

RESEARCH

Open Access

Real time reconstruction of quasiperiodic multi parameter physiological signals

Gartheeban Ganeshapillai* and John Guttag

Abstract

A modern intensive care unit (ICU) has automated analysis systems that depend on continuous uninterrupted real time monitoring of physiological signals such as electrocardiogram (ECG), arterial blood pressure (ABP), and photo-plethysmogram (PPG). These signals are often corrupted by noise, artifacts, and missing data. We present an automated learning framework for real time reconstruction of corrupted multi-parameter nonstationary quasiperiodic physiological signals. The key idea is to learn a patient-specific model of the relationships between signals, and then reconstruct corrupted segments using the information available in correlated signals. We evaluated our method on MIT-BIH arrhythmia data, a two-channel ECG dataset with many clinically significant arrhythmias, and on the CinC challenge 2010 data, a multi-parameter dataset containing ECG, ABP, and PPG. For each, we evaluated both the residual distance between the original signals and the reconstructed signals, and the performance of a heartbeat classifier on a reconstructed ECG signal. At an SNR of 0 dB, the average residual distance on the CinC data was roughly 3% of the energy in the signal, and on the arrhythmia database it was roughly 16%. The difference is attributable to the large amount of diversity in the arrhythmia database. Remarkably, despite the relatively high residual difference, the classification accuracy on the arrhythmia database was still 98%, indicating that our method restored the physiologically important aspects of the signal.

Introduction


A modern intensive care unit (ICU) employs multiple bedside monitors to track the state of patients. Automated analysis systems are used to analyze the signals collected by these monitors. Many of these systems run algorithms that are dependent on the morphologies of a signal or set of signals. For instance QRS detectors are used to identify beat boundaries [1], beat classifiers are used to identify ectopic beats [2,3], and morphological variability computes the morphological variations between successive heart beats [4]. These systems critically depend on continuous uninterrupted real time monitoring of the nonstationary physiological signals such as electrocardiogram (ECG), arterial blood pressure (ABP), and photo-plethysmogram (PPG). Unfortunately, the utility of these systems is compromised by the fact that the signals are often severely corrupted by noise, artifacts, and missing data.

In this article, we address the problem of identifying and reconstructing corrupted regions in a multi parameter quasiperiodic signal in a clinically useful way (Figure 1). Here, a multi parameter signal is a set of multiple time-aligned signals obtained from multiple sensor sources. They could be of the same type, e.g., two channels of an ECG, or of different signal types, e.g., ECG channels, ABP, and PPG signals. Our technique is based upon learning a patient-specific model of the relationships between time-aligned signals and then reconstructing corrupted segments using the information available in correlated signals.

This study provides a layer between the signal acquisition system and the automated analysis system. As such it is complementary to any signal enhancement (e.g., removal of baseline wander for ECG signals [1]) done during signal acquisition and any techniques for coping with low SNR in the analysis algorithms.

The examples in this article are drawn largely from signals related to cardiovascular activity. These signals are generated by the same underlying system (the heart). Therefore, the periodicity of these signals is related to

*Correspondence: garthee@mit.edu
Massachusetts Institute of Technology, CSAIL, 77 Massachusetts Avenue,
Cambridge, MA, USA



heart rate. Further, these signals also share the same repetitive structure, and hence, they are jointly quasiperiodic.

The types of noise artifacts and corruptions that these signals suffer include, missing data, physical activities, muscle artifacts (MA), electro-magnetic interference (EM), and baseline wander (BW) [1,5]. One or more signal can be corrupted at any time, but we consider the situation where at least one signal is uncorrupted. This is plausible because, unlike the signals, which are generally correlated as they share the same source (heart), corruptions have different sources and characteristics. For instance, electro-magnetic interference affects ECG signals but not the PPG signal.

In experiments carried out on the MIT-BIH Arrhythmia Database, a two-channel database with many clinically significant arrhythmias, our method improved a standard algorithm [6] for detecting a kind of ectopic beat by more than $7\times$ on a signal corrupted with white Gaussian noise, and increased the similarity to the original signal, as measured by the normalized residual distance, by more than $2.5\times$.

As discussed in Section ‘Offline learning’, the problem of reconstructing a corrupted multi-parameter physiological signal was formally posed in the 11th annual PhysioNet/CinC challenge [7]. On this CinC dataset, our method performed well relative to the winning entry. The methods used by the respondents to the challenge are substantially different from those described in this article. In particular, none of them attempted to segment the quasi-periodic signals before reconstructing them. This impaired their performance on correlated quasi-periodic signals (such as ECG and ABP), but allowed them to process independent signals (e.g., respiration).

Recent attempts to handle the corruptions by noise, artifact, and missing data in physiological signals focus on using redundant measurements, and fusing data from multiple sensors. Researchers have developed methods to robustly estimate heart rate (HR) by fusing information from multiple signals [8]. In addition to various ECG channels, researchers also make use of ABP and PPG signals for HR estimation [9]. In this context, researchers have proposed segmentation methods for ABP and PPG

that involve detecting the parts of the waveforms corresponding to the onset of a pulse [10,11]. These methods independently segment the physiological signals. Hence, they don’t preserve the alignment between the segments across different signals.

The organization of this article is as follows. In Section ‘Method’, we present our method and provide the mathematical framework of our study. In Section ‘Results’, we discuss the measures of performance used to evaluate our method, and present the results of a series of tests in which comparisons are made using each of the performance measures. Finally, in Section ‘Conclusion’, we summarize our study.

Method

We use a two-phase process designed to reconstruct corrupted signals in a way that improves the reliability of the automated systems that analyze these signals (Figure 2).

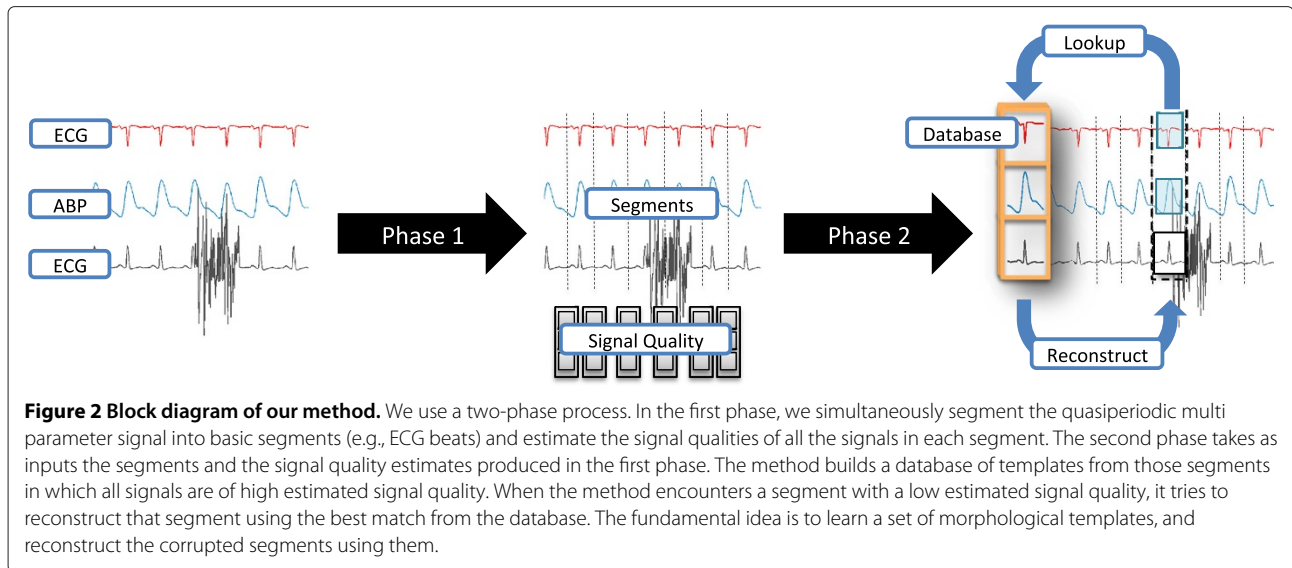
We use anonymized public medical data in our study.

Segmentation

The segmentation stage decomposes a continuous physiological signal into intervals with clinically relevant morphologies. We consider a multi parameter signal represented by a matrix \mathbf{S} , where each column represents an individual channel of the signal (e.g., an ECG channel, ABP or PPG) and each row represents a point in time. For simplicity, we assume that all the channels are sampled at the same rate so that the matrix \mathbf{S} has a repetitive structure that is shared by all the channels in the structure.

Physiological signals are typically segmented according to some well-defined notion. For example, QRS detectors that identify the R-peaks are typically used to segment ECG signals [12]. However, when there is significant corruption and noise in the signal, these methods fail to perform adequately. Use of correlated signals has been explored by researchers in the context of heart rate estimation, where a coarse estimation of segment lengths would suffice [8].

We jointly segment the multi parameter signal using template matching. The goal of the template-matching-based joint segmentation is to divide the multi parameter

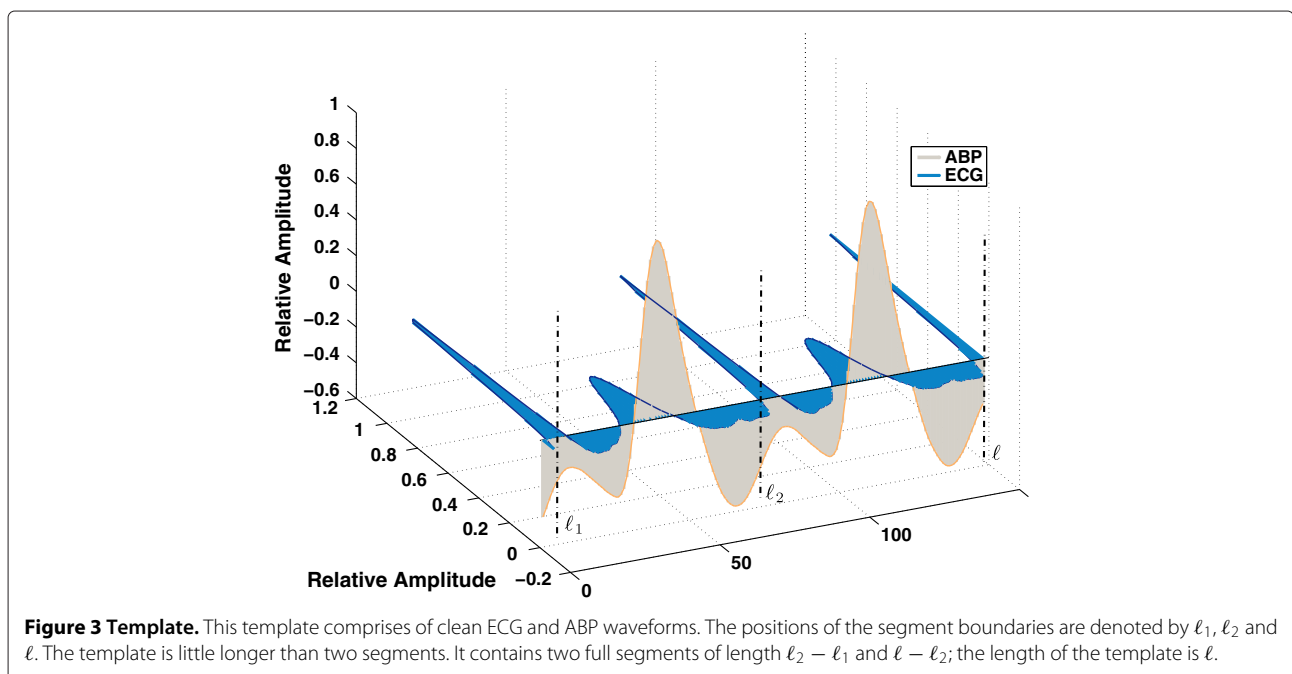


signal into quasiperiodic units (segments). The approach is different from typical ECG segmentation methods in two ways. First, it doesn't use any specialized knowledge of the characteristics of the signals. This allows the extension of the method to other types of signals (ABP and PPG). Second, a multi parameter signal is jointly segmented. Therefore, when one or more signals in the multi parameter signal are corrupted, uncorrupted signals are weighted more, and the joint segmentation is highly influenced by the uncorrupted signals.

The method uses an evolving template (Figure 3), a short multi parameter signal, and matches it with a

sliding window of the multi parameter signal. The method requires an initial template that can be provided manually or generated from the signal itself provided that the signal is uncorrupted at the beginning. In our experiments, we use the first two segments of the signal as identified by a widely used QRS detector [1]. Then the template is regularly updated to reflect the time evolution of the signal.

The algorithm continuously extracts a non-overlapping window from the signal, and identifies the segment boundary in the window by finding the prefix of the window that most closely matches the template.



The matching is done using weighted time warping (WTW) [13], which minimizes the weighted morphological dissimilarity across all the signals. The warped distance between two one-dimensional time series gives the morphological dissimilarity. Time warping is necessary, because the template and current window can be of different lengths. The weights represent the estimated quality of each signal, which is again computed by the morphological similarity between the signal in the window and its counterpart in the template.

To follow the gradual changes that are common in physiological signals, the template is updated regularly.

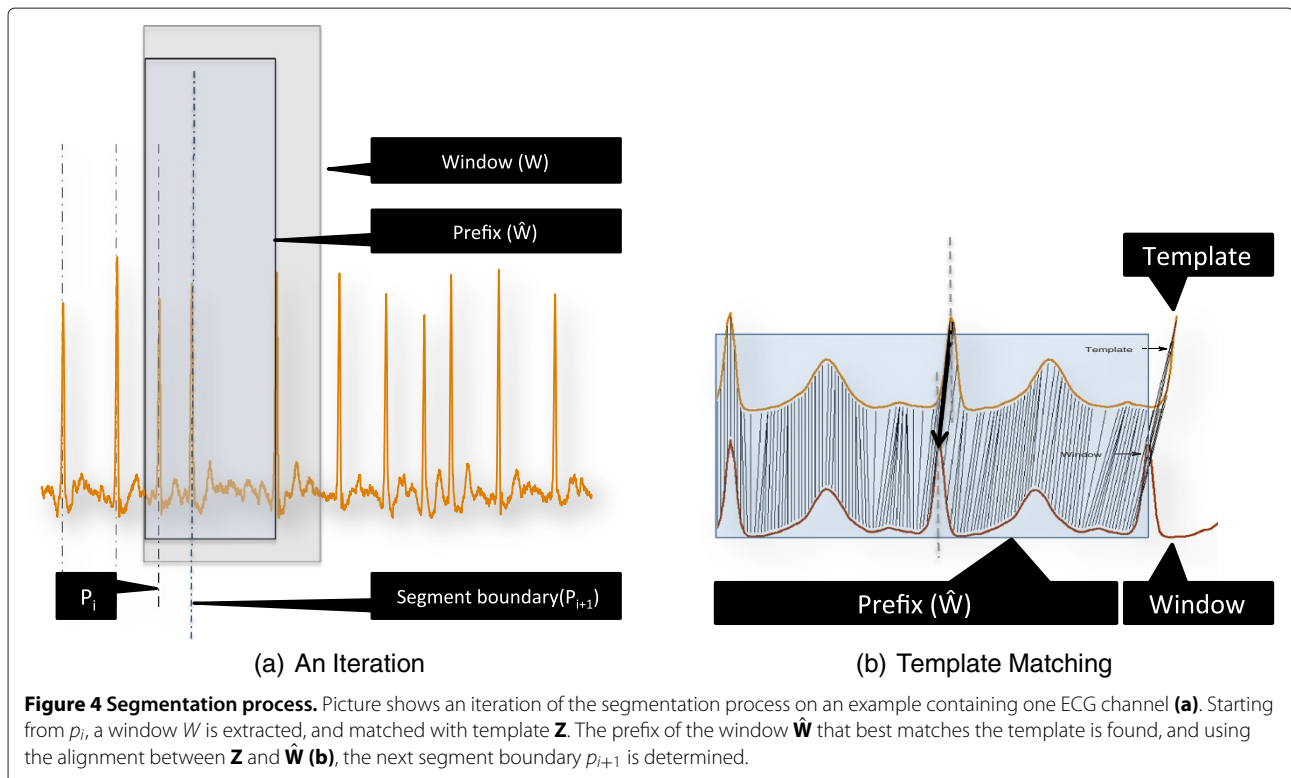
Weighted time warping based template matching

Let $\mathbf{S} \in \mathbb{R}^{n \times m}$ be a multi parameter time series consisting of m physiological signals, and $\mathbf{Z} \in \mathbb{R}^{\ell \times m}$ be the initial template. The goal is to segment \mathbf{S} into a set of segments $\mathbf{Y} = \{\mathbf{Y}_i\}$ where $\mathbf{Y}_i \stackrel{\text{def}}{=} \mathbf{S}_{[p_i, p_{i+1}]}$, and segment corresponds to a single heart beat. Here, $\mathbf{S}_{[p_i, p_j]}$ denotes the window in the target sequence \mathbf{S} from time $t = p_i$ to $t = p_j - 1$, where p_i and p_j are the row indices of the matrix \mathbf{S} .

As addressed later, we require the template (Figure 3) to be comprised of at least two segments to improve the accuracy of the matching. The template is used to find the segment \mathbf{Y}_i in \mathbf{Y} . We also assume that we know the locations of the segment boundaries $\ell_1, \ell_2, \dots, \ell$ in the template.

We start the process at some arbitrary point in time p_{start} on the signal that is to be segmented. This need not be an actual segment boundary. We run the algorithm starting at p_{start} , continuously segment \mathbf{S} , and add the segments to \mathbf{Y} . We also update \mathbf{Z} to reflect the evolution of the time series.

An iteration of our method on a single channel ECG signal is illustrated in Figure 4. We start each iteration with the extraction of a window $\mathbf{W} = \mathbf{S}_{[p_i, p_i+\nu]}$ from \mathbf{S} at p_i . Here, the window length is given by $\nu = \ell + e$, where ℓ is the length of the template and e is cushion length. We add a cushion length because the signal being processed may contain segments longer than those of the template, and we wish to ensure that the window obtained from the signal can be matched to the two segments that comprise the template (see Section ‘Window length’). In the following discussion, we use j to index individual signals in the multi parameter signal \mathbf{S} . The window is then detrended and normalized so that the corresponding signals have the same peak-peak distances, and the morphological quality estimates $\{q_j\}_m$ (Section ‘Signal quality estimation’) are computed, where q_j represents the morphological similarity between the channel W_j from \mathbf{W} and the corresponding channel Z_j from the template. For each channel j , a pairwise Euclidean distance matrix \mathbf{pD}^j is calculated between Z_j and W_j (Equations 2 and 1). The final distance matrix \mathbf{D} is obtained by weighting pairwise distance matrix \mathbf{pD}^j with q_j (Equation 3).



$$\mathbf{pD}^j = \begin{bmatrix} c_{1,1} & c_{1,2} & \cdots & c_{1,|Z_j|} \\ c_{2,1} & c_{2,2} & \cdots & c_{2,|Z_j|} \\ \cdots & \cdots & c_{x,y} & \cdots \\ c_{|W_j|,1} & \cdots & \cdots & c_{|W_j|,|Z_j|} \end{bmatrix} \quad (1)$$

$$c_{x,y} = (W_{x,j} - Z_{y,j})^2 \quad (2)$$

$$\mathbf{D} = \sqrt{\sum_{j=1}^m q_j \mathbf{pD}^j} \quad (3)$$

In Equation 3, we extend dynamic time warping [14] to higher dimensions by introducing a weighted norm over individual signals to vary the influence exerted by each signal.

For one-dimensional sequences $A = [a_1, a_2, \dots, a_x, \dots, a_n]$ and $B = [b_1, b_2, \dots, b_y, \dots, b_n]$, we use a recursion (Equation 4) with local continuity and global path constraints (proposed as Type III and Type IV local continuity constraints in [14]) to calculate the accumulated distance between subsequences $A = [a_1, a_2, \dots, a_x]$ and $B = [b_1, b_2, \dots, b_y]$. These constraints ensure that there are no long horizontal or vertical paths in the matrix \mathbf{aD} along the alignment. This prevents physiologically implausible alignments such as a long subsequence of a sequence being matched with a significantly shorter subsequence of another sequence.

$$\begin{aligned} \mathbf{aD}(A_x, B_y) &= \mathbf{D}(A_x, B_y) + \min\{\mathbf{aD}(A_{x-1}, B_{y-1}), \\ &\mathbf{aD}(A_{x-1}, B_{y-2}) + \mathbf{aD}(A_x, B_{y-1}), \mathbf{aD}(A_{x-1}, B_{y-3}) \\ &\quad + \mathbf{aD}(A_x, B_{y-1}) + \mathbf{aD}(A_x, B_{y-2}), \\ &\mathbf{aD}(A_{x-2}, B_{y-1}) + \mathbf{aD}(A_{x-1}, B_y), \mathbf{aD}(A_{x-3}, B_{y-1}) \\ &\quad + \mathbf{aD}(A_{x-1}, B_y) + \mathbf{aD}(A_{x-2}, B_y)\} \end{aligned} \quad (4)$$

Next, we trim \mathbf{W} to obtain the prefix $\hat{\mathbf{W}} = \mathbf{W}_{[1, \ell_w]}$ so that it contains only the portion of the signal that matches the template, where the prefix length ℓ_w is given by Equation 5.

$$\ell_w = \arg \min_{k_1} \frac{1}{k_1} \mathbf{aD}(k, |\mathbf{Z}|); \quad 1 \leq k_1 \leq |\mathbf{W}| \quad (5)$$

If the prefix $\hat{\mathbf{W}}$ doesn't contain two segments, the cushion length e is increased and the process is repeated. More detailed explanation is provided in Section 'Window length'.

Next, we find the optimal path alignment $(\phi(k) = (\phi_w(k), \phi_z(k)); 1 \leq k \leq K)$ between \mathbf{Z} and $\hat{\mathbf{W}}$ in the accumulated distance matrix \mathbf{aD} . Here, K is the length of the alignment [15]. From the alignment, we obtain the point f in $\hat{\mathbf{W}}$ that is matched with the segment boundary ℓ_2 in \mathbf{Z} .

$$f = \phi_w(k); \text{ where } \phi_z(k) = \ell_2 \quad (6)$$

This corresponds to the segment boundary we are interested in, because ℓ_2 marks the end of the first segment in the template.

Using a two segments long template, we align the ends of the second segment, backtrack to the end of the first segment in that alignment, and find the matching points for these segments. Using both sides of the segment boundaries increases the accuracy of the segmentation process (Figure 4(b)).

Then, using the corresponding length f , we update p_{i+1} to $p_i + f$. Finally $\mathbf{S}_{[p_i, p_{i+1}]}$ is added to \mathbf{Y} . Following the template update, the process is repeated to find the next segment boundary.

Template updating

To follow the gradual changes that are common in the physiological signals, we update the template regularly (Figure 5). We do this when the variation of the segment lengths in the neighborhood is small, and the quality estimates of all the channels are above a threshold. We require the difference between the maximum segment length and the minimum segment length to be less than 25% of the mean segment length over a moving time period of 60 seconds. We also require the minimum signal quality of any channel in the segment to be greater than a threshold.

We update the template by averaging the excerpt of the last two segments with the current template time-warped to match the excerpt length (Equation 8). The time warping is necessary, since the current template \mathbf{Z} and the excerpt of the last two segments \mathbf{Z}' can be of different lengths. New template \mathbf{Z}'' is given by,

$$\mathbf{Z}' = \mathbf{S}_{[(p_{i-1}-\epsilon), p_{i+1}]}; \epsilon \geq 1 \quad (7)$$

$$\mathbf{Z}'' = (1 - \eta) \text{warp}(\mathbf{Z}) + \eta \mathbf{Z}' \quad (8)$$

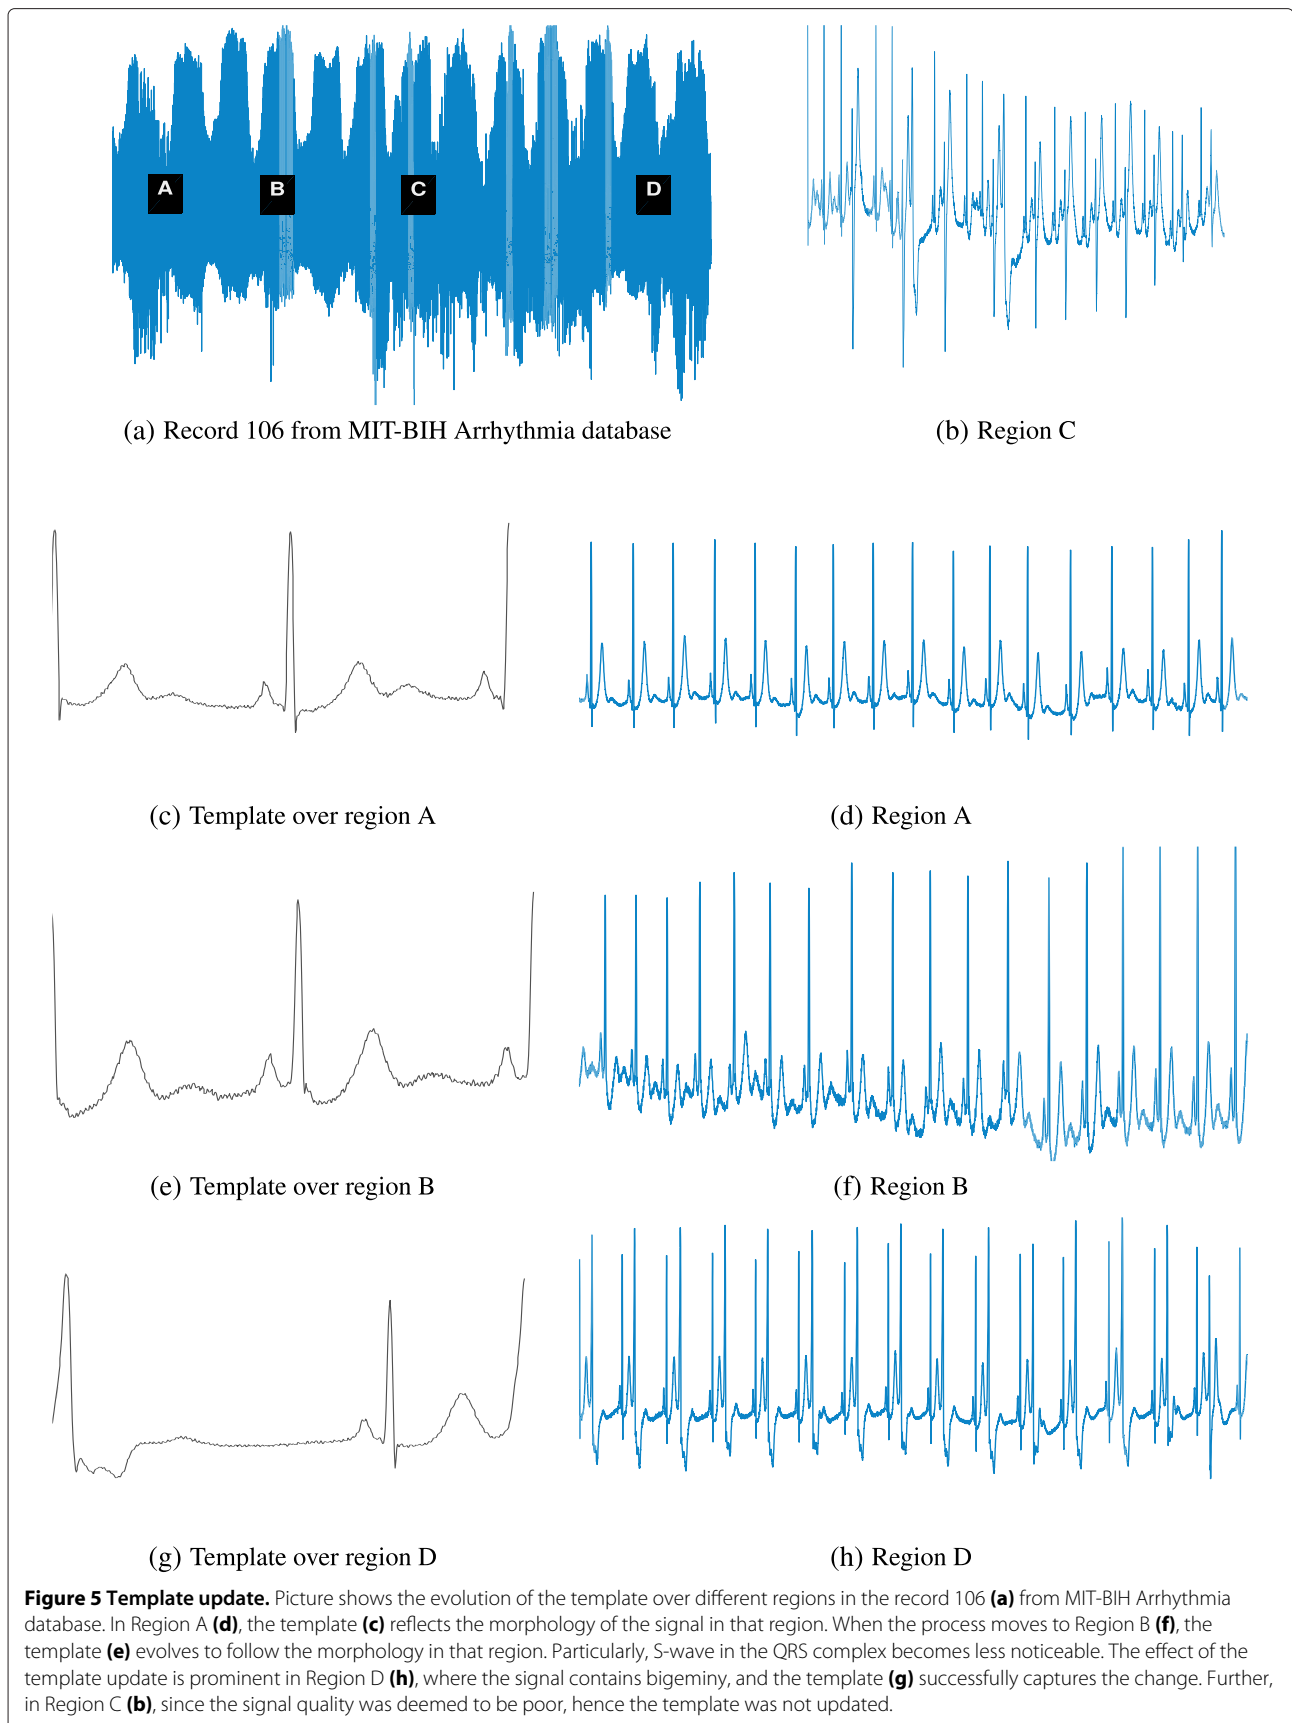
Here, ϵ ensures that the template consists of at least two segments. We vary the influence of the recent segments on the template through the constant $0 \leq \eta \leq 1$, where $\eta = 0$ implies a static template. We chose $\eta = 0.05$, but the results are not sensitive to the choice of η .

Window length

The template \mathbf{Z} contains two segments. We expect the window \mathbf{W} to contain at least two segments so that the prefix of the window would match the template. When the cushion length e is not long enough, the window $\mathbf{W} = \mathbf{S}_{[p_i, p_i+\nu]}$; $\nu = \ell + e$ would not contain two full segments. In such cases, we repeat the matching with a longer window.

Signal quality estimation

Our framework requires the assessment of the signal quality at multiple stages. In joint segmentation (Section 'Segmentation'), we weigh each signal's influence on the joint segmentation by the estimated quality of the signal. In reconstruction, we use signal quality estimates



to identify the corrupted segments. Further, in Section ‘Ensembles’, we use the signal quality of the matches to fuse the matches for reconstruction.

If we had a model of the source, the signal quality assessment would be straightforward, the signal to noise ratio (SNR) gives an estimate of the signal quality. However, for physiological signals, the source characteristics vary significantly over the time. Therefore, researchers use indirect measures to estimate the signal quality of the physiological signals. These methods assume that some characteristics of the signals are known *a priori*, and the noise artifacts cause any deviation present in the signal.

There is a lot of literature on signal quality estimation of physiological signals [8,16-18]. ECGSQI estimates the signal quality of a multi-channel ECG signal [8]. ABPSQI [8] combines wSQI and jSQI to obtain the ABP signal quality. Deshmane estimates PPG signal quality using Hjorth [19] parameters [10]. Each of these methods generates independent quality estimations on different kinds of signals using signal specific characteristics, and hence the indices are not comparable with each other.

To be useful in our framework, we need the signal quality estimate to have the following properties.

- Independent of the beat type and the location of the corruption within the segment,
- Comparable across different type of signals, such as ECG, ABP, and PPG signals, and
- Estimated at the granularity of a heartbeat.

In our method, we hypothesize that the corrupted regions of the signals will be morphologically dissimilar from clean signals, as represented by our evolving template.

We use longest common subsequences (LCSS) to estimate the morphological similarity. LCSS is a variation of edit distance that is extended to real sequences [20]. Given two time series, A and B , LCSS maximizes the similarity between A and B by finding the longest subsequences from both sequences whose elements fall within δ_x and δ_y of each other in X-axis and Y-axis. The following dynamic programming function on sub-sequences $A_x = a_1 a_2 \dots a_x$ and $B_y = b_1 b_2 \dots b_y$ is used to achieve this [20].

$$LCSS(A_x, B_y) = \begin{cases} 0 & \text{if } A \text{ or } B \text{ is empty} \\ 1 + LCSS(A_{x-1}, B_{y-1}) & \text{if } |a_x - b_y| < \delta_y \text{ and} \\ \max\{LCSS(A_{x-1}, B_y), LCSS(A_x, B_{y-1})\} & |x - y| < \delta_x \\ LCSS(A_x, B_{y-1}) & \text{otherwise} \end{cases} \quad (9)$$

Threshold δ_x controls the time warping permitted and δ_y bounds the pair wise differences of two similar elements. LCSS uses a bandwidth to estimate the similarity. This makes LCSS resilient against noise [20].

For the multi parameter signal \mathbf{S} and the template \mathbf{Z} , the signal quality of the k th signal in segment i is calculated from the morphological similarity between the window $A = S_{[p_{i-1}, p_{i+1}]}^k$, and the corresponding signal in the template $B = Z^k$. The similarity $L = LCSS(A, B) / \min(n, m)$ gives the signal quality estimate q^k of the k th signal in segment i . Here n and m are the lengths of the sequences A and B respectively. Figure 6 shows the signal qualities estimated by our method on an ECG signal.

To estimate the signal quality, we compare a window of the signal against the template. This would be valid if any difference between the template and the window were due to the noise or corruption present in the signal. However, the signal morphology evolves along the time. For instance, heart rate may increase, the amplitude of the ABP signal might drop, etc. We regularly update the template to accommodate the time evolution of the signal.

Comparison

Related methods such as ECGSQI and ABPSQI estimate the quality of the signal over a 10 s long region [8]. In contrast, the morphological similarity metric estimates the quality of the signal beat by beat. The ability to obtain a more fine-grained quality estimate helps us handle the bordering regions (start and end of the corrupted region) better.

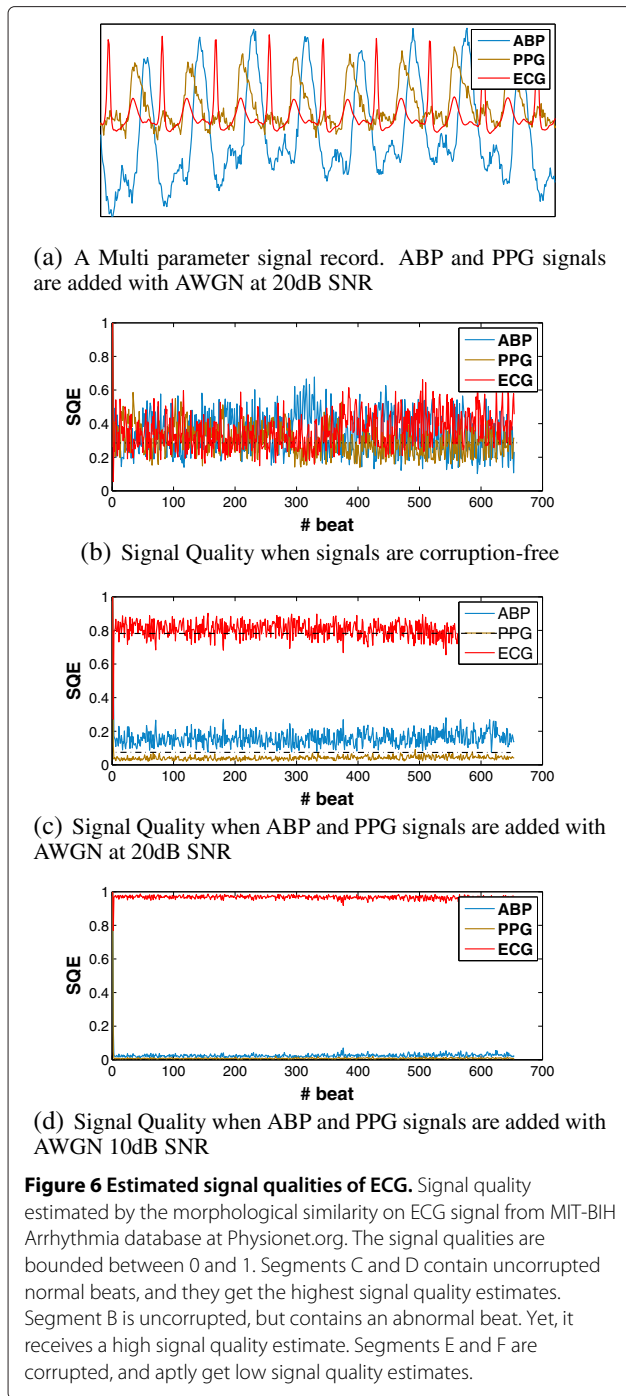
Unlike the related methods, morphological similarity is applicable to each of, ECG, ABP and PPG signals. Therefore, we are able to use a single metric across different types of signals (Figure 7). This allows us to use the signal quality estimates obtained from the morphological similarity to weigh the individual signal’s influence on the joint segmentation of a multi parameter signal.

Feature representation

In reconstruction, we search the database for the closest match to the current segment. Since we want to do it in real time with a growing database, we need a fast method of retrieval.

Since the segments are usually of different lengths, a direct comparison function, such as Euclidean distance is not suitable. On the other hand, variable length metrics such as DTW [14] and LCSS [20] are quadratic in the length of the sequence. Using one of these techniques would result in a complexity of $O(n\ell^2)$, where ℓ is the length of the sequence and n is the number of templates in the data base.

We use an intermediate feature-based representation of the segment to reduce the dimensionality of search. This represents a segment with a fixed length vector, hence two



sequences can be compared in linear time in feature space. Just as importantly, it provides a level of abstraction that preserves relevant physiological information. Table 1 lists the set of features in the feature representation $F_i = \{f_j | j = 1, 2, \dots, 30\}$. The first four of them are related to the segment intervals, which capture the contextual details of the beat, followed by the features that are both physiologically and morphologically representative of the signal in the segment.

Once we generate the features from the segments, we normalize them so that each feature has zero mean and unit variance. Next, using the standardized features, we perform a nearest neighbor search in the database to select the best $K = 20$ matches. The search is done using locality sensitive hashing as described in [21]. From the top K matches, we find the candidate for replacement \hat{Y}_i using both DTW distance and feature distance.

$$\hat{Y}_i = \arg \min_{k \in K} \left\{ \sum_j (F_{k,j} - F_{i,j})^2 + \lambda f_{\text{dtw}}(Y_k^a, Y_i^a) \right\} \quad (10)$$

$$c_i = \min_{k \in K} \left\{ \sum_j (F_{k,j} - F_{i,j})^2 + \lambda f_{\text{dtw}}(Y_k^a, Y_i^a) \right\} \quad (11)$$

Here, DTW distance (f_{dtw}) is computed between the clean channels (correlated signals) of the segments (Y_k^a and Y_i^a). The regularizer $\lambda = 25$ was empirically chosen. The summation is performed over all the features in the feature vector. The combination helps avoid over-fitting and balances the preservation of clinically relevant events and morphological similarity.

Reconstruction

Assume that we want to reconstruct the corrupted channel Y_i^1 of the current segment Y_i with the corresponding channel \hat{Y}_i^1 from the replacement candidate \hat{Y}_i . Further, Y_i^a is the correlated channel.

We first verify that the purportedly uncorrupted channel of the current segment is sufficiently similar to the corresponding channel in the replacement candidate. We perform reconstruction only if the cost of the match c_i is less than a threshold. If the cost c_i is greater, we flag the segment Y_i so that automated systems could avoid producing false alarms in those regions.

Since the length of the current segment Y_i , and the length of the candidate found (match) \hat{Y}_i are typically unequal, we next time warp the match with the current segment. Time warping is done by finding the optimal alignment $\phi(k)$ between the corruption-free channel of the current segment Y_i^a and the corresponding channel of the match \hat{Y}_i^a (Equations 12 and 13).

$$\phi(k) = (\phi_1(k), \phi_2(k)), 1 \leq k \leq K \quad (12)$$

$$C_\phi(\hat{Y}_i^a, Y_i^a) = \sum_{k=1}^K d(\hat{Y}_i^a[\phi_1(k)], Y_i^a[\phi_2(k)]) \quad (13)$$

$$C(\hat{Y}_i^a, Y_i^a) = \min_{\phi} C_\phi(\hat{Y}_i^a, Y_i^a) \quad (14)$$

We then replace each sample of the corrupted channel $Y_i^1[x]$ with the time-warped sample $\hat{Y}_i^1[\hat{x}]$, which is obtained from the median of the samples with which it is aligned.

$$\hat{x} = \text{median}(\phi_2(k)), 1 \leq k \leq K \text{ and } \phi_1(k) = x \quad (15)$$




Figure 7 Estimated signal qualities of a multi parameter signal. Signal quality estimates (SQE) of a multi parameter physiological signal. When all three signals are corruption-free, SQEs are approximately equal, and take an average of 0.33. When 20 dB AWGN is added to ABP and PPG signals **(a)**, it results in an SQE 5× higher for ECG than others (average SQE for ECG is 0.8 compared to 0.16 of ABP) **(b)**. At 10 dB AWGN the difference is more prominent **(c)**.

Table 1 Features

Feature	Description
$f_1 - f_4$	Pre, first-half, second-half, and post segment intervals
f_5	Square root of the total energy
$f_6 \dots f_{15}$	The fraction of the energy in the k th section (a segment is divided into 10 sections)
f_{16}	Kurtosis of the sample values
f_{17}	DTW distance between the signal in the segment, and the median of the same signal
$f_{18} \dots f_{27}$	DTW distance of k th subsequence (DTW alignment is divided into 10 sections)
f_{28}	Fraction of spectral energy in bandwidth 5–14 Hz.
f_{29}	The maximum sample value
f_{30}	The minimum sample value

Ensembles

Thus far, we have described applying our reconstruction framework in the context of a pair of correlated signals. In this section, we describe how to extend it to a situation in which there is more than one uncorrupted correlated signal. For each uncorrupted correlated signal, we find the entry from the corresponding database that most closely matches the current segment. We weight each entry by the estimated quality of the reconstruction, and combine the reconstructions.

For the segment Y_i , let \hat{Y}_i^1 be the corrupted channel, $\{\hat{Y}_i^j | j = 2, 3, \dots, m\}$ be the matches found using correlated signals $j = 2, 3, \dots, m$, and $\{c_i^j | j = 2, 3, \dots, m\}$ be the costs of the matches. The reconstruction \hat{Y}_i^1 is obtained by fusing the matches,

$$\hat{Y}_i^1 = \frac{\sum_{j=2}^m \hat{Y}_i^j (w_i^j)^p}{\sum_{j=2}^m (w_i^j)^p} \quad (16)$$

$$w_i^j = q_i^j r_i^j / c_i^j. \quad (17)$$

Here, q_i^j is the estimated signal quality of j th channel in i th segment, and r_i^j is the correlation coefficient between the 1st channel and j th correlated channel estimated around the segment \hat{Y}_i^j . The choice of p influences the performance of our method. For instance, when we have a diverse set of signals (e.g., ECG channels, ABP, and PPG), lowering p to one seems to improve the performance. However, when we have a set of highly correlated signals (e.g., several ECG channels), making $p \rightarrow \infty$ tends to perform better; this is similar to selecting the most useful signal $\hat{Y}_i^1 = \hat{Y}_i^k$, where $k = \arg \max_{j=2,3,\dots,m} w_i^j$.

Our approach is based on a “bucket of models” concept, and hence is an ensemble method. Typically, model-selection is performed by cross-validation in ensemble methods. However, because of the autoregressive nature

of the data, we use the hypothesized quality of the reconstruction (w_i^j) to perform the model-selection (or model-fusion depending on the choice of p), and determine the model weights, which get updated in real time to follow the signal evolution.

Results

The organization of this section is as follows. In Section ‘Segmentation’, we present the evaluation of the segmentation algorithm. In Section ‘Reconstruction’, we discuss the methodology that is used to evaluate the reconstruction, and present the results of a series of tests in which comparisons are made using this evaluation methodology.

In our experiments, we artificially corrupt one or more channels of each record by adding different kinds of synthetic corruptions. Table 2 lists these corruptions.

Segmentation

We compare our method’s effectiveness to that of a widely used QRS detection based segmentation method [1,25] that is known to be resilient against signal noise [12,26]. It is publicly available as open source software [6].

We use clean excerpts from the MIMIC data at Physionet.org [27] to evaluate our segmentation method. The database has 72 waveform records with several annotation sets including ECG beat labels. It includes recordings from multiple ECG channels, ABP, PPG, and other signal types (Figure 8). The signals are sampled at 125 Hz.

We randomly extracted 1000 high quality 5 min long excerpts from the MIMIC data, for which both our method and the QRS detector were 100% accurate in detecting the segment boundaries. First, we added additive white Gaussian noise (AWGN) at different Signal/Noise (SNR) levels to these excerpts and tested both the methods on them. We carried out three sets of tests. In the first set, we added white noise to all m channels, and in the second set to $m - 1$ channels. In the third set, we added transient corruption on five randomly selected non-overlapping 1 min long regions. We applied one of the following types of transient corruptions: signal interruption, exponential damping, overshooting and clipping, or superimposition of artificial low frequency signals and high frequency signals.

Table 3 presents the summary of the experimental results on the artificially corrupted data. Average segment length is 521 ms. The reported error is the mean distance over the 1000 excerpts between an actual segment boundary and the closest segment boundary found by either method. Our method was able to identify the segment boundaries with reasonable accuracy even in the presence of significant additive noise on all channels. If any one of the channels is free of noise, the performance is comparable to the performance on the uncorrupted data. Under transient corruption, the average error was

Table 2 Types of synthetic corruptions used in experiments

Experiment 1: segmentation	Experiment 2: reconstruction
Additive white Gaussian noise (AWGN): we use a standard AWGN algorithm [22]	
Transient corruption: We use the following types of transient corruptions. <ul style="list-style-type: none"> • Signal interruption: We make the amplitude of the signal zero. • Exponential Damping: For signal $x(t)$; $1 \leq t \leq \ell_x$, we damp the signal amplitude by applying the function $f(x(t)) = x(t) e^{(-\beta t)}$. Here, $\beta = 0.001$ is damping coefficient. • Overshooting and clipping: We amplify the signal by a randomly chosen factor (uniformly distributed) between 1 and 5, and clip at the maximum and minimum values of the original signal. • Super imposition of high frequency signal: We add white noise of the same power as the signal, high pass filtered at 150 Hz. • Super imposition of low frequency signal: We add white noise of the same power as the signal, low pass filtered at 0.05 Hz. 	Non-Gaussian corruption: We use the following types of structured (non-Gaussian) interferences in the evaluation of reconstruction. We use MIT-BIH noise stress test database [23] and nstdbgen [24] to generate these interferences. <ul style="list-style-type: none"> • Electromagnetic interference (EM) • Muscle artifact (MA) • Baseline wander (BW)

2.89 ms for our method. This is comparable to that of the data with AWGN at SNR 10 dB on all m channels. Under severe transient corruption the QRS detector becomes totally unusable, whereas our method finds the segment boundary with reasonable accuracy.

Reconstruction

Since our method is based on patient-specific learning, the method of training is an important part of the evaluation.

We used two approaches to the training, called offline and online. We use the following evaluation methodology on all subsequent experiments.

Evaluation methodology

In our experiments, we artificially corrupt one channel, for instance the first channel S^1 of the signal S , and try to reconstruct it. We use two criteria to evaluate the effectiveness of the reconstruction.

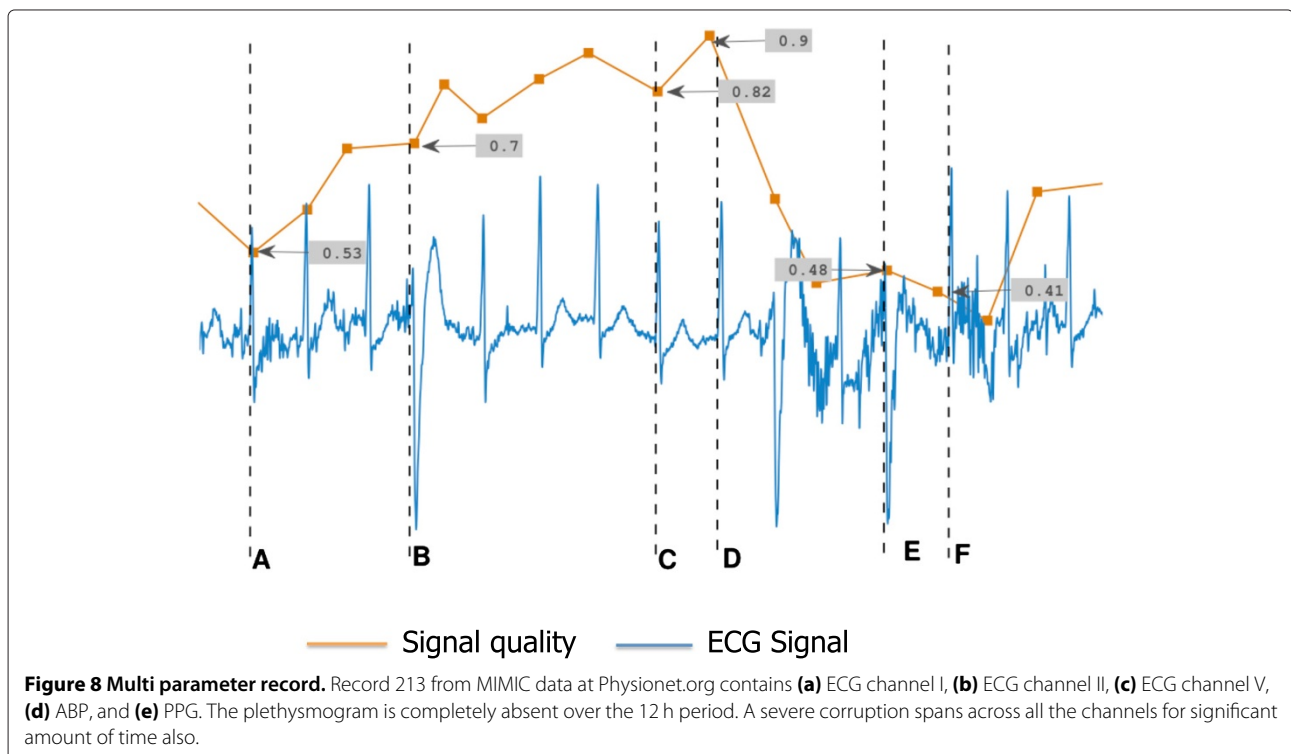


Table 3 Evaluation of segmentation algorithm

Type	Average error (ms) on artificially corrupted data			
	QRS	Our method		
		all m	any $m - 1$	
AWGN	20 dB	12.87	0.87	0.0097
(SNR)	10 dB	188.11	3.27	0.0011
	0 dB	303.48	5.81	0.008
Transient corruption		387.32	2.89	

1. Residual distance (r): We measure the similarity between two sequences, for instance the reconstructed data (\hat{S}^1) and the original uncorrupted data (S^1), by measuring the normalized Euclidean residual distance.

$$r(\hat{S}) = \sqrt{\frac{\sum_{k=1}^n (\hat{S}^1[k] - S^1[k])^2}{n \times \sigma_S^2}} \quad (18)$$

Here, n is the length of the signal S , and σ_S^2 is signal variance of the corrupted channel. Similarly, $r(\hat{S})$ can also be computed between the corrupted data \hat{S}^1 and the original data S^1 .

2. Classification accuracy (Δ): Our ultimate goal is to enable automated analysis systems to produce results that are more reliable. Hence, we test our method's ability to improve the classification accuracy of a clinically relevant task. We run a widely used premature ventricular contraction (PVC) detector [6], on the original data (S^1), the artificially corrupted data (\check{S}^1) and the reconstructed data (\hat{S}^1), and record their agreements. If the PVCs are detected within 150 ms on two signals, we consider it an agreement. The disagreement Δ is finally expressed in terms of the fraction between the total number of disagreements, and the total number of beats in the region.

We evaluated our study on two sets of data. Most of the experiments use the multi parameter physiological signal data from CinC Challenge 2010 Test Set at Physionet.org [27]. The diversity of heart beat types in the multi parameter dataset is limited, which makes it less than ideal for testing classification accuracy. Hence, we also evaluate our study on two-channel ECG data from the MIT-BIH arrhythmia database at Physionet.org [27], a widely used bench-mark database.

The CinC database has 100 waveform records of ECG, ABP, and PPG signals, each 10 min long, sampled at 125 Hz. We use only the 70 records that contain at least 3 channels that are relatively free of transient corruption and misalignment.

The arrhythmia database contains 48 half-hour ECG recordings, sampled at 360 Hz. The recordings were selected to include a variety of clinically significant arrhythmias. The database has 48 ECG waveform records; each contains two channels and is 30 min long. We use the 39 records from this set that are relatively free of significant corruption on both channels. An additional file lists the records from MIT-BIH Arrhythmia Database that were used in our experiments [see Additional file 1].

In our experiments, we artificially corrupt one channel of each record, and reconstruct it using our method. The original uncorrupted signal serves as a gold standard in our experiments. The method is evaluated by comparing both the corrupted channel and the reconstructed channel to the original channel. We use both additive white Gaussian noise (AWGN) and various kinds of non-Gaussian (structured) interference to generate synthetic corruption. We use MIT-BIH Noise Stress Test Database [28] and *nstdbgen*[23] to generate non-Gaussian interferences.

Offline learning

We build the template database using the first 80 % of each record, add AWGN at signal to noise ratio (SNR) 0 dB to last 20 % of the last channel of each record, and evaluate the effectiveness of reconstruction on this artificially corrupted data.

As Table 4 shows, our method improves the signal quality of the artificially corrupted multi parameter data by reducing the average residual distance by 33 \times and the average disagreement by 2.8 \times . Residual distance of the reconstructed data in the CinC multi parameter dataset is smaller than that in the two-channel arrhythmia dataset. We attribute this to the availability of multiple correlated

Table 4 Evaluation of offline learning reconstruction

Dataset	Residual distance (r)		Disagreement (Δ)	
	$r(\check{S})$	$r(\hat{S})$	$\Delta(\check{S})$	$\Delta(\hat{S})$
CinC 2010 (multi parameter)	1.00	0.03	0.14	0.05
MIT-BIH (two-channel)	1.01	0.40	0.14	0.02

signals in the multi parameter dataset and the utility of the ensemble approach that fuses the information from multiple correlated signals. On the other hand, the improvement in classification accuracy is higher for CinC dataset. We attribute this to the abundance of PVC beats in that dataset, thus making it easier to find morphologically similar segments for rare beat types.

CinC 2011 challenge

In the 11th annual PhysioNet/CinC [7] contestants were asked to reconstruct, using any combination of available prior and concurrent information, 30 s segments of ECG, continuous blood pressure waveforms, respiration, and other signals that had been removed from recordings of patients in ICU. The submissions to the contest used several methods such as neural networks, Kalman and adaptive filters, hidden Markov models, and principal component analysis [24,29-32].

The highest scored method [24] uses a multilayer perceptron neural network with a fixed length window to perform the regression. The method’s reported normalized residual distance for the contest is 0.17, compared to our method’s 0.02. To be fair, about 20 % of the contest involved signals that we did not test, a respiratory signal and two other pressure waveforms. These signals lack the quasiperiodic morphological structures, and have little correlation with the other signals, and therefore they are not suitable targets for our method.

Characteristics of interference

Although our method doesn’t make any assumptions about the type or level of corruption, it’s performance can depend upon the type and level of corruption. We use MIT-BIH Arrhythmia dataset for to test this. We corrupt the first 20 % of the first channel with AWGN at SNR levels of 10, 0, and –10 dB.

Table 5 summarizes the average disagreement (Δ), and the residual distance (r) for the reconstructed signal. At low SNR levels the performance doesn’t deteriorate with decreasing signal quality. But, somewhat surprisingly, we get the worst performance at the highest signal to noise ratio. The relatively poor performance at SNR 10 dB can be attributed to the fact that our algorithm makes a binary decision to reconstruct the signal or leave it as it is. This results in relatively poor performance, because when the

Table 5 Performance at different noise levels

SNR(dB)	$r(\hat{S})$	$\Delta(\hat{S})$
10	0.410	0.031
0	0.401	0.021
-10	0.402	0.022

Table 6 Performance with structured interference

	$r(\hat{S}0)$	$\Delta(\hat{S})$
AWGN	0.410	0.031
EM	0.36	0.023
MA	0.19	0.003
BW	0.05	0.001

signal is only mildly corrupted our algorithm chooses not to reconstruct it.

Next, we test our method against different structured interferences. Again, we build our database from the first 80 % of each record. We corrupt the last 20 % of the first channel with the following types of corruptions at SNR 10 dB: additive white Gaussian noise (AWGN), electromagnetic interference (EM), muscle artifact (MA), and baseline wander (BW). We chose 10 dB because that was the noise level at which our method exhibited the worst performance in the test reported above. We use MIT-BIH Noise Stress Test Database [28] and *nstdbgen*[23] to generate non-Gaussian interferences. Table 6 summarizes the average disagreement (Δ), and the residual distance (r) for the reconstructed signal. The worst performance was observed for AWGN, and EM noise.

Online learning

In this experiment, we add artificial corruption to the last minute of the data and to 1 min long regions starting at the 2nd, 4th, or 6th minute. The idea is to keep the span of the corrupted region to 20 % of the data, and evaluate the performance of our method. The algorithm builds an initial database from the first 2 min of the data, and adds new entries to the database whenever it encounters a segment that it judges to be relatively clean.

Table 7 presents the results of the experiment. Our framework’s performance seems to be independent of the location of corruption. Though, we start with a database built from only 20 % of the signal, we are able to achieve performance similar to the offline approach, where we built the database from the first 80 %. The only difference is a slight decline in classification accuracy. We attribute this to an increased probability that the training data doesn’t include PVC’s with morphologies similar to those in the test data. This is also consistent with the fact that

Table 7 Evaluation of online learning reconstruction

Corruption starts at (mins)	Residual distance (r)		Disagreement (Δ)	
	$r(\hat{S})$	$r(\hat{S})$	$\Delta(\hat{S})$	$\Delta(\hat{S})$
2	0.97	0.042	0.15	0.08
4	0.97	0.041	0.14	0.06
6	0.97	0.038	0.14	0.05

our algorithm performs better when the training data is more proximate in time to the test data.

Conclusion

The goal of this study was to find a way to recover clinically useful information in corrupted quasiperiodic physiological signals.

The fundamental premise of this study is that given a multi parameter signal, information in some signals can be used to infer the value of other signals. Based on this premise, we approached the problem of reconstruction as a learning problem. The key technical idea is to use templates to capture the relationships between the morphologies of correlated signals. These templates are learned on a patient-specific basis, and they continue to evolve as data is processed. A level of abstraction is provided by converting variable length segments of a physiological signal into fixed length feature vectors that capture relevant properties of the segment and the context in which it occurs.

The use of a feature vector also plays a vital role in speeding up lookups so that reconstruction can be done in real time. Our method has a running time that is linear on the number of samples n , and hence is suitable for real time applications. Using two levels of indirection (the fixed length feature vector and locality sensitive hashing based nearest neighbor lookup) we achieve a running time that is independent of the size of the database.

We evaluated our method on MIT-BIH arrhythmia data, a two-channel ECG dataset with many clinically significant arrhythmias, and on the CinC challenge 2010 data, a multi-parameter dataset containing ECG, ABP, and PPG. For each, we evaluated both the residual distance between the original signals and the reconstructed signals and the performance of a heartbeat classifier on a reconstructed ECG signal. At an SNR of 0 dB, the average residual distance on the CinC data was only roughly 3% of the energy in the signal, but on the arrhythmia database it was roughly 16%. The difference is attributable to the large amount of diversity in the arrhythmia database. Despite the relatively high residual difference, the classification accuracy on the arrhythmia database was still 98%, indicating that our method did restore the physiologically important aspects of the signal. The results strongly suggest that our initial premise is valid and that templates are indeed a good way to capture the relationships among correlated signals.

While we have tested our method only on ECG, ABP, and PPG signals, we believe that it could be useful in other multi-signal settings in which one or more signals are corrupted and at least one correlated signals is transiently uncorrupted [29,31]. Going forward, we plan to test our algorithm on a database containing simultaneous

recordings of other signals such as CVP, PAP, and respiratory signals.

Because of inter-patient variation the use case for the method involves learning the patient's baseline before starting reconstruction. It continues to learn as the patient's baseline changes over time. If there is an abrupt change in the baseline, however, the method will not work. This might not be a practical problem, since it is probably appropriate that an abrupt change in the baseline triggers an alarm.

In other future study, we plan to explore using transfer learning to bootstrap the process. For instance, we would like to test the effectiveness of beginning the reconstruction process with a generic template database, and then adapt it to the current record by learning new morphological relationships in the current record. This is useful at the beginning of the reconstruction process when the database would be otherwise empty.

Additional file

Additional file 1: MIT-BIH Arrhythmia Database Records.

Competing interests

Both authors declare that they have no competing interests.

Acknowledgements

We would like to thank Collin Stultz, Alaa Kharbouch, and Zeeshan Syed for their valuable suggestions, and Jessica Liu for her help with running the experiments. We would also like to thank the anonymous reviewers and the editor whose many constructive suggestions led to considerable improvement in the manuscript. This study was supported by Quanta Computer Inc.

Received: 31 October 2011 Accepted: 19 July 2012

Published: 16 August 2012

References

1. J Pan, WJ Tompkins, A real-time QRS detection algorithm. *IEEE Trans. Biomed. Eng.* **BME-32**(3), 230–236 (1985)
2. YH Hu, S Palreddy, W Tompkins, A patient-adaptable ECG beat classifier using a mixture of experts approach. *IEEE Trans. Biomed. Eng.* **44**(9), 891–900 (1997)
3. P de Chazal, R Reilly, A patient-adapting heartbeat classifier using ECG morphology and heartbeat interval features. *IEEE Trans. Biomed. Eng.* **53**(12), 2535–2543 (2006)
4. Z Syed, *Computational Methods for Physiological Data* Ph.D. dissertation (MIT, Cambridge, MA 2009)
5. JY Wang, A new method for evaluating ECG signal quality for multi-lead arrhythmia analysis. *Comput. Cardiol.* **29**, 85–88 (2002)
6. EP Limited. <http://www.eplimited.com/software.htm>. Accessed on 30 Oct 2011
7. G Moody, The physionet/computing in cardiology challenge 2010: mind the gap. *Comput. Cardiol.* **37**, 305–308 (2010)
8. Q Li, RG Mark, GD Clifford, Robust heart rate estimation from multiple asynchronous noisy sources using signal quality indices and a Kalman filter. *Physiol. Meas.* **29**, 15–32 (2008)
9. Q Li, RG Mark, GD Clifford, in *ICBBE'2008: 2nd international conference on bioinformatics and biomedical engineering, 2008* Suppress false Arrhythmia alarms of ICU monitors using heart rate estimation based on combined arterial blood pressure and ECG analysis. (2008), pp. 2185–2187
10. A Deshmane, *False Arrhythmia Alarm Suppression using ECG, ABP, and Photoplethysmogram* M.S. thesis (MIT, Cambridge, MA, 2009)

11. A Aboukhalil, L Nielsen, M Saeed, R Mark, G Clifford, Reducing false alarm rates for critical arrhythmias using the arterial blood pressure waveform. *J. Biomed. Inf.* **41**(3), 442–451 (2008)
12. BU Kohler, C Hennig, R Orglmeister, The principles of software QRS detection. *IEEE Eng. Med. Biol. Mag.* **21**, 42–57 (2002)
13. G Gartheeban, J Gutttag, in *BIO SIGNALS'2011 International Conference on Bio-Inspired Systems and Signal Processing* Weighted time warping for temporal segmentation of multi-parameter physiological signals. (Rome, Italy, 2011), pp. 125–131
14. C Myers, L Rabiner, A Rosenberg, Performance tradeoffs in dynamic time warping algorithms for isolated word recognition. *IEEE Trans. Acoustics, Speech Signal Process.* **28**(6), 623–635 (1980)
15. Z Syed, J Gutttag, C Stultz, Clustering and symbolic analysis of cardiovascular signals: discovery and visualization of medically relevant patterns in long-term data using limited prior knowledge. *EURASIP J. Adv. Signal Process.* **2007**, 1–17 (2007)
16. D Clifford, F Azuaje, *Advanced Methods and Tools for ECG Data Analysis* (Artech House, Boston, 2006)
17. M Oefinger, Monitoring transient repolarization segment morphology deviations in Mouse ECG. *Ph.D. thesis* (MIT, Cambridge, MA, 2006)
18. W Zong, G Moody, R Mark, Reduction of false blood pressure alarms by use of electrocardiogram blood pressure relationships. *Comput. Cardiol.* **26**, 305–308 (1999)
19. B Hjorth, The physical significance of time domain descriptors in EEG analysis. *Electroencephalogr. Clin. Neurophysiol.* **34**(3), 321–325 (1973)
20. M Vlachos, M Hadjieleftheriou, D Gunopulos, E Keogh, Indexing multidimensional time-series. *VLDB J.—Int. J. Very Large Data Bases.* **15**, 1–20 (2006)
21. A Andoni, Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM.* **51**, 117–122 (2008)
22. K McClaning, T Vito, *Radio Receiver Design* (Noble Pub, Atlanta, 2001)
23. Noise Stress Test Data Generator. <http://www.physionet.org/physiotools/wag/nst-1.htm>. Accessed on 30 Oct 2011
24. R Rodrigues, Filling in the gap: a general method using neural networks. *Comput. Cardiol.* **37**, 453–456 (2010)
25. J Urrusti, W Tompkins, in *EMBC'1993:15th Annual International Conference of the IEEE* Performance evaluation of an ECG QRS complex detection algorithm. (1993), pp. 800–801
26. G Friesen, T Jannett, M Jadallah, S Yates, S Quint, H Nagle, A comparison of the noise sensitivity of nine QRS detection algorithms. *IEEE Trans. Biomed. Eng.* **37**, 85–98 (1990)
27. AL Goldberger, LAN Amaral, L Glass, JM Hausdorff, PC Ivanov, RG Mark, JE Mietus, GB Moody, CK Peng, HE Stanley, PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation.* **101**(23), e215–e220 (2000)
28. MIT-BIH Noise Stress Test Database. <http://www.physionet.org/physiobank/database/nstadb>. Accessed on 30 Oct 2011
29. I Silva, PhysioNet 2010 challenge: a robust multi-channel adaptive filtering approach to the estimation of physiological recordings. *Comput. Cardiol.* **37**, 313–316 (2010)
30. Z Fejzo, H Lev-Ari, Adaptive Laguerre-lattice filters. *IEEE Trans. Signal Process.* **45**(12), 3006–3016 (1997)
31. A Hartmann, Reconstruction of missing cardiovascular signals using adaptive filtering. *Comput. Cardiol.* **37**, 321–324 (2010)
32. Y Li, Y Sun, P Sondhi, L Sha, C Zhai, Reconstructing missing signals in multi-parameter physiologic data by mining the aligned contextual information. *Comput. Cardiol.* **37**, 449–452 (2010)

doi:10.1186/1687-6180-2012-173

Cite this article as: Ganeshapillai and Gutttag: Real time reconstruction of quasiperiodic multi parameter physiological signals. *EURASIP Journal on Advances in Signal Processing* 2012 **2012**:173.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com
