32nd Annual International Conference of the IEEE EMBS
Buenos Aires, Argentina, August 31 - September 4, 2010

The effect of lossy ECG compression on QRS and HRV feature extraction

Niall Twomey[1], Noel Walsh[2], Orla Doyle[1], Brian McGinley[2], Martin Glavin[2], Edward Jones[2], W. P. Marnane[1]

Abstract— **This paper describes the performance of beat detection and heart rate variability (HRV) feature extraction on electrocardiogram signals which have been compressed and reconstructed with a lossy compression algorithm. The set partitioning in hierarchical trees (SPIHT) compression algorithm was used with sixteen compression ratios (CR) between 2 and 50 over the records of the MIT/BIH arrhythmia database. Sensitivities and specificities between 99% and 85% were computed for each CR utilised. The extracted HRV features were between 99% and 82% similar to the features extracted from the annotated records. A notable accuracy drop over all features extracted was noted beyond a CR of 30, with falls of 10% accuracy beyond this compression ratio.**

I. Introduction

The ECG signal can provide useful diagnostic information. Currently, most monitoring is performed within the hospital environment. This situation is undesirable, due to the scarcity of expensive hospital resources. In addition, ambulatory monitoring outside the clinical environment provides real life data over a longer time-frame, thereby providing the clinician with more relevant data with which to make a diagnosis.

Ambulatory ECG monitors have existed for some time, with current trends towards achieving a longer operating life and greater processing capabilities at minimal power levels. As information storage and transmission are major consumers of power in embedded devices, data compression will result in improved transmission efficiency, reduced storage requirements, lower power consumption and longer battery life. However, an important consideration is the effect of lossy compression on the fidelity of the signal.

SPIHT compression is a set partitioning coding method that was originally proposed for image compression. It provides good compression performance with a reasonably efficient algorithm. In previous research, Lu and Kim employed SPIHT for ECG signal compression [1]. The algorithm was chosen for use in this paper as it is considered the state-of-the-art method for ECG compression.

QRS detection provides a well-documented platform for diagnosing numerous heart rate conditions, such as arrhythmia, etc [2]. Based on the canonical shape of the well-behaved heart beat very high quality automatic extraction of heart rate variability data has been successfully demonstrated [3]. An equally high performance in QRS extraction for non-standard ECG traces is also important, but due to the diversity of the shape of the ECG trace this is more difficult. For this reason, the MIT-arrhythmia database [4] was used in this study as it provides a rich catalogue of heart beat classes over which the QRS detection will be performed.

In total the MIT-arrhythmia database provides 48 30-minute ECG signal recordings that are sampled at 360 Hz. Annotations are also provided which classify many features of the heart rate signal.

Afonso *et al.*'s [5] ECG beat detection algorithm was used in this study. It was chosen due to its excellent published performance over the MIT-arrhythmia database, and because it has a low subjective comparison with respect to the computational load [3].

This study demonstrates the effect of the SPIHT compression algorithm on QRS detection and feature extraction using the MIT-arrhythmia database.

II. Methods

A. Compression

The effect of lossy ECG compression can be quantified by comparing metrics extracted from the original and reconstructed signals (data that resulted from a lossy compression and subsequent reconstruction will be referred to as lossy data).

The encoding process begins with the discrete wavelet transform being applied to the ECG signal. The SPIHT algorithm is an efficient method of compressing quantised wavelet coefficients to generate an embedded output bitstream. The main principles of the algorithm are based on zero-tree coding, which SPIHT extends with a set partitioning sorting algorithm to process the wavelet coefficients in order of importance[1]. This progressive encoding enables the compression process to be terminated at any point to meet a target bit rate. The same partitioning rules and decoding path are implemented in the decoder, to reverse the process.

The SPIHT compression algorithm was used to compress the ECG traces of the MIT-arrhythmia database with the following compression ratios: $\{2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 25, 30, 35, 40, 50\}$. Once compressed and reconstructed, the lossy data pertaining to each compression ratio was passed to an automatic QRS detection algorithm.

B. QRS detection

Afonso *et al.*'s [5] beat detection algorithm was used for this study. It decomposes the ECG into uniform frequency bandwidths and uses a heuristic detection strategy to fuse decisions from multiple one-channel beat detection algorithms together [5].

This work is supported by Science Foundation Ireland (SFI/07/SRC/I1169), and an enterprise partnership scheme between IRCSET and the TRIL Centre.

[1] Department of Electrical and Electronic Engineering, University College Cork, Ireland. {niallt,orlad}@rennes.ucc.ie, l.marnane@ucc.ie

[2] College of Engineering and Informatics, National University of Ireland, Galway {n.walsh7, brian.mcginley, martin.glavin, edward.jones}@nuigalway.ie

The MIT-arrhythmia database provides beat and non-beat annotations (non-beat annotations are annotations flagging non-beat events, e.g. the start of a rhythm change, etc). The beat annotations are not always localised on the apex of the QRS complex. A QRS complex may be in the form of a spike in the up or downward direction, so a positive and negative peak detection algorithm was executed over a small range about each annotated QRS position (± 5 samples). The local maximum (or minimum) about this range was used as the comparative QRS point.

These fixed QRS points were used as a baseline in order to quantify the accuracy of the automatically extracted QRS points of Afonso *et al.*'s algorithm over the original (non-compressed) dataset. The beat detection algorithm provides a set of data points (times at which a QRS complex was deemed to have occurred). Each point in the extracted QRS complex time array was compared to the annotated data. The constraints of positive QRS complex identification were set with an 88 ms tolerance time window which allowed QRS points localised within the QRS onset and offset to be flagged as correctly extracted QRS points, as discussed by Friesen *et al.* [6]. Each point in the extracted QRS complex time array was compared to the annotated data and if an extracted point was found to be within the tolerance the point was deemed successfully extracted. Additional post-processing was not performed on the extracted QRS points.

C. Lossy analysis

For each quantitative measurement below, the extracted QRS points were compared to the annotated QRS points. The similarity to the results tabulated from the annotated dataset were computed with: $\%\ \text{similarity} = 100\% \times (\text{annotatedValue} - \text{lossyValue})\,/\text{annotatedValue}$.

1) Sensitivity/specificity: Sensitivity (Se) and specificity (Sp) are defined by:

$$\text{Se} = \frac{\text{TP}}{\text{TP} + \text{FN}}; \quad \text{Sp} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{1}$$

where TP is the total number of beats detected correctly (i.e. true positives), FN is the number of false negatives of the algorithm, and FP is the number of false positives of the algorithm.

These metrics gauge the accuracy of the beat detection algorithm. The sensitivity reports the percentage of true beats that were correctly recognised by the algorithm, while the specificity is the percentage of beat detections that are true beats.

The QRS complex times returned by the algorithm for each lossy dataset were then compared to the annotated beat times, and via equation (1) the relative sensitivities and specificities were computed.

2) PRD: The Percentage Root Means Squared Difference (PRD) is a measure of the percentage difference between two signal traces. It is defined by:

$$\text{PRD}(x, \widehat{x}) = 100\% \times \sqrt{\frac{\sum_{n=1}^{N} \left(x(n) - \widehat{x}(n)\right)^2}{\sum_{n=1}^{N} x(n)^2}} \tag{2}$$

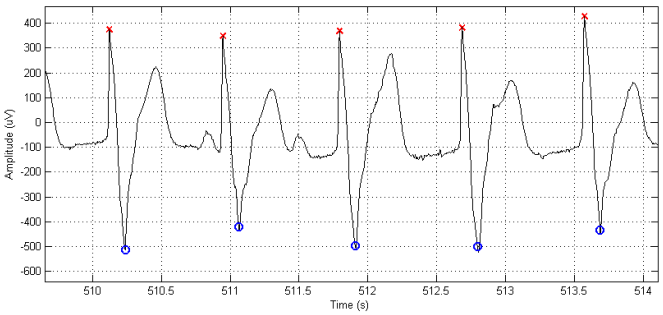

Fig. 1. Location of annotated ($\times$) and automatically extracted ($\circ$) QRS point locations

where x is the original ECG trace, $\widehat{x}$ is the lossy ECG trace and both signal traces are of the same length, N.

3) Time-domain analysis: Two time-domain heart rate variability features were extracted from the Heart Rate Trace ($\text{HRT} = 60 \times \text{fs}/\text{beatIntervals}$):

- Mean heart rate: HR
- Standard Deviation of heart rate: SD

4) Frequency-domain analysis: As the times at which the heart beats are not periodic, traditional frequency-domain analyses of the QRS complex series cannot be used. The Lomb periodogram [7] was adopted to estimate the power spectral density of the QRS beat signal $x(t_n)$ of length N. It is defined by (3).

$$P_x(f) = \frac{1}{2\sigma^2}\left\{ \frac{\left[\sum_{n=1}^{N}(x(t_n)-\overline{x})\cos(2\pi f(t_n-\tau))\right]^2}{\sum_{n=1}^{N}\cos^2(2\pi f(t_n-\tau))} + \frac{\left[\sum_{n=1}^{N}(x(t_n)-\overline{x})\sin(2\pi f(t_n-\tau))\right]^2}{\sum_{n=1}^{N}\sin^2(2\pi f(t_n-\tau))} \right\} \tag{3}$$

Where $\bar{x}$ and σ^2 are the mean and variance of the series. τ makes the series insensitive to time shift, and is defined by (4).

$$\tau = \frac{1}{4\pi f}\arctan\left(\sum_{n=1}^{N}\sin(4\pi f t_n) \,/\, \sum_{n=1}^{N}\cos(4\pi f t_n)\right) \tag{4}$$

The features that are extracted from the PSD were the total respective powers in the low frequency (LF: $0.04 - 0.15$ Hz) and the high frequency (HF: $0.15 - 0.4$ Hz) bands. The ratio LF/HF is a measurement of the sympathetic and parasympathetic response of the nervous system [8].

III. Results

A. QRS detection

The QRS detection algorithm was used over each recording of the MIT-arrhythmia database and the extracted QRS times were compared to those on the manually composed annotation table. The majority of the QRS points were within the 88 ms tolerance window specified by [6].

Figure 1 shows a particular example where the annotated QRS points (shown by the $\times$) and the extracted QRS points

Fig. 2. Sensitivity, specificity and PRD with respect to data volume savings

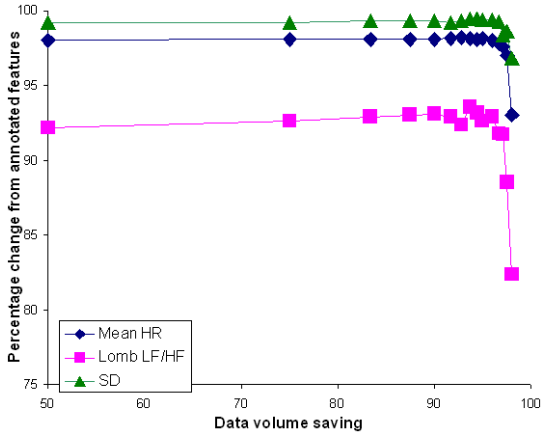

Fig. 3. Mean, standard deviation of heart rate and LF/LF ratio

(shown by the ○) are marked on the same signal trace. The time difference between the × and ○ is longer than the tolerance window so while a beat was detected it was flagged as both a false negative and a false positive. In three (recordings: 107, 119, 121) cases this effect dominated the QRS detection, resulting in very poor sensitivities and specificities (both approximitely 50%) and these recordings were omitted from further analysis. In other recordings this mismatch was also present, and the effect of this on QRS extraction figures will be discussed in the subsequent sections.

Typically, if this detection artefact was present in the analysis of the original dataset, it followed through to the lossy datasets too.

B. Lossy analysis

The results are presented with respect to the Data volume Saving (DS = $100\% \times (1 - \frac{1}{CR})$.

1) Sensitivity/specificity: Afonso *et al.* used the MIT-arrhythmia database annotations in order to compute the performance of their algorithm. The values for sensitivity and specificity (when averaged over each recording in the MIT-arrhythmia database) were given as: 99.59% and 99.56%[5].

The sensitivity and specificity were computed for the extracted QRS points of the algorithm over the original dataset and were computed as 99.086% and 99.133%, see Figure 2. These figures are marginally lower than those computed by Afonso *et al.* due to the points discussed in Section III-A.

Figure 2 plots the sensitivity and the specificity of the algorithms against the percentage saving that the compression ratio provides. A corner point can be seen at a saving of 96.66% (CR = 30). Beyond this point the sensitivity and specificity begin to drop sharply.

2) PRD: The PRD was calculated on a per-recording and per-compression ratio basis. Figure 2 plots 100% − PRD against the data volume savings. As expected the trend of the trace tends to fall away from 100% with a growth in savings.

3) Time-domain analysis: Both time domain features were extracted over one minute epochs.

The heart rate is plotted against the data volume savings in Figure 3. It maintains a close tie (> 97%) to the heart rate that was extracted from the annotated data. As the compression ratio increases this value begins to drop. The percentage similarity drops below 97% at a data saving of 97.5% (CR = 40), and it continues to drop to 93% at a DS of 98% (CR = 50).

The standard deviation of the heart rate averaged over each test on an epoch-by-epoch basis is plotted in Figure 3. A consistent level of percentage similarity is plotted until a data saving of 96.66% (CR = 30) is encountered when it begins to drop.

Both of the time domain features extracted show closest similarities (> 97%) between DS values of 50% and 97% ($2 \leq$ CR ≥ 50). However, the highest similarity was found when no compression ratio was applied to the dataset.

4) Frequency-domain analysis: The Lomb periodogram was used to calculate the ratio between the low frequency band and the high frequency band and is plotted against the data volume savings in Figure 3.

The LF/HF ratio is the most dissimilar of the features to those extracted from the annotated data over all compression ratios (the highest level of similarity, 93%, to the annotated levels is found at a data storage saving of 91% (CR = 25)). This feature also demonstrates the downward trend in performance that has been characteristic of the other features investigated in this study at data savings of above 96% (CR = 30).

IV. Discussion

A. Sensitivity/specificity

Figure 2 plots the sensitivity and specificity of Afonso *et al.*'s QRS detector over each available dataset. The baseline performance of the detector is comparable to the results of [5]. The slight discrepancies can be attributed to incorrect beat localisation. This study did not perform any additional post-processing to the extracted QRS complex times, but with appropriate techniques (such as widening the detection window) the effect shown in Figure 1 could be rectified.

Intuitively, as the storage saving grows, the sensitivity and specificity generally tends to fall. At a DS of 97.5% (CR = 40) the sensitivity and specificity are still over 94% and 96% but when the compression ratio reaches DS of 98% (CR = 50) the sensitivity falls to 85% while the specificity still remains at nearly 93%.

The reason for this drop in sensitivity when the savings reaches 98% (CR = 50) is similar to the reason behind the baseline sensitivity and specificity being slightly lower than [5] (when averaged over each recording): Heart beats are detected but are incorrectly localised on the ECG trace. This effect begins to become increasingly prevalent as the compression ratio approaches ~ 40 after which it begins to dominate the detection results.

It should be noted that the specificity rises after an initial drop from a DS between 75% and 96% (CR between 4 and 30) to 98.24% and 98.36%. The rise in specificity means that the number of false positives drops. However, the sensitivity also drops during this period. Figure 2 demonstrates that beyond a DS of 96% (CR $=$ 30) the sensitivity and specificity begin to drop dramatically.

B. PRD

As the compression ratio increases the PRD increases. In Figure 2, it is interesting to note that the changes in sensitivity and specificity are not influenced by the changes in PRD until a PRD of 8% is reached.

C. HRV features

It was stated previously that no additional post-processing was performed on the results of the QRS detector. It was also observed that Afonso *et al.*'s QRS detector did not always localise the QRS point on a peak of a heart beat.

The heart rate was calculated by averaging the beats over an epoch. An effect of this averaging is that the slight disturbances mentioned above are minimised due to a filtering effect. This makes the heart rate signal trace perform well up to high data saving ratios; the heart rate is the most accurate feature in comparison to the annotation with 93.09% accuracy at a DS of 98% (CR $= 50$).

The Lomb periodogram is, however, sensitive to the time of each QRS point and the difference in the positioning of the QRS peaks might contribute to the lower accuracy of this feature against the reference data. This can be seen by the lower overall accuracy of the algorithm over the data set and the sharp drop after the compression ratio mentioned before.

D. ECG recordings

The MIT-arrhythmia database stores a wide collection of heart beat events. The performance of the beat detection algorithm, over the more "normal" ECG signal traces, is not so adversely affected by the changes in DS or CR, and in some cases (recordings 217, 231) a sensitivity or specificity of 100% is maintained between the DS of 50% to 98% (CR $=$ 2 to 50), with other specific cases (recording 100) demonstrate very slight performance change ($< 0.5\%$).

That every feature trace demonstrates a corner at or near a DS of 96% (CR $=$ 30) indicates that compression ratios beyond this point perhaps shoud not be considered for beat detection and heart rate variability applications where high degrees of accuracy are important.

V. Conclusion

This paper has examined the effect of lossy compression of ECG signals on the performance of QRS detection. Sensitivities and specificities of up to and over 99% can be obtained with Afonso *et al.*'s algorithm with lossy data at a DS of 50% (CR $= 2$). Both sensitivity and specificity remain in the 98% bracket up to a DS of 95% (CR $= 20$) and PRD of 5.8%. The HRV features extracted show a similar trend maintaining high levels ($> 90\%$) of similarity until a SD of 96% (CR $= 25$).

The stable QRS detection results provides the foundation for accurate HRV feature extraction ($> 90\%$ for DS of 97%, CR $= 35$) with the time domain features consistently being more accurate ($> 93\%$ for DS of 98%, or CR $= 50$).

Therefore, if high QRS detection accuracy is necessary for an application, but low volumes of data storage are also required, this paper demonstrates that these can be achieved with the SPHIT compression algorithm and Afonso *et al.*'s beat detection algorithm. Figures 2 and 3 demonstrate that high accuracies and high compression savings are obtained with a CS of 96.6% (CR < 30) and below. Beyond this point compression losses begin to heavily corrupt the sensitivity and specificity recordings.

References

[1] Lu Z., Kim D.Y., Pearlman W.A., "Wavelet compression of ECG signals by the set partitioning in hierarchical trees (SPIHT) algorithm," IEEE Transactions on Biomedical Engineering, vol. 47, 2000, pp. 849–856.

[2] G.D. Clifford, F. Azuaje, and P.E. McSharry, Advanced Methods and Tools for ECG Data Analysis, Artech House, 2006.

[3] Kohler, B.-U.; Hennig, C.; Orglmeister, R.; , "The principles of software QRS detection," IEEE transactions on Engineering in Medicine and Biology Magazine, vol.21, no.1, pp.42-57, Jan.-Feb. 2002

[4] Massachusetts Institute of Technology. MIT-BIH ECG database. Available: http://ecg.mit.edu/.

[5] Afonso, V.X.; Tompkins, W.J.; Nguyen, T.Q.; Shen Luo; , "Filter bank-based ECG beat detection," Engineering in Medicine and Biology Society, 1996. Bridging Disciplines for Biomedicine. Proceedings of the 18th Annual International Conference of the IEEE , vol.3, no., pp.1037-1038 vol.3, 31 Oct-3 Nov 1996

[6] Friesen, G.M.; Jannett, T.C.; Jadallah, M.A.; Yates, S.L.; Quint, S.R.; Nagle, H.T.; , "A comparison of the noise sensitivity of nine QRS detection algorithms," IEEE Transactions on Biomedical Engineering, vol.37, no.1, pp.85-98, Jan. 1990

[7] N. Lomb. Least-squares frequency analysis of unequally spaced data. Astrophysics and Space Science, vol. 39, pp. 447462, 1976.

[8] Malik, M, Camm, AJ (Eds), Heart Rate Variability. Futura Publishing Company, Armonk NY 1995.