Model-based filtering, compression and classification of the ECG

GD Clifford[1,2], A Shoeb[1,2], PE McSharry[3,4,5], BA Janz[1,2,6]

[1]Harvard-MIT Division of Health Sciences & Technology, 45 Carleton St., Cambridge, MA, USA
[2]Massachusettes Institute of Technology, Cambridge, USA
[3]Department of Engineering Science, University of Oxford, Oxford, UK
[4]Mathematical Institute, University of Oxford, Oxford, UK
[5]Centre for the Analysis of Time Series, London School of Economics, London, UK
[6]Beth Israel Deaconess Medical Center, Boston, MA, USA

Abstract

Extensions are presented to a previously described realistic nonlinear model of the electrocardiogram to account for T-wave asymmetry. By fitting the parameters of this model using a nonlinear optimization, we demonstrate that an arbitrary ECG can be modeled and consequently in-band noise can be completely removed. We also show that the fitting procedure effects a compression at a rate of ($\frac{F_s}{3(n+m)}$:1) per beat or ($\frac{\overline{\mathrm{RR}}}{3}\frac{F_s}{(n+m)}$:1), where $\overline{\mathrm{RR}}$ is the reciprocal heart rate, F_s is the sampling frequency, and n is the number of symmetric features (or turning points) and m the number of asymmetric features used to fit the beat morphology. Performance tests show that the algorithm can run in real time on a modern desktop PC. Finally we demonstrate that by clustering the parameters, waveform classification is possible.

1. Introduction

Conventional ECG filters are limited by their generic applicability, in that they use only a vague knowledge of the expected frequency band of interest and use almost no information concerning either the general morphology of an ECG, or a patient specific template. Adaptive filters have been proposed [1, 2] which require another reference signal, or some *ad-hoc* generic model of the signal as an input. In this paper an alternative filtering paradigm is proposed which uses a patient specific model of the ECG, yet requires no prior knowledge of the morphology and only one channel of the ECG.

By fitting a modified version of a previously described model [3] to each beat, and constraining the fit with a time-averaged template, a filtering of each beat is performed. The model consists of a sum of Gaussians centered on each wave of the ECG (P, Q, R, S and T). Each Gaussian is fully specified by three parameters; location in time, amplitude and broadness. Therefore, the representation of the ECG as a series of Gaussians is also a form of (lossy) compression. Finally, the parameters for each beat can be compared to a normal set of parameters and a classification made.

This paper presents the general framework for the methodology described above, illustrated with the fitting of both real and artificial beats in noiseless and noisy conditions. A discussion of compression and classification applications is given and future research directions, including real time considerations.

2. Methods

2.1. The signal model

The general model under consideration has previously been described by the authors in [3, 4] in its application to modeling the ECG, blood pressure and respiratory waveforms. In essence, each feature of the ECG ($PQRS$ & T) is described by a Gaussian with three parameters; the amplitude a_i, width b_i and phase $\theta_i = 2\pi/t_i$ (or relative position with respect to the R-peak). The vertical displacement of the ECG, z, is described by an ordinary differential equation,

$$\dot{z}(a_i,b_i,\theta_i) = -\sum_{i\in\{P,Q,R,S,T^-,T^+\}} a_i\Delta\theta_i e^{\left(\frac{-\Delta\theta_i^2}{2b_i^2}\right)} \quad (1)$$

where $\Delta\theta_i = (\theta - \theta_i)$, the relative phase. Note that no z-offset exists as the model-fit assumes $z = 0$ at the iso-electric level. Numerical integration of this equation using appropriate set of a_i, b_i and θ_i leads to the familiar ECG waveform. In this paper we present a formal method for deriving these parameters on a beat-by-beat basis. The times and angles are specified relative to the position of the R-peak and are detailed in Table 1.

Furthermore, an optional extra parameter has been added to the T feature, denoted by a superscripted $-$ or $+$, to indicate that they are located at values of θ (or t) slightly either side of the original θ_i. By using two sets of $\{a_i,b_i,\theta_i\}$ to represent a particular feature, an asymmetric turning point may be formed. Although this is particularly important for the T-wave on the ECG, it is of negligible importance for the other four features in the ECG. This paper therefore adopts the convention that only 6 features are required for the ECG; P, Q, R, S, T^- and T^+.

Table 1. Initial (superscript zero) and final-fit model parameters (1). $\alpha = \sqrt{\overline{\mathrm{RR}}}$. Random offsets not shown.

Index (i)	P	Q	R	S	T^-	T^+
Time (s)	-0.14$\alpha^{\frac{1}{2}}$	-0.05α	0	0.03α	0.20	0.28
θ_i^0 (rads)	-0.90$\alpha^{\frac{1}{2}}$	-0.31α	0	0.20α	1.25	1.74
a_i^0	1.2	-5.0	30.0	-7.5	1.2	1.2
b_i^0	10	10	10	10	10	10
Time (s)	-0.14$\alpha^{\frac{1}{2}}$	-0.05α	0	0.03α	0.20	0.28
θ_i (rads)	-0.90$\alpha^{\frac{1}{2}}$	-0.31α	0	0.20α	1.25	1.74
a_i	1.25	-4.7	30.1	-6.9	1.65	1.98
b_i	11.6	14.8	13.6	12.1	13.7	11.3

2.2. Fitting parameters to the model

One efficient method of fitting the ECG model described above to an observation $s(t)$, is to minimize the squared error between the $s(t)$ and z. That is, we wish to find

$$\min_{a_i,b_i,\theta_i} \|s(t) - z(t)\|_2^2 \qquad (2)$$

over all six i, with $t_i = 2\pi/\theta_i$. Fortunately, we can analytically integrate (1) to give $z(a_i,b_i,t_i) = \sum_i 2a_i\Delta\theta_i \exp(-\Delta\theta_i^2/2b_i^2)$. Equation (2) can then be solved using an eighteen-dimensional gradient descent in the parameter space. The Matlab function *lsqnonlin.m* performs the required implementation of this nonlinear least-squares optimization.

2.3. Preprocessing and initialization

To minimize the search space for fitting the parameters ($a_i b_i$, and θ_i), a simple peak-detection and time-aligned averaging to form an average beat morphology template is formed over at least the first 60 beats centered on their R-peaks. (The template window is unimportant, as long as it contains all the $PQRST$ features and does not extend into the next beat). Cross correlation is then performed between each beat and the template to remove outliers (with a linear cross-correlation coefficient less than 0.95). (See [5] for details and a justification of this method). If more than 20% of the beats are removed, then another 60 beats are allowed into the average template, and the outlier rejection procedure is re-iterated. When less than 20% of the beats are discarded, another average template is then made of the remaining beats. Peak and trough detection is then performed on this template (using refactory constraints for each wave) to find the relative locations of the turning points in time (and hence the θ_i). T^- and T^+ are initialized $\pm$40ms either side of θ_T. By measuring the heights of each peak (or trough) an estimate of the a_i can also be made. Each b_i is initialized with a value $10 + 5\mu$.where μ is a uniform distribution on the interval [0,...,1]. Each of the a_i and θ_i, shown in table 1, were initialized with random perturbations of μ and 20μ respectively.

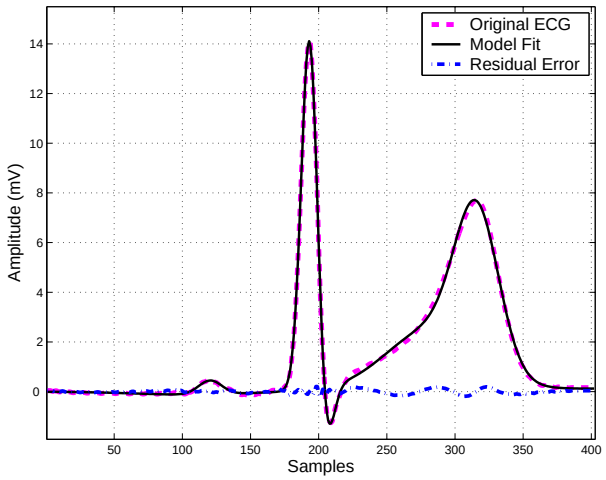

Figure 1. Original ECG, model fit, and residual error.

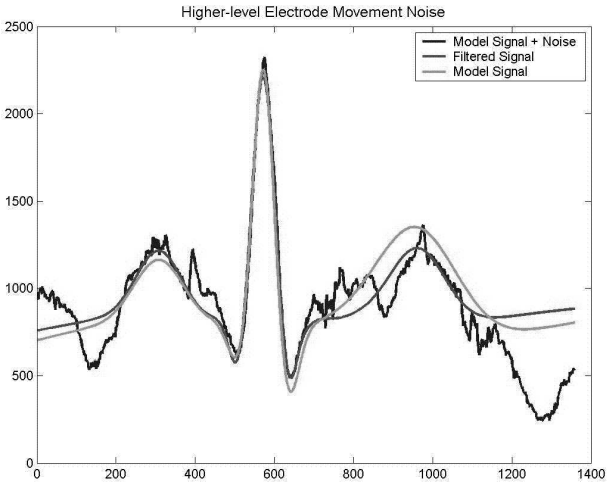

Figure 2. Movement noise filtered by model.

Note that it is important that salient features that one wishes to fit (the P-wave and QRS segment in the case of the ECG) are sampled at a high enough frequency to allow them to contribute sufficiently to the optimization. In empirical tests we found that all ECGs below approximately 512Hz required upsampling (with an appropriate anti-aliasing filter). This corresponds to about 30 sample points in the QRS complex. Using less than 30 samples in a wave can lead to some extremely bizarre fits that fulfill the optimization criteria.

Figure 1 illustrates a real beat (recorded from a V5 lead on the first author), a typical fit to a template of real beat, and the residual error.

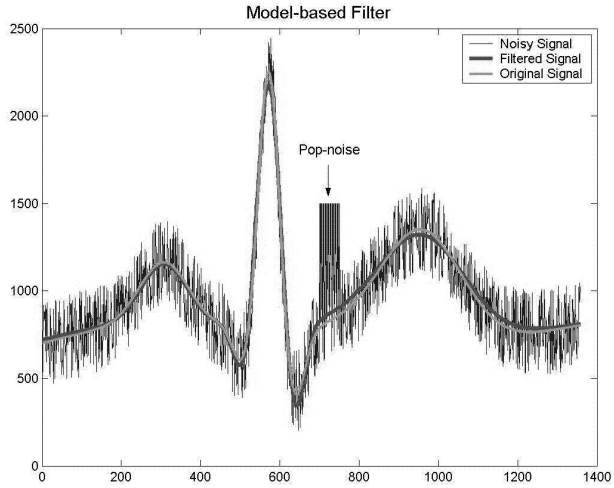

Figure 3. Electrode *pop* noise filtered by model.

3. Results

Fig. 2 illustrates the results of fitting the model to a segment of ECG cleanly recorded from the author, and contaminated by electrode motion noise (taken from the Physionet Noise Stress Test Data Base [6]). Note that despite the significant waveform distortion, the locations of the P, Q, R, S and T peaks match the underlying (uncorrupted signal) to sub-sample precision, even with ($F_s > 1$kHz). Note also that the the error around the iso-electric point and ST-level are negligible in a clinical sense (< 0.1mV, or about 5% to 10% of the QRS amplitude for a sinus beat on a V5 lead). (Amplitudes have been scaled by an arbitrary, but consistent factor). Fig. 3 illustrates the performance of this model-based filtering method in the presence of a nonstationarity. In this case the noise is white and Gaussian, except for a series of impulses (*pop* noise) on the ST-segment with an amplitude half that of the QRS complex. Note that the residual error around the ST-segment is extremely small. In contrast to this, conventional filters are unable to effectively remove such impulses, leading merely to a smoothing of the sharp edges. Although such filters perform the filtering much more rapidly, execution time on a moder desktop PC is less than half a second, making this a possible real-time technique.

4. Discussion

4.1. Filtering

Filtering of the ECG by fitting equation 1 to small segments of the ECG around each QRS-detection fiducial point is an excellent way to provide an idealistic (zero-noise) representation of the morphology which captures much of the clinical information of that beat. In fact, this approach could be generalized to any band-limited waveform with fewer than F_s oscillations per sample. In particular, the signal we are representing does not need to be periodic and is therefore particularly suited to biomedical signals. Since the model is a compact representation of oscillatory signals with few turning points compared to the sampling frequency, it therefore has a band pass filtering effect leading to a lossy transformation of the data into a set of integrable Gaussians distributed over time.

4.2. Compression performance

The fitting procedure effects a (lossy) compression at a rate of ($\frac{F_s}{3k}$:1) per beat or ($\frac{\overline{\mathrm{RR}}}{3}\frac{F_s}{k}$:1), where $\overline{\mathrm{RR}}$ is the reciprocal of the (average) heart rate, F_s is the sampling frequency, and $k = n + 2m$ is the number of features or turning points used to fit the heart beat morphology (with n symmetric and m asymmetric turning points). For a low ECG sampling rate of 128Hz, this translates into a compression ratio greater than 7:1 at a heart rate of 60bpm. However, for high sampling rates ($F_s = 1024$) this can lead to compression rates of almost 57:1. Reducing k from the full representation of $k = 18$ is often appropriate for tasks which require only the QRS complex ($k = 9$) or the ST-segment ($k = 12$) to be analyzed. Of course, high heart rates will reduce this compression unless the dynamic properties of the model are used to encode the heart rate-dependent variations through dynamic shifts in the values of the a_i, b_i and θ_i. For a given segment of τ seconds with an average heart rate of $\frac{60}{\mathrm{RR}}$, the compression ratio rises by a factor $\frac{\tau}{\mathrm{RR}}$. Of course, the model is just an approximation and therefore the compression becomes even more lossy. One should also note that no explicit accounting of abnormal beats has been made in these calculations and a new set of parameters must be derived, possibly for each new abnormal beat encountered in the ECG record.

4.3. Parameter clustering for classification

Although classification of the waveform in terms of the values of a_i, b_i, and θ_i has not been explored in-depth in this paper (that is a full paper in its own right, which will follow in the future), preliminary results indicate that the clustering of normal beats in this 18-dimensional space is *tight* enough to allow separation between beat types (and even artifacts). However, two important questions remain unresolved. Firstly, we must normalize for heart-rate dependent morphology changes. This may prove relatively trivial since the model itself incorporates an heart rate-dependent shift in the $\{a_i, b_i, \theta_i\}$ in the factor, α (see table 1). Secondly, clustering for beat typing is dependent on population morphology averages for a *specific* lead configuration. Not only would different configurations lead

to different clusters in the 18-dimensional parameter space, but small differences in the exact lead placement relative to the heart would cause an offset in the cluster. A method for determining just how far from the standard position the recording is, and a transformation to project back onto the correct position would be required. One possibility could be to use a procedure designed Moody *et al.* [7] for their ECG classifier *Aristotle*. In their approach, the beat clusters are defined in a Karhunen-Loève (KL)-space and therefore an estimate of the difference between the classified KL-space and the observed KL-space is made. Classification is then made after transforming from the observation to classification space on which the training was performed.

5. Conclusion

A technique for simultaneously filtering, compressing and classifying the ECG has been described which can work in real time on a modern desktop PC. By fitting a set of six Gaussians, each specified by three parameters in an ordinary differential equation, and performing a constrained nonlinear optimization, we demonstrate that in-band noise can be completely removed. One advantage of using prior knowledge concerning beat morphology is that a fitting error can be calculated with respect to the model, and thus we have an in-line measure of how well the procedure has filtered the ECG segment. By measuring the distance between the fitted parameters and pre-trained clusters in the 18-dimensional parameter space, classification is possible. However, as with all classifiers, there are two potential problems with the method detailed in this paper. Firstly, if a non-parameterized beat is encountered, it will be considered to be an artifact. Secondly, if an artifact closely resembles a known beat, it will obviously have a good fit to the known beat. Therefore setting tolerances on the acceptable magnitude of the error will prove crucial, and a test over a set of labeled databases is required. This need is closely allied to the requirements of a classifier in general.

It should be noted that the real test of the filtering properties is not the residual error, but how distorted the clinical parameters of the ECG (such as the ST-level and QT-interval) really are and whether they cause a normal beat to be abnormally classified as a normal beat. This analysis will be included in a follow-up paper. The method of producing confidence intervals for a particular fit, or classification is an important step in determining the performance of particular algorithm (see Hughes *et al.* [8]). In-line methods such as these will facilitate the robust interpretation of data and algorithms, reducing the number of false alarms that are triggered. In particular, the smooth nature of the fitted waveform allows for simple and robust detection of clinical features such as the iso-electric point, QT-interval, and ST-level. The residual error from the fitting procedure then provides a confidence measure for the model-derived values of these features.

Our previously published model has been generalized to allow modeling of turning points which exhibit asymmetries (such as the T-wave) by allowing such a feature to be described by two Gaussians. The model as such, can now be used to represent any waveform. However, the model complexity increases considerably for stochastic processes which inherently have many fluctuations compared to the sampling frequency. The main utility of the method detailed in this paper lies in the fact that the model represents smooth oscillations with few turning points compared to the sampling frequency, and therefore has a band pass filtering effect leading to a lossy transformation of the data into a set of integrable Gaussians distributed over time. Each clinical feature of the ECG waveform is represented by a known and limited set of parameters. This allows for a very compact representation of the ECG morphology and makes the description mathematically tractable and completely generalizable to any semi-periodic signal.

References

[1] Biomedical Signal Analysis; A case study approach. IEEE Press, John Wiley & Sons, 2002; 137–176.

[2] Barros A, Mansour A, Ohnishi N. Removing artifacts from ECG signals using independent components analysis. Neurocomputing 1998;(22):173–18.

[3] McSharry PE, Clifford GD, Tarassenko L. A dynamical model for generating synthetic electrocardiogram signals. IEEE Trans Biomed Eng 2003;50(3):289–294.

[4] Clifford GD, McSharry PE. A realistic coupled nonlinear artifi cial ECG, BP, and respiratory signal generator for assessing noise performance of biomedical signal processing algorithms. Proc of SPIE International Symposium on Fluctuations and Noise 2004;5467(34):290–301.

[5] Clifford GD, Tarassenko L, Townsend N. Fusing conventional ECG QRS detection algorithms with an auto-associative neural network for the detection of ectopic beats. In 5^{th} International Conference on Signal Processing. IFIP, Beijing, China: World Computer Congress, August 2000; 1623–1628.

[6] www.physionet.org/physiobank/database/nstdb/.

[7] Moody GB, Mark RG. QRS morphology representation and noise estimation using the Karhunen-Loève transform. Computers in Cardiology 1989;269–272.

[8] Hughes NP, Tarassenko L. Automated QT interval analysis with confi dence measures. Comput in Cardiol 2004;31.

Address for correspondence:

Gari D. Clifford
Harvard-MIT Division of Health Sciences & Technology
45 Carleton St., / Cambridge MA 02139, / USA
gari@mit.edu