbf{g}}_{(\tau_h,\sigma_j)}$ characterized by different time shifts, but with a common value of parameter σ_j. The size of the full dictionary $\mathbf{D}$ is $N \times M$, with $M = HJ$. Because of the way $\mathbf{D}$ is built, the indication of column index m suffices to determine values of *both* τ and σ, that is: $\underline{\mathbf{g}}_{(\tau_h,\sigma_j)} = \underline{\mathbf{d}}_m$ for some m. Since a dictionary only allows discrete sets of values for parameters σ and τ, correct design and a suitable trade-off between the size of the dictionary and waveform approximation accuracy are essential.

The analysis algorithm picks the index values of dictionary columns that best match model (2) for the given segment $\underline{\mathbf{x}}$. Indicating by S the selected set of dictionary column indices, the signal estimate is then: $\underline{\hat{\mathbf{x}}} = \mathbf{D}_S \underline{\hat{\mathbf{a}}}_S$, where $\mathbf{D}_S = \left[\underline{\mathbf{d}}_m\right]_{m \in S}$. Given $\mathbf{D}_S$, amplitude coefficients $\underline{\hat{\mathbf{a}}}_S$ are obtained by straightforward least-squares solution of (3). The

acquired ECG trace is first pre-processed to remove low frequency noise and baseline wander due to respiration and motion artifacts. These components are extracted by applying two cascaded median filters. The first filter of 200-ms width removes QRS complexes and P waves, its output is then processed by a median filter of 600-ms width to remove T waves [23]. The resulting signal contains baseline wander, that is then subtracted from the original ECG trace. The pre-processed trace is then divided into a sequence of constant-length segments for analysis.

Dictionary-based signal analysis centres on finding a sparse solution to a matrix–vector equation, formally:

$$\hat{\underline{\mathbf{a}}} = \arg\min_{\underline{\mathbf{a}}} \|\underline{\mathbf{a}}\|_0 \quad \text{subject to:} \quad \|\underline{\mathbf{x}} - \mathbf{D}\underline{\mathbf{a}}\|_2 < \epsilon \tag{6}$$

where $\underline{\mathbf{a}}$ is the vector of Gaussian kernel coefficients a_i and ϵ is a threshold value associated to the energy of the residual $\underline{\mathbf{r}} = \underline{\mathbf{x}} - \mathbf{D}\hat{\underline{\mathbf{a}}}$.

To find all the components of model (2) we apply an *orthogonal matching pursuit* (OMP) recursive greedy algorithm [24]. OMP iterations can be summarized as follows. After initializing the set of selected dictionary column indices to the empty set, $S = \varnothing$ and the signal estimate to $\hat{\underline{\mathbf{x}}} = \underline{\mathbf{0}}$, the following steps are applied iteratively:

1. compute $\underline{\mathbf{r}} = \underline{\mathbf{x}} - \hat{\underline{\mathbf{x}}}$, then select dictionary index: $m^* = \arg\max_m |\underline{\mathbf{d}}_m^T \underline{\mathbf{r}}|^2$
2. accordingly update the selected index set: $S = S \cup m^*$ and the dictionary submatrix $\mathbf{D}_S$;
3. compute a new amplitude estimate: $\hat{\underline{\mathbf{a}}}_S = \left(\mathbf{D}_S^T \mathbf{D}_S\right)^{-1} \mathbf{D}_S^T \underline{\mathbf{x}}$;
4. calculate the new signal estimate: $\hat{\underline{\mathbf{x}}} = \mathbf{D}_S \hat{\underline{\mathbf{a}}}_S$.

Iterations are stopped either when the maximum allocated number of Gaussian kernels I_{max} has been reached, or when percent relative deviation (PRD) of $\hat{\underline{\mathbf{x}}}$ from the analysed segment $\underline{\mathbf{x}}$ drops to 1% or lower. The final value, defined for the kth segment as $r(k) = \sqrt{\frac{\|\hat{\underline{\mathbf{x}}}_k - \underline{\mathbf{x}}_k\|^2}{\|\underline{\mathbf{x}}_k\|^2}}$, is computed and associated to the segment as an indication of estimation quality.

4. Fixed segmentation

Segmentation is essential to the approach presented in this paper, and although fixed segmentation may appear straightforward, it calls for great care in dealing with edge effects. In fact, it may happen that the larger elementary waves fall close to either end of the analysed segment, resulting in significant truncation of the relevant waveform.

Another issue is the choice of I_{max} for Gaussian model (2) when applied to an ECG trace segment, since in this case the number of components is unknown and the position of QRS complexes relative to segment bounds is also unknown.

The number of heart beats within a segment can vary from one segment to the next. For a given length, upper bounds can be determined based on the expected range of physiological heart rate values, and it is advisable to also allow for possible arrhythmia events. On the other hand, the number of OMP iterations per segment is upper bounded by I_{max}, that should not be set too high.

4.1. Overlap

A Gaussian model is quite adaptable and can fit a variety of waveforms, including truncated ones, which unfortunately means that the analysis algorithm is easily affected by edge effects. It will attempt to replicate truncation by concentrating as many Gaussian kernels as needed close to the edge of the analysed segment, but the resulting local overfitting may cause sub-optimal allocation and underfitting elsewhere in the segment, because of the limited number of kernels allotted for each segment.

The problem has been addressed by segment overlapping as follows:

- head and tail edge extensions to the segment are introduced to help deal with truncation effects. Accordingly, dictionary column vector length is increased to $N + 2P$, where P is the edge width and corresponds to the maximum value of the dispersion parameter σ within the dictionary;
- an overlap, the same length as the extension, is introduced between consecutive segments. As a consequence, elementary waveforms subjected to truncation are considered in both segments, but only the better fitting estimate is kept. *Effective* segment length is thus reduced to $N - 2P$.

The segment structure is shown in Fig. 1, where the cross-hatched parts at the two ends represent the edge allowance and single-hatched rectangles represent the overlap areas, that are the same width as the edges. The ratio $\eta = \frac{N-2P}{N+2P}$ can be taken as an indication of the effect of fixed-size overlapping on processing efficiency.

4.2. Location parameter range

Two different criteria were tested to determine the range of location parameter values within the Gaussian dictionary. In the first case, the range of variation for τ coincides with the segment length, that is $H = N$, and dictionary matrix size is $(N+2P)\times NJ$. Gaussian kernels located close to the bounds of segment edge extensions can be truncated, but even in this condition they are ensured to be correctly reproduced up to at least $\pm\sigma$ of their position. This is shown by the small subset of dictionary vectors, with a single value of the dispersion parameter, plotted in Fig. 2(a), where hatched and cross-hatched parts have the same meaning as in Fig. 1.

Fig. 2(b) refers to one of the ECG traces analysed for this work (trace no. 100 of the MIT-BIH Arrhythmia Database [16]). Estimated Gaussian kernel amplitudes α_i are plotted versus time location, normalized and indicated by the sample index, that is allowed to vary over the full segment length. For a segment length of 2 s this corresponds to $N = 720$ samples as shown, since the sampling rate is 360 Hz. Edge extensions and overlap areas are both $P = 48$ samples long. This trace covers a 30-minute recording time and is partitioned into 1042 overlapping 2-s segments, that are all superposed in this figure. Effective analysed segment length is actually 624 samples (1.73 s), that is, the blue part of the plot in Fig. 2(b).

Although overlap accounts for little more than 13% of the segment length, discarded coefficient estimates in the leading and trailing overlap areas, shown respectively in green and red, are about 25% of the total provided by the algorithm. This confirms the clustering effect caused by over-fitting at the segment ends and suggests that the overlap-and-discard approach described above can deal with this undesired aspect effectively. Estimates from the blue part of the segment are kept and can be streamed out for wireless transmission.

It can be noticed in Fig. 2(b) that some amplitude coefficients in the overlap areas are significantly larger, as a consequence of the attempt to fit truncated waveform segments. In a few cases, this also leads to the use of too many Gaussian kernels for this purpose, so that the bound I_{max} is reached before a satisfactory fit is achieved in the useful part of the analysed segment.

This led to the introduction of a second, different criterion for the range of τ_i, that in this case is allowed to reach out a further P points beyond the segment edge extensions on either side, that is, $H = N+4P$. In this way, a variety of truncated Gaussian kernels are introduced in the dictionary close to the edges, as illustrated in Fig. 3(a), providing scope for better fit at the segment boundaries with fewer elements. The result is shown in Fig. 3(b), where over-size values of α_i are no longer present resulting in consistently better fits, as will be shown later. Dictionary matrix is now a larger size, $(N + 2P) \times (N + 4P)J$, but it must be emphasized that with this choice any selected truncated kernel will fall in areas where it will eventually be discarded, without affecting ECG modelling accuracy.

Fig. 1. Segment structure, with edges, overlap areas and two ranges of variation for τ as discussed in Section 4.

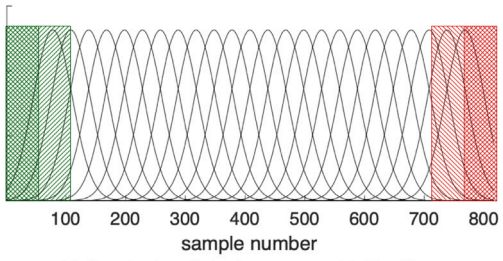

(a) Gaussian kernels for 2-s segments with $H = N$.

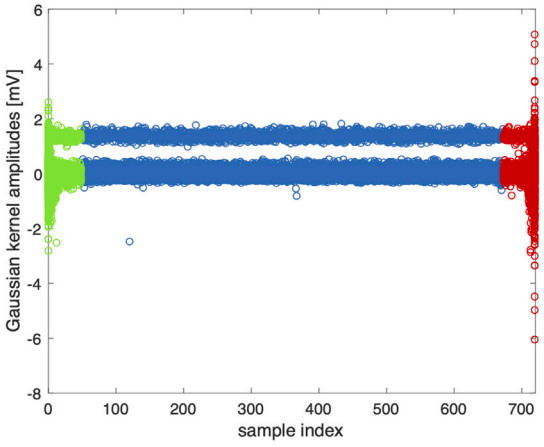

(b) Gaussian kernel amplitudes versus locations with $H = N$, dictionary size 816×32400. All segments of a single ECG trace are superposed.

Fig. 2. Location parameter τ varies over the entire length of the segment. Green (leading) and red (trailing) areas show overlaps where estimated coefficients are discarded.

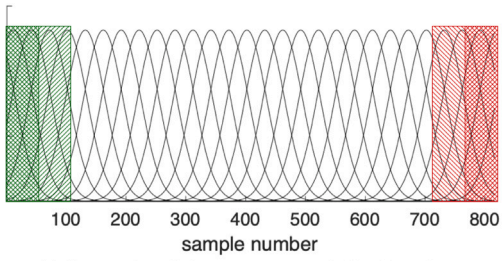

(a) Gaussian kernels for 2-s segments with $H = N + 4P$.

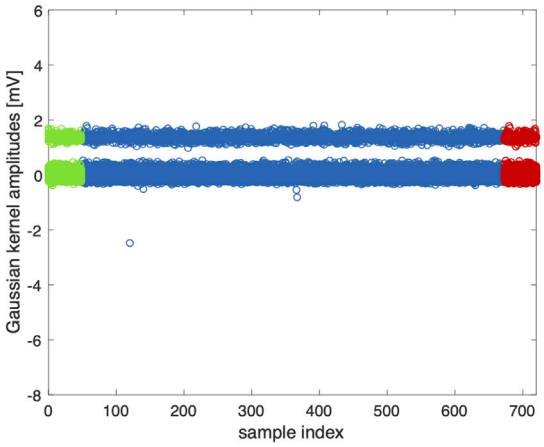

(b) Gaussian kernel amplitudes versus locations with $H = N + 4P$, dictionary size 816×41040. All segments of a single ECG trace are superposed.

Fig. 3. Range of variation of location parameter τ extends beyond the segment edge extensions.

In the following we compare 0.5-s, 1-s and 2-s segmentation, noting the terms of trade-off involved in a choice. Table 1 evidences how the two criteria discussed above and the choice of different segment lengths are related to dictionary size and to the inefficiency introduced by overlap. Common parameters are the sampling rate of 360 Hz employed for all traces considered in this work [16] and the number of values in the dispersion grid, that is set to $J = 45$.

4.3. Segment size and number of components

As already noted, fixed-size overlapping affects processing efficiency, its increasing impact indicated by a decreasing value of η as segment length N gets shorter. However, dictionary size also drops with N, relieving to some extent the computational effort. Finally, the choice of a suitable upper bound to the number of Gaussian components per segment, I_{max}, can affect both accuracy and computation.

In the following, we discuss these aspects considering a fixed value $I_{max} = 24$, that we originally set for 2-s intervals as it enables to deal with up to three QRS complexes in a segment, equating to a maximum heart rate of 90 beats-per-minute (bpm). On average, this is adequate for the kind of ECG traces we analyse in this work, but could be increased to deal with faster heart beats. By keeping I_{max} constant, we are also able to consider its relation to segment length in the analysis

Table 1
Key data for different segmentation lengths, with sampling rate: 360 Hz and $J = 45$ dispersion values.

	Segment length		
	0.5 s	1 s	2 s
Number of samples, N	180	360	720
Effective length $N - 2P$	84	264	624
Percent overlap	53.3%	26.6%	13.3%
$\eta = (N - 2P)/(N + 2P)$	0.3	0.58	0.76
Dictionary column size	276	456	816
Range for τ_i (1)	180	360	720
dictionary row size (1)	8100	16 200	32 400
Range for τ_i (2)	372	552	912
Dictionary row size (2)	16 740	24 840	41 040

Table 2
Summary of coefficient analysis on a single trace.

	Segment length		
	0.5 s	1 s	2 s
Accepted estimates	19 926	19 026	17 177
Estimates per segment	2.58	7.73	16.48
Discarded, head	26 737	6647	2688
Discarded, tail	30 869	7101	2733
Total discarded	74%	42%	24%

Table 3
Average computing results for the set of 46 traces from the MIT-BIH Arrhythmia Database.

	Segment length		
	0.5 s	1 s	2 s
Total time [s]	11 664	8408	13 550
Average per trace [s]	254	183	295
Total no. of segments	355 856	113 252	47 932
Segments per trace	7736	2462	1042

summarized in Table 2, where the same trace referred to in Fig. 3(b) is considered.

Results show that when the value of I_{max} is large in relation to segment length, further inefficiency arises from the excessive freedom allowed the OMP algorithm in pursuing the target of PRD $\leq$ 1. This is particularly noticeable from the wild overfits occurring in the edge extensions of the 0.5-s segment. The bound is obviously tighter for the longer 2-s segment, whereas for the intermediate 1-s length the total number of accepted coefficient estimates settled to around 19,000. Medical annotations for the analysed trace show it comprises a total of 2273 cardiac cycles, which means the average number of Gaussian elementary waveforms per cycle varies between 7.5 and 8.8, depending on segment length. This appears in good agreement with the literature value given in Section 2 since, in practice, it was noticed that R peak waveforms are seldom symmetric and in this case at least one more Gaussian kernel is employed to model them accurately.

5. Performance aspects

To characterize the proposed approach we considered the full set of 30-minute ECG recordings provided by the MIT-BIH Arrhythmia Database. Specifically, we selected only traces from modified limb lead II, as they were available for 46 records. Results thus refer to the analysis of over 100,000 heart beats, some of which are labelled as anomalies.

Specifically, a first group of 21 traces is representative of the variety of waveforms and artifact that an arrhythmia detector might encounter in routine clinical use. Anomalies, when present, are limited to ectopic beats, that are not necessarily related to abnormal conditions. Subjects are all adult, almost evenly divided between female and male, and under-sixty and over-sixty years olds. A second group of 25 subjects is characterized by features of the heart rhythm, QRS morphology variation, or signal quality that may be expected to present significant difficulty to arrhythmia detectors because of complex arrhythmias and conduction abnormalities [25]. For the purpose of our study, we are interested in determining whether such differences affect algorithm performances.

5.1. Computing time and memory requirements

In our trials, MatLab running on a 2,6 GHz Intel Core i7 quad-core processor was employed to process the ECG traces, while the MatLab Profiler tool was used to analyse performance. Our analysis is summarized in Table 3.

It must be remembered that each OMP iteration step involves matrix–vector products, the computation of a pseudo-inverse of progressively larger size and, above all, the search for a peak value over a vector the same size as the dictionary column. Computational cost increases with the number of components modelled by (2). Shorter segments can be processed faster, but the number of segments gets larger and, with fixed overlap length, efficiency decreases.

It can be noticed that 1-s segmentation appears to provide a good trade-off, resulting in the faster execution time, which translates into a 3-minute time for a 30-minute trace. Even allowing for a less performing processor, there is good reason to believe that real-time analysis is achievable in practice.

As far as memory requirements are concerned, dictionary matrix **D** must necessarily be stored into the edge device, accounting for most of the space. For 1-s segmentation matrix size is 456 × 24840 at most, each matrix cell being a 32-bit floating point number. Total read-only memory size is then about 45 Mbyte, which is not particularly demanding and might be compared with 500 Mbyte per channel per day for data acquisition by a standard Holter recorder.

5.2. Data rate and compression

For each Gaussian kernel modelling the trace, the algorithm provides the amplitude estimate and the corresponding dictionary column index. Each segment is in turn associated to an absolute position index within the trace. The actual position of an elementary Gaussian waveform is then the sum of the segment start index and relative position τ_i within the segment. Since amplitude accuracy is important we use for this a 32-bit floating point format, while dictionary column indices are represented by unsigned 16-bit values, yielding a total of 6 bytes per estimate.

To complete the information for a segment, a further 32 bits are needed for the position index, such large size being enough to cover over 4 months of continuous recording at the 360 Hz sampling rate considered in this paper. A single byte is employed to provide the number of kernels employed to reproduce the segment, and another byte contains the segment PRD, that gives an indication of local estimation quality, yielding an additional 6 bytes per segment.

With I_{max} = 24, the maximum data size per segment is 150 bytes, regardless of segment length. This would also be the required payload size, assuming Gaussian estimates are streamed out of the sensing unit at one packet per segment. This is mostly in excess of the average needs, that can be assessed taking the average numbers of estimates per segment, as provided in Table 2. Approximating to the larger nearest integer to determine how many bytes per segment are needed, this works out to 24 bytes for a 0.5-s segment, 54 bytes for a 1-s segment and 108 bytes for a 2-s segment. Since the aim of the approach is to achieve continuous streaming of ECG data, it can be seen that corresponding data rates are less than 500 bits per second, which is compatible with typical low-power wide-area network and narrow-band IoT 4G/5G data throughputs. If packet rates are considered instead, it should be reminded that payloads might cover multiple segments, trading data latency for efficiency.

Since heart rate and ECG trace shape can vary, an easy way to assess compression is by comparing the size of acquired traces with the corresponding sequence of parameter estimates. This shows that a 30-minute ECG trace taking about 1 Mbyte is converted into a data sequence of approximately 100 kbytes. The resulting compression ratio is around 90% and again corresponds to a data rate estimate of less than 500 bit/s.

5.3. Waveform reconstruction accuracy

If the Gaussian model that describes the ECG signal is accurate enough, the proposed algorithm enables to reconstruct a trace without introducing artifacts. Examples of reconstruction for one segment in an ECG trace are shown in Fig. 4, where it can be seen that some mild

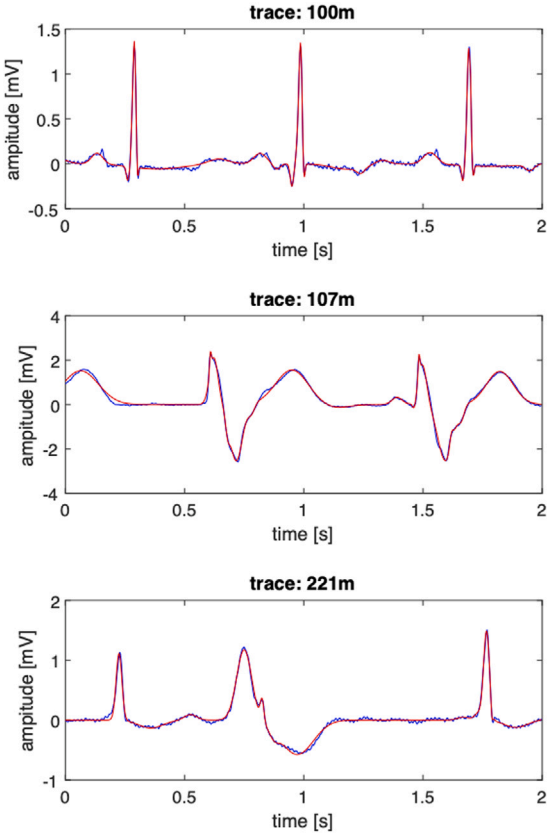

Fig. 4. Comparison between original (blue) and reconstructed (red) trace for representative 2-s segments of three ECG traces in the MIT-BIH Arrhythmia Database.

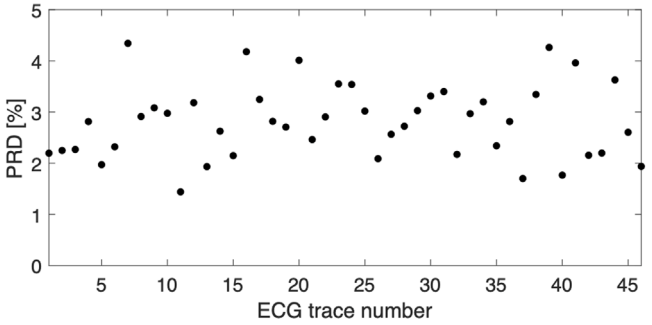

Fig. 5. Percent relative deviation (PRD) for 46 analysed traces from MIT-BIH Arrhythmia Database — 2-s overlap, dictionary size 816 × 32400.

smoothing is introduced in the reconstructed traces. A considerable variety is presented, evidencing the adaptability of the dictionary to different waveform shapes.

Percent Root mean square Difference (PRD) is an index of distortion caused by model approximation, defined as:

$$\mathrm{PRD} = 100 \cdot \sqrt{\frac{\sum \left[\hat{x}(nT_s) - x(nT_s)\right]^2}{\sum x^2(nT_s)}} \tag{7}$$

where $\hat{x}(nT_s)$ is the signal reconstructed by the proposed Gaussian model and summations extend over a whole trace. The relevance of PRD might be argued since, being defined by a summation, it tends to smooth out the effect of local discrepancies, that on the contrary might have particular significance in an ECG trace. In this case however, local accuracy for matched elementary waveforms is ensured by dictionary design, as noted in Section 3. Therefore assessment by PRD is appropriate to evidence the general quality of agreement.

PRD values are plotted in Figs. 5 and 6 for each of the reconstructed traces in the MIT-BIH Arrhythmia Database, allowing a comparison between the two approaches to segmentation presented in Section 4. It should be considered that, for trace interpretation, the assessment associated to the PRD metric is usually 0%–2% = *very good*, 2%–9% = *good* [26].

With the first approach (Fig. 5) PRD values are always below 5%, considered to correspond to *good* trace quality, even for the pathological traces contained in the database. In this case edge effects related to truncation and overfitting are not totally eliminated, contributing to a slight increase in PRD that results in a common range of values for all kinds of traces in the database.

The second approach achieves even better results although a significant increase in the size of the dictionary is required, about 50% for 1-s segmentation as can be noticed from Table 1. Accuracy gains are shown by Fig. 6, where PRD values are mostly below 2%. In this case trace quality is regarded as *very good*, and it is then possible to differentiate performance for individual traces. For this reason, on the abscissa of Fig. 6 the generic trace index is replaced by the MIT-BIH Database numerical identifier associated to each trace [25].

It can be seen that the 2% PRD threshold is exceeded only in two cases, namely trace nos. 108 and 222. Nevertheless, in the former case (a female subject of 87 years) the trace is mostly a clean signal, except some noise bursts and a few reported borderline arrhythmia effects. Trace 222 (a female subject of 84 years, with atrial flutter and fibrillation events) is also reported as mostly clean, although in this case several intervals of high-frequency noise/artifacts are reported. On the other hand, traces with significant arrhythmia events can all be reconstructed from their feature stream with PRD < 2%. It appears therefore that algorithm performances are indeed largely independent of subject conditions, and a key goal of the approach has been achieved.

6. Conclusions

In the context of this work, a cardiac monitor is seen as the edge device in a networked, possibly cloud-based healthcare system. Advanced cardiac monitors can often run rather sophisticated signal analysis algorithms, suggesting that computing power is generally available to support the monitoring framework we assume. Here, it is proposed to concentrate that power on the analysis effort, enabling the monitoring device to provide a continuous stream of model parameters, whereas diagnostic interpretation, as well as waveform reconstruction where needed, are tasked to core devices within the system.

Our model-based approach makes use of an algorithm where accuracy and computational complexity can be tuned to meet the needs of long-term ECG monitoring, overcoming issues about acquisition and wireless transmission of signals by wearable devices.

Results presented so far demonstrate that an ECG trace can be decomposed by Gaussian dictionary-based signal analysis and accurately represented by an information stream, while ECG-specific references such as R-peak position are unnecessary for trace segmentation.

A peculiar aspect of the approach is the fact that part of the measurement information is implicitly contained in the dictionary column index. In fact, each dictionary column contains the samples of the relevant Gaussian kernel, and no explicit information about location and dispersion parameters is needed at the wireless edge device that actually monitors the heart.

On the other hand, the receiving network node can reconstruct the full waveform given the amplitude estimate and the column index, using the latter as a search key to access location and dispersion parameter information. Although somewhat naïve, this mechanism can be seen as an elementary form of data protection.

The proposed signal analysis approach achieves significant data compression, allowing to send ECG trace data as a stream at around 500

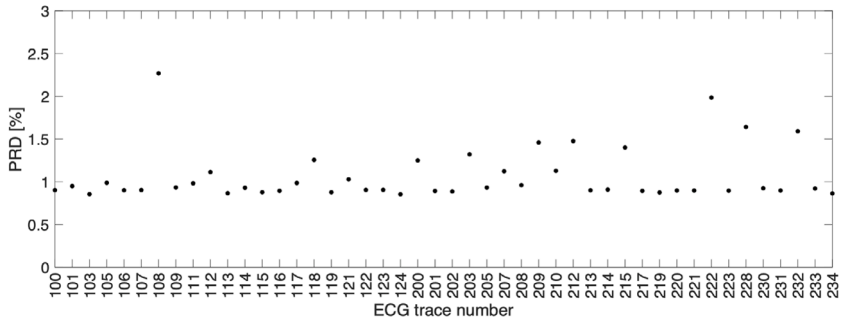

Fig. 6. Percent relative deviation (PRD) for 46 analysed traces from MIT-BIH Arrhythmia Database, 2-s overlap, dictionary size 816 × 41040.

bit/s, which is sustainable by low-power wide-area network devices in mobile applications. Gaussian kernels achieve accurate morphological modelling of ECG traces, allowing accurate waveform reconstruction. Results presented in the paper confirm that algorithm performances are mostly unaffected by any particular health condition for the range of patients considered in the database.

CRediT authorship contribution statement

Alessandra Galli: Methodology, Validation, Investigation, Writing – review & editing. **Giada Giorgi:** Methodology, Validation, Formal analysis, Writing – review & editing. **Claudio Narduzzi:** Conceptualization, Methodology, Software, Writing – original draft.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data publicly available on http://physionet.org.

References

[1] S. Yin, et al., Wearable physiological multi-vital sign monitoring system with medical standard, IEEE Sens. J. 21 (23) (2021) 27157–27167.
[2] I. Villanueva-Miranda, H. Nazeran, R. Martinek, CardiaQloud: A remote ECG monitoring system using cloud services for ehealth and mhealth applications, in: 2018 IEEE 20th Int. Conf. on e-Health Net. App. and Services, Healthcom, 2018, pp. 1–6.
[3] G. Giorgi, A. Galli, C. Narduzzi, Smartphone-based IOT systems for personal health monitoring, IEEE Instrum. Meas. Mag. 23 (4) (2020) 41–47.
[4] J. Lloret, L. Parra, M. Taha, J. Tomás, An architecture and protocol for smart continuous ehealth monitoring using 5G, Comput. Netw. 129 (2017) 340–351.
[5] H. Ali, H.H. Naing, R. Yaqub, An IoT assisted real-time high CMRR wireless ambulatory ECG monitoring system with arrhythmia detection, Electronics 10 (16) (2021) 1871.
[6] M. Abdelazez, S. Rajan, A.D. Chan, Detection of abnormal heartbeats in compressed electrocardiograms, in: 2018 IEEE Int. Sym. on Medical Measurements and Applications, MeMeA, 2018, pp. 1–6.
[7] A.B. Jambek, N.A. Khairi, Performance comparison of Huffman and Lempel–Ziv Welch data compression for wireless sensor node application, Am. J. Appl. Sci. 11 (1) (2014) 119–126.
[8] C.K. Jha, M.H. Kolekar, Electrocardiogram Data Compression Techniques for Cardiac Healthcare Systems: A Methodological Review, IRBM, 2021.
[9] G. Giorgi, A combined approach for real-time data compression in wireless body sensor networks, IEEE Sens. J. 17 (18) (2017) 6129–6135.
[10] M. Elgendi, A. Mohamed, R. Ward, Efficient ECG compression and QRS detection for E-health applications, Sci. Rep. 7 (1) (2017) 1–16.
[11] A. Servati, et al., Novel flexible wearable sensor materials and signal processing for vital sign and human activity monitoring, Sensors 17 (7) (2017) 1622.
[12] A. Galli, G. Frigo, G. Giorgi, Robust ECG denoising for ehealth applications, in: 2018 IEEE Int. Sym. on Medical Measurements and Applications, MeMeA, 2018, pp. 1–6.
[13] A. Galli, et al., Denoising ECG signal by CSTFM algorithm: Monitoring during motorbike and car races, IEEE Trans. Instrum. Meas. 68 (7) (2019) 2433–2441.
[14] A. Galli, C. Narduzzi, G. Giorgi, ECG monitoring and anomaly detection based on compressed measurements, in: Proc. of the 2018 3rd Int. Conf. on Biomed. Imag. Sig. Proc., 2018, pp. 10–17.
[15] A. Galli, G. Giorgi, C. Narduzzi, Gaussian-based analysis for accurate compressed ECG trace streaming, in: Proc. of 25th IMEKO TC4 International Symposium, 2022.
[16] G.B. Moody, R.G. Mark, The impact of the MIT-BIH arrhythmia database, IEEE Eng. Med. Biol. 20 (3) (2001) 45–50.
[17] P.E. McSharry, et al., A dynamical model for generating synthetic electrocardiogram signals, IEEE Trans. Biomed. Eng. 50 (3) (2003) 289–294.
[18] G.D. Clifford, et al., Model-based filtering, compression and classification of the ECG, Int. J. Bioelectromagn. 7 (1) (2005) 158–161.
[19] G. Da Poian, R. Bernardini, R. Rinaldo, Gaussian dictionary for compressive sensing of the ECG signal, in: Proc. of 2014 IEEE Workshop on Biometric Measurements and Systems for Security and Medical Applications, BIOMS, 2014, pp. 80–85.
[20] O. Sayadi, M.B. Shamsollahi, A model-based Bayesian framework for ECG beat segmentation, Physiol. Meas. 30 (3) (2009) 335–352.
[21] E. Kheirati Roonizi, R. Sassi, A signal decomposition model-based Bayesian framework for ECG components separation, IEEE Trans. Signal Proc. 64 (3) (2016) 665–674.
[22] A. Galli, G. Giorgi, C. Narduzzi, Standardized Gaussian dictionary for ECG analysis – A metrological approach, IEEE Open J. Instrum. Meas. 1 (2022) 4000209, http://dx.doi.org/10.1109/OJIM.2022.3196703, 1–9.
[23] P. De Chazal, et al., Automated processing of the single-lead electrocardiogram for the detection of obstructive sleep apnoea, IEEE Trans. Biomed. Eng. 50 (6) (2003) 686–696.
[24] J.A. Tropp, A.C. Gilbert, Signal recovery from random measurements via orthogonal matching pursuit, IEEE Trans. Inform. Theory 53 (12) (2007) 4655–4666.
[25] MIT-BIH arrhythmia database directory, 2023, (https://physionet.org/physiobank/database/html/mitdbdir/mitdbdir.htm#toc) – Hypertext edition, (Accessed 21 July 2023).
[26] Y. Zigel, A. Cohen, A. Katz, The weighted diagnostic distortion (WDD) measure for ECG signal compression, IEEE Trans. Biomed. Eng. 47 (11) (2000) 1422–1430.