

MIT Open Access Articles

*Weighted Time Warping for Temporal Segmentation
of Multi-Parameter Physiological Signals*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: "Weighted Time Warping for Temporal Segmentation of Multi-parameter Physiological Signals." Proceedings of the International Conference on Bio-inspired Systems and Signal Processing, Rome, Italy, (January, 2011) 125-131.

As Published: http://www.biosignals.biostec.org/Abstracts/2011/BIOSIGNALS_2011_Abtracts.htm

Publisher: Biosignals

Persistent URL: <http://hdl.handle.net/1721.1/73944>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



WEIGHTED TIME WARPING FOR TEMPORAL SEGMENTATION OF MULTI-PARAMETER PHYSIOLOGICAL SIGNALS

Gartheeban Ganeshapillai, John Guttag

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

Keywords: temporal segmentation, multi-parameter physiological signals, time warping, time series.

Abstract: We present a novel approach to segmenting a quasiperiodic multi-parameter physiological signal in the presence of noise and transient corruption. We use Weighted Time Warping (WTW), to combine the partially correlated signals. We then use the relationship between the channels and the repetitive morphology of the time series to partition it into quasiperiodic units by matching it against a constantly evolving template. The method can accurately segment a multi-parameter signal, even when all the individual channels are so corrupted that they cannot be individually segmented.

Experiments carried out on MIMIC, a multi-parameter physiological dataset recorded on ICU patients, demonstrate the effectiveness of the method. Our method performs as well as a widely used QRS detector on clean raw data, and outperforms it on corrupted data. Under additive noise at SNR 0 dB the average errors were 5.81 ms for our method and 303.48 ms for the QRS detector. Under transient corruption they were 2.89 ms and 387.32 ms respectively.

1 INTRODUCTION

In this paper, we address the problem of segmenting a quasiperiodic multi-parameter physiological signal in the presence of noise and transient corruption.

Early warning systems in an ICU require continuous uninterrupted real-time monitoring of the physiological signals (Tarassenko et al., 2006; Mark and Shavdia, 2007; Cao et al., 2008; Henriques and Rocha, 2009; Chen et al., 2009). Unfortunately, these signals often suffer transient corruption. An algorithm that estimates the correct values of the corrupted signals can help automated systems produce more reliable results, and make them more amenable for visual inspection. An accurate segmentation makes the estimation task easier.

We represent a multi-parameter signal by a matrix $S_{n \times m}$, 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. If the matrix represents a quasiperiodic multi-parameter signal (QPMS), it will have




Figure 1: The Electrocardiogram (ECG), Arterial Blood Pressure (ABP) and the Photo Plethysmogram (PPG) extracted from MIMIC, record 039.

a repetitive structure that is shared by all the channels in the structure. This is common in situations where the signal is generated by the same underlying sys-

tem. In the case of the cardiovascular system, the periodicity of all the ECG channels, the blood pressure waveforms, and the plethysmogram are related to the heart rate.

In practice, samples may be corrupted in an unknown fashion. We would like to estimate the actual values of the samples in row i of the matrix \mathbf{S} that represents a QPMS, using the corrupted values in that row and the estimates of prior values up to row $(i - 1)$. However, many physiological signals (e.g., those related to cardiovascular activity) are quasiperiodic. For such signals, the estimation process starts by identifying the segment boundaries.

This paper presents a novel approach to identifying segment boundaries in the presence of significant amounts of transient corruption spanning multiple columns and rows of the matrix \mathbf{S} . The key idea is that by simultaneously considering all the channels one can segment them more accurately than would be possible by considering each channel independently. The method is based on matching a sliding window of the QPMS to a template. The template is a short multi-parameter signal that is regularly updated based upon recent estimates. The initial template is derived from an archived QPMS.

Segmentation is done by finding the prefix of the window that most closely matches the template. The matching is done using a new method, Weighted Time Warping (WTW) that minimizes the weighted morphological dissimilarity between template and the prefix of the window across all the channels. The morphological dissimilarity is given by the warped distance between a channel in the window and the corresponding channel in the template. The weight represents the estimated quality of the channel.

Experiments carried out on MIMIC (Goldberger et al., 2000), a publicly available multi-parameter physiological dataset recorded from ICU patients, demonstrate the effectiveness of the method. We considered additive noise and the following types of corruptions: signal interruption, exponential damping, overshooting, clipping, and superimposition of artificial low frequency signals and high frequency signals. Here, the corruptions across the channels are not perfectly correlated.

We compared our method to a widely used QRS detection based segmentation method. Our method performs as well as the QRS detector on the raw data and significantly outperforms it on the data synthetically altered with additive noise and transient corruption. When additive noise at SNR 0 dB was applied to all m channels of the signal, the average errors were 5.81 ms for our method and 303.48 ms for the QRS detector. The average errors were 2.89 ms and 387.32

ms respectively when transient corruption was added. In the last two cases, the QRS detector was totally unusable, whereas our method was able to do the segmentation with reasonable accuracy.

The primary contributions of this work are

- Formulation of the problem of segmentation as a dissimilarity minimization problem, and the use of a new dissimilarity metric, weighted time warping.
- The use of a multi-parameter signal template that eliminates the need of any prior knowledge of the specific properties of the signal, and
- The dynamic adaptation of the template, which allows us to accommodate the time evolution of the signal.

2 BACKGROUND AND RELATED WORK

Physiological signals in the ICU are often severely corrupted; corruptions by noise, artifact and missing data lead to serious errors in automated medical systems and early warning systems. (Li et al., 2008) provides a survey of strategies used to address this problem. Recent attempts to mitigate this problem focus on using redundant measurements, and fusing data from multiple sensors (Li et al., 2008; Aboukhalil et al., 2008; Li and Clifford, 2008; Deshmene, 2009). Typically they employ independent methods on different channels and combine the results only at the final stage. In contrast, we process all channels simultaneously.

Dynamic Time Warping (DTW) is increasingly being used in temporal segmentation problems (Kovar and Gleicher, 2004; Zhou et al., 2008). (Park and Glass, 2006) uses DTW to segment speakers in an audio signal. DTW has also been used in locating motion clips (Gleicher and Kovar, 2004) and temporal segmentation of human motions (Zhou et al., 2008). (Vlachos et al., 2006) provides an overview on the implementation of multidimensional Dynamic Time Warping (DTW) over the L_p norm. Based on this work, we propose Weighted Time Warping (WTW), a novel method that uses time warping to perform segmentation.

WTW generalizes multidimensional DTW by weighing the individual signals by the signal quality. In WTW, we also apply the local continuity and global path constraints that reduce the influence of outliers. The use of such constraints was proposed in (Myers et al., 1980).




Figure 2: This template comprises of clean ECG and Arterial Blood Pressure (ABP) waveforms. The positions of the segment boundaries are denoted by ℓ_1 , ℓ_2 and ℓ . The template is little longer than two segments. It contains two full segments of length $\ell_2 - \ell_1$ and $\ell - \ell_2$; the length of the template is ℓ .

3 METHOD

First we present the overview of the algorithm, followed by the discussion of a single iteration of the loop. Additional details are provided later in this section.

3.1 Overview

Goal : Let $\mathbf{S} \in \mathbb{R}^{n \times m}$ be a multi-parameter time-series consisting of m and $\mathbf{Z}_{\ell \times m} = \{Z_j \in \mathbb{R}^\ell\}_m$ an initial template.¹ The goal is to segment \mathbf{S} into a set of quasiperiodic units $\mathbf{Y} = \{\mathbf{Y}_i\}$ where $\mathbf{Y}_i \stackrel{\text{def}}{=} \mathbf{S}_{[p_i, p_{i+1})}$. Here, $\mathbf{S}_{[p_i, p_j)}$ denotes the window in the target sequence \mathbf{S} from time $t = p_i$ to $t = p_j - 1$.

We require the template (Figure 2) to be comprised of at least two segments. These segments are used to find the quasiperiodic unit \mathbf{Y}_i in \mathbf{Y} . We also assume that we know the locations of the segment boundaries $\mathbf{Z}.\ell_1$, $\mathbf{Z}.\ell_2$ and $\mathbf{Z}.\ell$ in the template. Here, $\mathbf{Z}.k$ denotes the k^{th} row in \mathbf{Z} . It is a vector of samples corresponding to time k . The prototypical template is initially obtained from an archive.

Procedure : 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 then run the algorithm starting at p_{start} , continuously segment \mathbf{S} , and add the segments to \mathbf{Y} . We also

¹ Bold capital letters represent a matrix \mathbf{S} , non-bold capital letters denote a column vector A , and lower-case letters denote a scalar w .

update \mathbf{Z} to reflect the evolution of the time series. This enables us to accommodate the gradual changes in the morphology of the signal.

An iteration : We start the iteration with the extraction of a window $\mathbf{W} = \mathbf{S}_{[p_i, p_i+v)}$ from the time series data \mathbf{S} at p_i . Here, the window length is given by $v = \ell + e$, where e is the buffer length. In the following discussion, we use j to index the single parameter signals. It is preprocessed, and the morphological quality estimates $\{q_j\}_m$ are computed, where q_j represents the morphological similarity between the channel W_j from \mathbf{W} and the corresponding channel from the template Z_j . For each channel j , a pairwise Euclidean distance matrix \mathbf{pD}^j is calculated between Z_j and W_j (Equation 1-2). The final distance matrix \mathbf{D} is obtained by weighting the pairwise distance matrix \mathbf{pD}^j with q_j (Equation 3).

$$c_{x,y} = W_{x,j} - T_{y,j} \quad (1)$$

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

$$\mathbf{D} = \sqrt{\sum_{j=1}^m q_j |\mathbf{pD}^j|^2} \quad (3)$$

The accumulated distance matrix \mathbf{aD} is then computed from \mathbf{D} using dynamic programming. For two single parameter signals $A \in \mathbb{R}^{l_a}$ and $B \in \mathbb{R}^{l_b}$, the accumulated distance can be calculated by the following recursion

$$\begin{aligned} \mathbf{aD}(A_i, B_i) &= \mathbf{D}(A_i, B_i) + \min\{\mathbf{D}(A_{i-1}, B_{i-1}), \\ &\quad \mathbf{aD}(A_{i-1}, B_i), \mathbf{aD}(A_i, B_{i-1})\}. \end{aligned} \quad (4)$$

We next check \mathbf{aD} for such spurious matchings. When a spurious matching is encountered, the buffer length e is increased and the process is repeated. Otherwise we trim \mathbf{W} to obtain \mathbf{W}^* so that it contains only the portion of the signal that matches the template.

We next use the accumulated distance matrix \mathbf{aD} to find the optimal path alignment between \mathbf{Z} and \mathbf{W}^* , as in DTW. From the alignment, we extract the point $\mathbf{W}.f^*$ in \mathbf{W}^* that is matched with the segment boundary $\mathbf{Z}.\ell_2$ in \mathbf{Z} . This corresponds to the segment boundary we are interested in, because $\mathbf{Z}.\ell_2$ marks the end of the first segment in the template. Then, the corresponding length f^* is used to update p_{i+1} with $p_i + f^*$. A second pass of the algorithm is used to fine tune the results. For ECG, a peak detector is used in the second run. 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.

3.2 Weighted Time Warping

We introduce a weighted norm (Equation 3) over the parameter signals to vary the influence exerted by each parameter. The morphological quality metric q_j captures the morphological similarity between the parameter signal W_j and the template Z_j .

We hypothesize that the corrupted regions of the signals are morphologically dissimilar to the clean signals. We estimate the dissimilarity using the warped distance cd_j .

$$cd_j = \min_k \frac{1}{k} \mathbf{aD}^j(\ell, k) \quad (5)$$

$$k_j = \operatorname{argmin}_k \frac{1}{k} \mathbf{aD}^j(\ell, k) \quad (6)$$

Here, \mathbf{aD}^j is the accumulated distance matrix obtained from \mathbf{pD}^j using dynamic programming, k_j is the alignment length, and ℓ is the length of the template.

By simply inverting cd_j we could obtain the morphological quality metric q_j . Because we use a single metric across different channels such as ECG, ABP and PPG, sometimes two morphologically similar but time warped sequences from one channel produce higher cost than two morphologically dissimilar sequences from another channel. To address this, we incorporate additional information such as the alignment length k_j and the difference in the standard deviations of W_j and T_j to obtain q_j . Shorter alignment lengths indicate the partial matchings that usually result in lower costs. Also a significant difference in the standard deviation between W_j and T_j implies that they represent two dissimilar processes. We also use a non-linear transformation to amplify the dynamic range of the values in the following equation that estimates the morphological quality.

$$q_j = \max \left\{ 0, \exp \left(1 - \frac{\lambda \times cd_j}{k_j \times \log(\frac{std_{W_j} + 1}{|std_{T_j} - std_{W_j}|})} \right) - 1 \right\} \quad (7)$$

Here, std_{T_j} and std_{W_j} denote the standard deviation of T_j and W_j respectively. Further, k_j is the length of the path alignment corresponding to cd_j (Equation 5). λ is a normalization coefficient which was empirically chosen to be 0.05.

3.3 Path Constraint

In a typical formulation of DTW, the distance function that is used to solve DTW (Equation 4), allows any path to be taken from (A_0, B_0) to (A_{l_a}, B_{l_b}) . This makes DTW susceptible to degenerate matchings, especially in the presence of noise. For example, a long

subsequence of a signal might be matched with a significantly shorter subsequence of another signal.

Therefore, we use local continuity and global path constraints (proposed as Type III and Type IV local continuity constraints in (Myers et al., 1980)) to prevent such physiologically implausible alignments by updating Equation 4 with

$$\begin{aligned} \mathbf{aD}(A_i, B_i) = & \mathbf{D}(A_i, B_i) + \min \{ \mathbf{D}(A_{i-1}, B_{i-1}), \\ & \mathbf{aD}(A_{i-1}, B_{i-2}) + \mathbf{D}(A_i, B_{i-1}), \\ & \mathbf{aD}(A_{i-1}, B_{i-3}) + \mathbf{D}(A_i, B_{i-1}) + \mathbf{D}(A_i, B_{i-2}), \\ & \mathbf{aD}(A_{i-2}, B_{i-1}) + \mathbf{D}(A_{i-1}, B_i), \\ & \mathbf{aD}(A_{i-3}, B_{i-1}) + \mathbf{D}(A_{i-1}, B_i) + \mathbf{D}(A_{i-2}, B_i) \}. \end{aligned} \quad (8)$$

This ensures that there are no long horizontal or vertical paths in the matrix \mathbf{aD} along the alignment. It also results in global path constraints, by excluding excluding certain parts on the accumulated distance matrix in which optimal warping paths could lie.

3.4 Templates

The templates (Figure 2) are initially derived from an archive of the prototypical multi-parameter signal. They are then updated using the recent signal estimates.

3.4.1 Template Length

When searching for the segment boundaries we only assume the approximate location of the starting point (p_i). This allows us to start the algorithm at an arbitrary location (p_{start}). It also makes the detection of the segment boundary p_{i+1} less sensitive to an inaccuracy in determining the previous segment boundary (p_i). To accommodate this we use templates that are more than two segments long.

3.4.2 Template Update

To follow the gradual changes that are common in the physiological signals, we update the template regularly. However, we do this only when the behavior is consistent, i.e., the recent estimates of the segment lengths have gradual variation. In that case, if all the channels are also clean enough, i.e., $q_j > \beta; \forall j$, we update the template by averaging the excerpt of the last two segments with the time warped version of the current template (Equation 10). Here, 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. Finally we normalize the template.

$$\mathbf{Z}' = \mathbf{S}_{[(p_{(i-1)} - \epsilon), p_{(i+1)}]}; \epsilon \geq 1 \quad (9)$$

$$\mathbf{Z}'' = \text{warp}(\mathbf{Z}) + \eta \mathbf{Z}'; \eta \geq 0 \quad (10)$$

$$\mathbf{Z}^* = \text{normalize}(\mathbf{Z}'') \quad (11)$$

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 η , where $\eta = 0$ implies a static template. We chose $\eta = 0.6$ empirically.

3.5 Long Segments

Spurious matchings occur when the segments in the window \mathbf{W} are significantly longer than the segments in the template \mathbf{Z} ; therefore the current buffer length e is not long enough to span at least two segments in the window. In such cases, the window $\mathbf{W} = \mathbf{S}_{[p_i, p_i+v]}$ would not contain two full segments and does not give the expected matching. In such situations, we have to repeat the matching with a longer window. We increase v through e and repeat the process with a longer window. Otherwise we trim the window and obtain \mathbf{W}^* , the portion of the window that matches the template.

4 EXPERIMENTAL RESULTS

We applied our method to multi-parameter physiological signal data from MIMIC at Physionet.org (Goldberger et al., 2000). The database has 72 waveform records with several annotation sets including ECG beat labels. It includes recordings from multiple ECG channels, Arterial Blood Pressure (ABP) and Photo Plethysmogram (PPG). The signals are sampled at 125 Hz. From this database, we selected 70 records that contained one continuous hour of at least one ECG channel, and ABP, PPG or both.

We compared our method to a widely used QRS detection based segmentation method (Pan and Tompkins, 1985; Urrusti and Tompkins, 1993). This method provides high specificity and sensitivity with low computational load (Kohler et al., 2002). Further it is publicly available as open source software.²

We conducted the following experiments : 1) on the raw MIMIC data we compared our algorithm to the QRS detector, and 2) we selected clean excerpts from the raw data, artificially corrupted them, and compared our algorithm to the QRS detector on this corrupted data.

²<http://www.eplimited.com/>




Figure 3: Record 213 contains a) ECG channel I, b) ECG channel II, c) ECG channel V, d) ABP and e) PPG. The plethysmogram is completely absent over the 12 hour period. A severe corruption spans across all the channels for significant amount of time also.

We evaluate the performance using the number of errors and the average error in milliseconds. The number of errors is defined as the total number of the estimated segment boundaries that are different from actual boundaries by more than one sample. This also includes the cases where a segment boundary is completely missed and an erroneous boundary is found. The average error is the mean distance between an actual segment boundary and the closest segment boundary found by the method.

4.1 Raw MIMIC Data

We performed two tests on the raw data. In the first test, we randomly selected 10 records and ran the tests for the minimum of 12 hours and till the end of the record. We compared the results of our method and the QRS detector to the segment boundaries marked in the data.

The results are presented in Table.1. The results exhibit significant inter-record deviations. This is mainly due to the presence of sporadic corruptions in certain records, which results in a burst of errors in these records. For example, in record 213 (Figure 3), the corruption is severe, and it spans all the channels. It causes significantly large number of errors for both our method and the QRS detector. The median number of errors for our algorithm was 65, whereas it was 330 for the QRS detector.

In the second test, we ran the test across 70 records for an hour. WTW was able to segment 4.121×10^5 segments accurately while the QRS detector was able to segment 4.111×10^5 segments correctly. Although our method made fewer mistakes compared to that made by the QRS detector, the difference in the accuracies of both the methods is not significant because

Table 1: Comparison of WTW against the QRS detector on selected records.

Rec ID	avg Sgmt len (ms)	WTW		QRS	
		Err #	avg Err (ms)	Err #	avg Err (ms)
39	506.25	6	0.01	30	0.25
41	737.91	2	0.00	31	0.47
55	671.13	19	0.05	385	28.27
208	642.67	215	20.63	699	135.39
216	771.33	293	1.10	330	1.64
220	632.88	5	1.65	4	1.20
253	864.40	35	0.15	346	35.13
262	824.93	65	0.30	249	3.55
284	651.70	119	0.23	191	0.96
226	532.17	311	0.55	675	6.25
213	660.31	1921	3.19	2927	42.22

Median number of beats in one record is 56007

of the sparsity of the occurrence of corruption in the raw data.

4.2 Artificially Corrupted MIMIC Data

4.2.1 Additive Noise

In this experiment we evaluated our algorithm’s performance on noisy data. We obtained clean excerpts from the raw MIMIC data and corrupted it with artificial noise. We randomly extracted 1000 high quality 5 minute excerpts for which both our method and the QRS detector were 100% accurate. We added AWGN noise at different Signal/Noise (SNR) levels to these excerpts and tested both the methods on them.

We carried out two tests. In the first test, we added noise to all m channels, whereas in the second test, we added noise to only $m - 1$ channels.

The results of these tests are presented in Table 2. Our method was able to identify the segment boundaries with reasonable accuracy even in the presence of significant additive noise; further, if any one of the channel is free of noise, the performance is comparable to the performance on the clean raw data.

4.2.2 Transient Corruption

In this experiment we evaluated our algorithm’s performance on data altered, by transient corruption. We randomly extracted 1000 high quality 5 minute excerpts and artificially corrupted them. On each of these excerpts, we corrupt $m - 1$ single parameter signals. For each of these single parameter signals, we randomly selected 5 non-overlapping 1 minute long

Table 2: Summary of the experimental results on the artificially corrupted data.

Average error (ms) on artificially corrupted data				
Type		QRS	WTW	
			all m	any $m - 1$
AWGN (SNR)	20 dB	12.87	0.87	0.0097
	10 dB	188.11	3.27	0.0011
	0 dB	303.48	5.81	0.008
Transient Corruption		387.32	-	2.89
Average segment length is 521 ms.				

regions and applied one of the following types of corruptions: signal interruption, exponential damping, overshooting and clipping, or superimposition of artificial low frequency signals and high frequency signals. The ECG channels and either ABP or PPG were corrupted.

The results are presented in Table 2. Under transient corruption, the average error was 2.89 ms for WTW. This is comparable to that of the data with AWGN at SNR 10 dB on all m parameters. Under severe transient corruption the QRS detector becomes totally unusable.

5 CONCLUSION

We presented a novel online method for temporal segmentation of quasiperiodic multi-parameter physiological signals in the presence of noise and transient corruption. Our method uses Weighted Time Warping to exploit the relationship between the partially correlated channels and the repetitive morphology of the time series.

Our method has a greater constant overhead in computational complexity relative to QRS detection based segmentation algorithm. For a window of length ℓ , our method uses dynamic time warping which runs in $O(\ell^2)$ time and space. A QRS detection based segmentation runs in $O(\ell)$ time. In the examples used in this paper ℓ is less than 500.

Our method is particularly useful when the system suffers noise and transient corruption. For corruptions, we tested our method on few artificial corruptions; but the real world corruptions could be different, and perhaps adversarial. We haven’t tested our method on these specific types of corruptions. In the case of additive noise, AWGN is the most difficult to handle because it spans the entire frequency spectrum and has the highest uncertainty. Our method performs well against AWGN.

We chose ABP, PPG and ECG for testing, because

they are commonly available in the recordings of ICU patients. However, our method should be applicable to any set of correlated physiological signals such as Central Venous Pressure (CVP), Pulmonary Arterial Pressure (PAP) and respiratory signal.

We showed that our method performs as well as an excellent QRS detector on relatively clean ECG data. On 1 hour long test data across 70 records, our method achieved 99.56% accuracy, whereas the QRS detector achieved 99.41% accuracy. When AWGN is synthetically added, the difference in the performance between our method and the QRS detector becomes significant. Our method was able to limit the average error to 5.81 ms when all m parameters were corrupted with AWGN at SNR: 0dB, and to 0.008 ms when $m - 1$ parameters were corrupted at the same SNR. The average error for the QRS detector rose to 303.48 ms when the ECG channel is corrupted with the same AWGN. Similar results were observed on transient corruption.

REFERENCES

- Aboukhalil, A., Nielsen, L., Saeed, M., Mark, R., and Clifford, G. (2008). Reducing false alarm rates for critical arrhythmias using the arterial blood pressure waveform. *Journal of biomedical informatics*, 41(3):442–451.
- Cao, H., Eshelman, L., Chbat, N., Nielsen, L., Gross, B., and Saeed, M. (2008). Predicting icu hemodynamic instability using continuous multiparameter trends. *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE DOI - 10.1109/IEMBS.2008.4650037*, pages 3803–3806.
- Chen, X., Xu, D., Zhang, G., and Mukkamala, R. (2009). Forecasting acute hypotensive episodes in intensive care patients based on a peripheral arterial blood pressure waveform. *Computers in Cardiology, 2009 DOI - UR -*, pages 545–548.
- Deshmane, A. (2009). False arrhythmia alarm suppression using ecg, abp, and photoplethysmogram.
- Friesen, G., Jannett, T., Jadallah, M., Yates, S., Quint, S., and Nagle, H. (1990). A comparison of the noise sensitivity of nine qrs detection algorithms. *Biomedical Engineering, IEEE Transactions on DOI - 10.1109/10.725330*, 37(1):85–98.
- Gleicher, M. and Kovar, L. (2004). Automated extraction and parameterization of motions in large data sets. *ACM transactions on graphics*.
- Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., and Stanley, H. E. (2000). Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220.
- Henriques, J. and Rocha, T. (2009). Prediction of acute hypotensive episodes using neural network multi-models. *Computers in Cardiology*, 36.
- Kohler, B.-U., Hennig, C., and Orglmeister, R. (2002). The principles of software qrs detection. *Engineering in Medicine and Biology Magazine, IEEE DOI - 10.1109/51.993193*, 21(1):42–57.
- Kovar, L. and Gleicher, M. (2004). Automated extraction and parameterization of motions in large data sets. *ACM Transactions on Graphics (TOG)*.
- Li, Q. and Clifford, G. D. (2008). Suppress false arrhythmia alarms of icu monitors using heart rate estimation based on combined arterial blood pressure and ecg analysis. pages 1–3.
- Li, Q., Mark, R., and Clifford, G. (2008). Robust heart rate estimation from multiple asynchronous noisy sources using signal quality indices and a kalman filter. *Physiological measurement*.
- Mark, R. and Shavdia, D. (2007). Septic shock: providing early warnings through multivariate logistic regression models.
- Myers, C., Rabiner, L., and Rosenberg, A. (1980). Performance tradeoffs in dynamic time warping algorithms for isolated word recognition. *Acoustics, Speech and Signal Processing, IEEE Transactions on DOI -*, 28(6):623–635.
- Pan, J. and Tompkins, W. J. (1985). A real-time qrs detection algorithm. *Biomedical Engineering, IEEE Transactions on DOI - 10.1109/10.725330*, BME-32(3):230 – 236.
- Park, A. and Glass, J. (2006). A novel dtw-based distance measure for speaker segmentation. *Spoken Language Technology Workshop, 2006. IEEE DOI - 10.1109/SLT.2006.326807*, pages 22–25.
- Tarassenko, L., Hann, A., and Young, D. (2006). Integrated monitoring and analysis for early warning of patient deterioration. *Medical Applications of Signal Processing, 2005. The 3rd IEE International Seminar on (Ref. No. 2005-1119) DOI -*, pages 64–68.
- Urrusti, J. and Tompkins, W. (1993). Performance evaluation of an ecg qrs complex detection algorithm. *PROC ANNU CONF ENG MED BIOL*.
- Vlachos, M., Hadjieleftheriou, M., Gunopulos, D., and Keogh, E. (2006). Indexing multidimensional time-series. *The VLDB Journal — The International Journal on Very Large Data Bases*, 15(1).
- Zhou, F., Torre, F., and Hodgins, J. (2008). Aligned cluster analysis for temporal segmentation of human motion. *Automatic Face & Gesture Recognition, 2008. FG '08. 8th IEEE International Conference on DOI - 10.1109/AFGR.2008.4813468*, pages 1–7.