A NOVEL FRAMEWORK FOR SIGNAL REPRESENTATION AND SOURCE SEPARATION; APPLICATIONS TO FILTERING AND SEGMENTATION OF BIOSIGNALS

GARI D. CLIFFORD

Harvard-MIT Division of Health Sciences & Technology, Massachusetts Institute of Technology,
77 Massachusetts Av., Cambridge, MA 02139, USA gari@mit.edu
http://alum.mit.edu/www/gari/

A general technique for representing quasi-periodic oscillations, typical of biomedical signals, is described. Using energy thresholding and Gaussian kernels, in conjunction with a nonlinear gradient descent optimization, it is shown that significant noise reduction, compression and turning point location is possible. As such, the signal representation model can be considered a form of correlated source separation. Applications to filtering, modelling and robust ECG QT-analysis are described.

Keywords: ECG; Gradient Descent; Model-based Filtering; Nonlinear Least Squares Optimization; QT-analysis; Segmentation; Source Separation.

1. Introduction

In general, techniques for filtering and segmenting features of interest in a time series use either *ad-hoc* derived basis functions (such as in wavelets[1], principal component[2] or independent component analysis[3,4]) or a very general model of a feature's frequency content (such as in linear or Wiener filtering[5]). The lack of an explicit signal-specific or individual-specific model of the features of interest can severely reduce the ability of an algorithm to isolate the required signal sources from contaminants or to identify onsets and offsets of the features.

The method presented in this paper provides a general framework for deriving models of quasi-stationary signals for robust filtering, compression and segmentation of a signal and for the location of regions of change. As such, the method presented here can be viewed as a type of adaptive filter or as a method for correlated source separation[6] in the time domain. In particular, the approach described in this paper is suited to biomedical signals, that are often characterised by oscillations at specific frequencies and contaminated by *in-band* noise (which is both periodic and statistical)[7]. Identifying the underlying signal is therefore extremely difficult.

Biomedical adaptive filters traditionally use either a generalized knowledge of the spectrum of a signal source [5], an *ad-hoc* reference signal[8] derived from the

2   *Gari D. Clifford*

observable signal features or statistically derived basis functions that are either non-individual specific[2] or their interpretation is unknown[3,4] (since it is a function of the ever-changing data). The approach described in this article differs from previous approaches in three important ways. Firstly, a signal- or subject-specific model of the features to be extracted is created (through an averaging process). Secondly, the model is then adaptively fitted to each of the features in a sequential manner (on-line) using a nonlinear least squares optimization procedure. The output of the model fitting procedure is the extracted source, segmented in time. Thirdly, the use of Gaussians to describe each turning point in the model, allows a statistically meaningful derivation of onset and offset points for subtle smoothly varying changes for high levels of noise (in the original signal).

In this paper, this general framework for filtering and change point localization is described. After presenting the general signal model, an application to QT-interval analysis of the electrocardiogram (ECG) is explored. Finally, applications to signal filtering and modelling in general are discussed.

2. Method

The assumption in the following method is that the time series under analysis is composed of a set of distinct, yet transient (although not necessarily independent) morphologies. Examples of these include the set of features used to classify sleep from the electroencephalogram, (such as *K complexes* and *sleep spindles*[9]), the heart sounds recorded in the phonocardiogram, or the waves in a pulsatile blood pressure waveform[7]. Once a set of general features are identified, a template of each feature can be formed and a mixture of temporally shifted basis functions (such as Gaussians) can be fitted to each major turning point in the signal using an optimization procedure.

The signal model is a modification of a previously described dynamic model[10], where each turning point in a signal is represented by a Gaussian of varying width and amplitude, centered at different points in time. This novel implementation extends the model by adding a new Gaussian for each asymmetric turning point, then adaptively modifying the parameters to fit a distinct observation. Previously, the original model has been adapted to produce realistic arrhythmias[11], abnormal features[12], blood pressure signals[13] and respiration[13]. Here, the concept is generalized to model any signal and provide an automatic method for deriving the model parameters.

2.1. *A General Gaussian Signal Model*

If we assume a transient feature (such as a K complex) is smoothly varying and composed of M symmetric and N asymmetric turning points, then $M + 2N$ Gaussians are required to described the feature (since a Gaussian is symmetric). I.e., an asymmetric turning point requires two Gaussians to be accurately represented. The

segment of the signal z, which describes the feature under analysis is given by

$$z = \sum_{i=1}^{M+2N} \kappa \exp(\Delta t_i^2/2b_i^2) + z_i t \qquad (2.1)$$

where $\Delta t_i = (t - t_i)$, is the relative position of each turning point from the location in time t, of a reference point (fiducial marker), $\kappa = a_i/2b_i$ is a normalization constant (chosen for consistency with the original dynamic model[10]), and the z_i are baseline offset parameters for each of the turning points. The coefficients a_i govern the magnitude of the turning points and the b_i define the width (time duration) of each turning point. The model is therefore fully described by $3(M+2N)$ parameters.

2.2. *Fitting the model to features*

In order to fit Eq. 2.1 to a feature, an approximate template must be constructed. A general method for this is to apply a coarse matched filter (such as cross correlation with a population-independent general template) or an energy thresholding technique (which is common in ECG analysis) to the signal in question. The selection of one technique for this process over another depends on the distribution of the energy of the observation over time. If the signal energy is evenly distributed over time, some *a priori* knowledge of the features must be used to form a simple template for a matched filter.

Fiducial markers can then be located at various points in time that provide time-specific reference markers for each candidate feature (segment of signal). By segmenting the time series around each fiducial point, and performing a temporal average, a first template class is generated. By comparing each candidate feature to the first template class, possible artifacts or patterns belonging to other feature classes can be rejected (using a suitable threshold such as a cross-correlation below 0.9; see Clifford *et al* [14]). The first feature class can then be modified to be the average of the non-rejected individual features (to construct a more specific feature). The rejected candidate can then be averaged to form a second feature class template and the process repeated until the number of possible remaining candidates (which were not included in the previous classes) are below some pre-defined threshold, or the inter-pattern variance between the remaining candidate patterns becomes too high to allow the formation of any more distinct groups.

In the case of an ECG, the first feature class is likely to be a sinus beat (as long as it is the dominant morphology in the time series[14]). Abnormal beats will be rejected and the dominant abnormal beat will become the second feature class. High correlations between the average of this *rejected* set and each member of the set will identify the new members of the set. Rejected beats will cascade down to the next candidate class.

For each template class, an initial model must then be derived. The model order $\mathcal{O} = M+2N$, the number of symmetric plus twice the number of asymmetric turning points in the class. Often, this is a known quantity for most biomedical features,

4   *Gari D. Clifford*

but in some circumstances, an unsupervised method for determining the model order is required. One possible method is as follows: if there are enough feature candidates to form a smooth, low noise template, the number of turning points can be calculated by numerically differentiating the feature and locating the zero crossing points (after allowing for delays in the numerical differentiation function). The degree of asymmetry for each turning point can then be found by squaring the resultant differential and comparing the resultant two peaks (one for the upslope and one for the downslope). If a given pair of peaks are similar in height and width, then the peak is symmetric and only one set of a_i, b_i, and t_i are required for the peak. If the peaks in the squared differential, for a given pair, differ sufficiently (by some predefined threshold that depends on the feature class and signal amplitudes) then the peak is deemed asymmetric and two Gaussians are required to describe the turning point.

It should be noted that this procedure effectively determines the approximate starting points for fitting the model to each feature candidate. However, the height (a_i) and width (b_i) of each Gaussian in the initial model remain to be determined. For most applications, (as long as the t_i are initially limited so that they do not vary significantly) the initialization of the a_i and b_i do not affect the final outcome, and random small values are sufficient. However, in some situations abnormal local minima in the model fitting procedure are possible and the use of an estimate of the width and height of the turning points not only helps to avoid this, but also allows a significant acceleration in the time for fitting each feature candidate.

The residual error between the result of the model fitting procedure (described below) and the original feature provides a facility to reject particular fits. Fig. 1 illustrates the logic flow in this problem. It should also be noted that a classification can be performed by initialising with each possible class (variant of the model) and picking the class with the minimum residual error, or the smallest distance (in parameter space) between a given fit and a cluster center of representative candidates in the same parameter space.

2.3. *Nonlinear least squares gradient descent*

An efficient method of fitting the signal model (Eq. 2.1) to a candidate vector $s(t)$, is to minimize the squared error between s and the model output, z. In other words, we wish to find error

$$\varepsilon_r = \min_{a_i, b_i, \theta_i} \|s(t) - z(t)\|_2^2 \tag{2.2}$$

over all of the $3(M + 2N)$ parameters in the model. Eq. (2.2) can be solved using an $(3M + 6N)$-dimensional nonlinear gradient descent on the parameter space[15]. In general, the problem of multidimensional nonlinear least-squares fitting requires the minimization of the squared residuals of n functions, f_j, in p parameters, x_j,
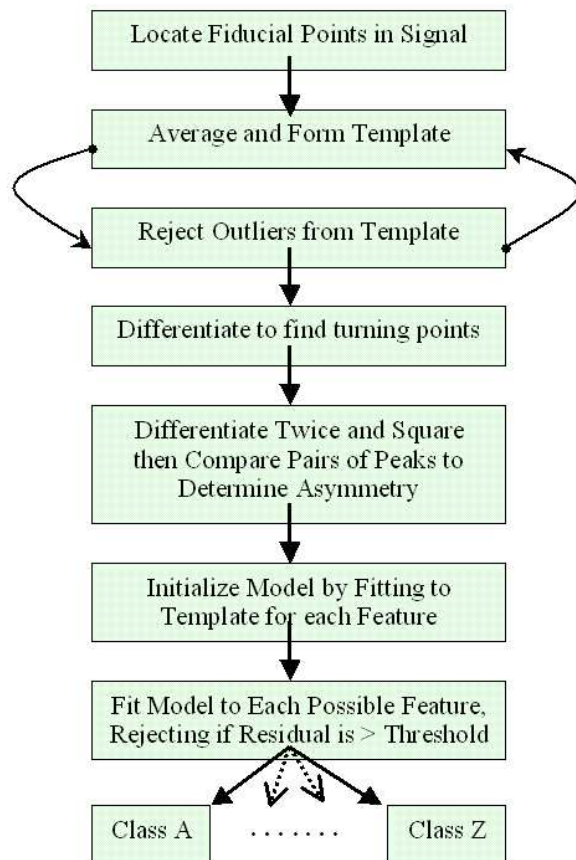
Fig. 1. Generalized procedure for constructing a model for each class and fitting each possible model to each feature candidate.

$$\Phi(x) = (1/2)\sum_{j=1}^{n} f_j(x_1,...,x_p)^2 \qquad (2.3)$$

$$= (1/2)\|F(x)\|^2 \qquad (2.4)$$

All algorithms for achieving this minimization must proceed from an initial guess using the linearization,

$$\psi(p) = \|F(x+p)\| \approx \|F(x) + Jp\| \qquad (2.5)$$

6   *Gari D. Clifford*

where x is the initial point, p is the next step and J is the Jacobian matrix $J_{jk} = df_j/dx_k$. Additional strategies are used to enlarge the region of convergence and include requiring a decrease in the norm $\|F\|$ on each step or using a *trust region* to avoid steps which fall outside the linear régime. This procedure has been implemented in two different libraries; the Gnu Scientific Libraries (GSL) in C, and in MATLAB using the function *lsqnonlin*.
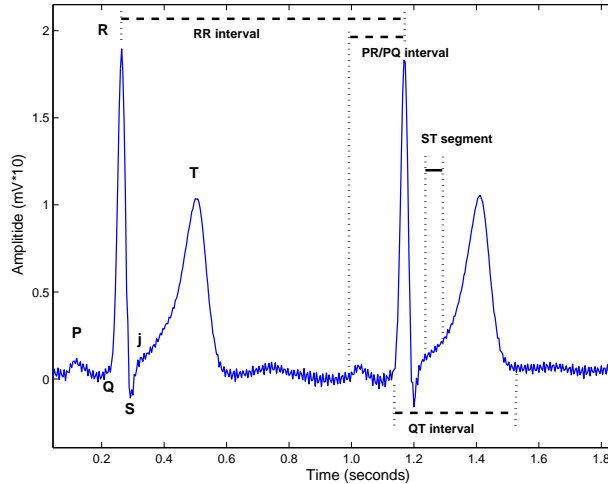

Fig. 2. Typical features in the ECG.

3. Application to the ECG

In this section an example of the above described procedure is applied to the ECG to derive a robust segmentation and derivation of the typical clinical features found in the ECG. The ECG is usually described in terms of its major turning points (labelled P, Q, R, S and T - see Fig. 2). The timing between the onset and offset of particular features of the ECG (referred to as an interval) provides a measure of the state of the heart and can indicate the likelihood of various cardiological conditions. Perhaps the two most important intervals in the ECG waveform are the PR interval and the QT interval (see Fig. 2). The PR interval is defined as the time from the start of the P wave to the start of the QRS complex, and corresponds to the time from the onset of atrial depolarization to the onset of ventricular depolarization. The QT interval is defined as the time from the start of the QRS complex (ventricular depolarization) to the end of the T wave (ventricular repolarization). Since the QT interval corresponds to the total duration of electrical activity (both depolarization and repolarization) in the ventricles, changes in the QT interval are currently the gold standard for evaluating the effects of drugs on ventricular repolarization. Accurate measurement and assessment of the QT interval is therefore extremely

important for the assessment and validation of new drugs in clinical trials. Unfortunately, a precise mathematical definition of the end of the T wave does not exist and therefore T wave end measurements are inherently noisy or subjective and QT interval measurements often suffer from relatively high inter- and intra-expert (or algorithm) variability. In particular, the smooth, almost asymptotic, transitions of the end of the T wave into the noisy baseline make locating a single stable end point extremely difficult, if not impossible.

The method presented in this paper resolves these issues by characterizing the changes of the waveforms in terms of Gaussians, which have a mathematical (probabilistic) definition for their initial and terminal points. Furthermore, the use of a model with no explicit noise variable, results in a waveform that is smooth and noise-free, and therefore amenable to noise-sensitive processes such as differentiation.

3.1. *The ECG model*

The ECG typically contains 5 turning points (P, Q, R, S and T), which are all approximately symmetric except for the T wave (at low to medium heart rates). Therefore, 18 parameters (6 Gaussians) are required to accurately describe the ECG with Eq. 2.1, which can be rewritten as:

$$z = \sum_{i \in \{P,Q,R,S,T^-,T^+\}} (a_i/2b_i) \, e^{\frac{\Delta t_i^2}{2b_i^2}} + z_i t \qquad (3.1)$$

where T^- and T^+ are the two Gaussians used to describe the asymmetric T wave. (This variant of the model was first proposed by the author in[16].) The application to the localization of the PR and QT intervals is a simple extension of this method.

For any given ECG lead, a QRS detector[17] is used to locate the fiducial points (e.g. R-peaks) of each beat. Each beat is then segmented so that no trailing T wave enters in the front of a window or any following beat's P wave enters into the latter part of the segmenting window. In practice this involves calculating the minimum RR interval (highest instantaneous heart rate) and using the Bazzet or Fridericia correction factor[18] (QTc = αQT where $\alpha = $ RR$^{-1/2}$ or RR$^{-1/3}$ respectively) to calculate the longest possible window from the R-peak. Extreme QTc elongation can vary between QT$_{max}$ = 0.44s to QT$_{max}$ = 0.9s[19], and therefore a window of QTc$_{max}$/RR$_{min}^{1/3}$ (of up to 0.9s) extending from the Q-onset is appropriate. The front end of the window can be calculated in a similar manner, with the PR interval being modulated by the local RR interval[20] so that PR$_{max}$/RR$_{min}^{-1/3}$ extending from the start of the window to the Q point is appropriate (with a maximum time of about 200ms from the P-onset to the R-peak). The window in which a fit is performed is therefore asymmetric around the R-peak. The application of the fitting procedure to a typical ECG can be seen in Fig. 3, which illustrates the initial conditions in the upper left frame, four intermediate steps, and the final fit in the lower right frame.
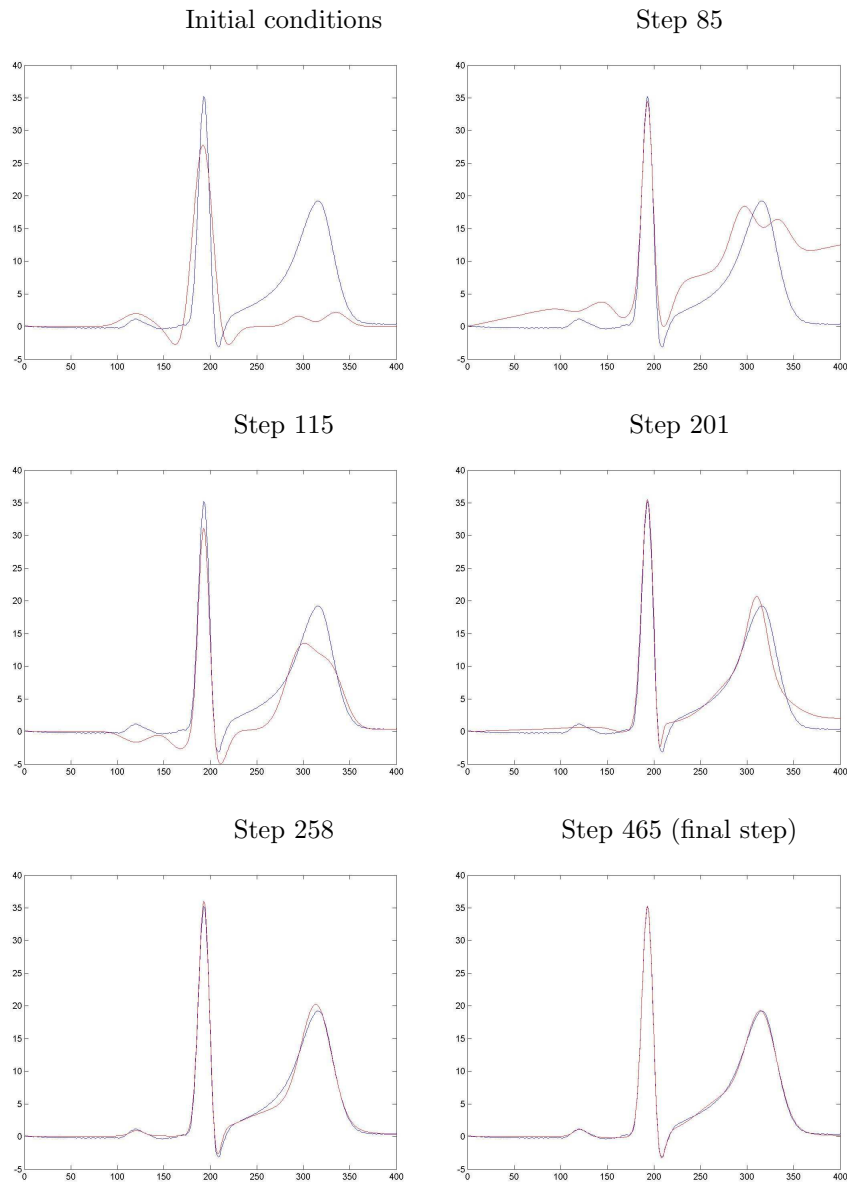
8   *Gari D. Clifford*

Initial conditions

Step 85

Step 115

Step 201

Step 258

Step 465 (final step)

Fig. 3. Six different steps in the fitting procedure. The target waveform (the dark blue line) does not vary from step to step whereas the Gaussians being fitted vary in location, amplitude and width from step to step.
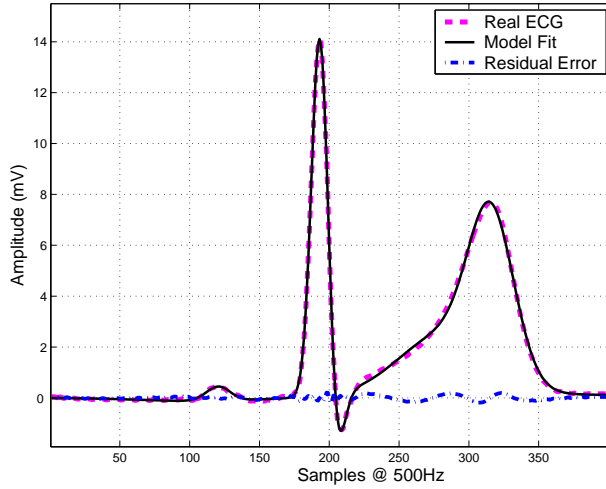
Fig. 4. Original ECG, nonlinear model fit, and residual error.

To minimize the search space for fitting the parameters (a_i, b_i, and t_i), a simple peak-detection and time-aligned averaging technique is performed to form a low-noise template from the average beat morphology using approximately the first 60 beats centered on their R-peaks. The exact number of beats in the template is a trade-off between adapting to the heart rate-induced nonstationarities yet including sufficient beats to average out noise and morphological differences. Furthermore, the template window length is unimportant, as long as it contains all the $PQRST$ features and does not extend into the next beat). This method, including outlier rejection is detailed in Clifford *et al* [14].

T^- and T^+ are initialized ±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, were initialized with random perturbations of μ and 20μ respectively. Fig. 4 illustrates an example of a template ECG, the result of the model-fit, and the residual error, ε_r.

3.2. *QT analysis; waveform boundary localization*

For QT analysis, it is necessary to locate the onset of the Q wave and the offset of the T wave. In the presence of noise, this is an extremely difficult problem. However, using the model-fitting approach, the resultant ECG is noise-free and therefore thresholding on the differential of the signal is sufficient to localize the QT-interval. Ideally one would choose the points where $dz/dt = 0$ (just before θ_Q and just after θ_{T+}, parameters we know precisely from the model). In practice, finite sampling rates mean that we must choose a sampling rate-dependent threshold ϵ. (10^{-4} is sufficient for a normalized ECG sampled at 256Hz.) Fig. 5 illustrates this process;

10   *Gari D. Clifford*

QRS onset and T wave end points are marked with a cross and vertical line. The differential, dz/dt, is shown as a dotted line. Fig. 6 illustrates the same process for a series of 256 beats, with a histogram of all the resultant QT interval distributions.

In fact, clinicians rarely choose the lowest gradient point after the T wave as the offset since inherent noise in the ECG makes this impossible. The T wave end is usually placed by experts about 40 to 60ms earlier than the minimal gradient. This *clinical* offset, which can be adjusted for, is a function a_{T-}, b_{T-}, a_{T+}, a_{T+}, $\theta_{T+} - \theta_{T-}$, sampling frequency and amplitude resolution (since all these features affect the gradient of the T wave downslope that clinicians use to extrapolate the T wave end-point).
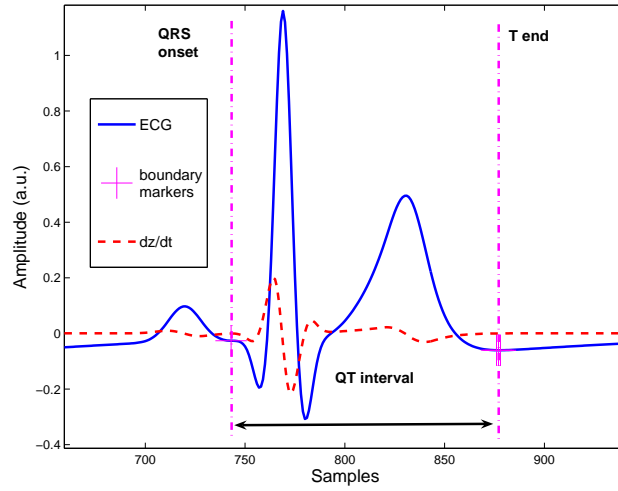

Fig. 5. ECG with QRS onset and T wave end point marked (+). The differential of the ECG (dz/dt) is also shown. Waveform boundaries are located where $dz/dt < \epsilon$.

3.3. *Translation of parameters back into the dynamic model*

Once the parameters of the model have been fitted to an observation, the dynamic equations of motion for a synthetic ECG[10] can be used to generate a realistic representation of a particular individual's ECG. The equations are given by a set of three ordinary differential equations

$$\dot{x} = \rho x - \omega y,$$
$$\dot{y} = \rho y + \omega x,$$
$$\dot{z} = -\sum_{i \in \{P,Q,R,S,T\}} a_i \Delta\theta_i \exp(-\Delta\theta_i^2/2b_i^2) - (z - z_0), \qquad (3.2)$$

where $\rho = 1 - \sqrt{x^2 + y^2}$, $\Delta\theta_i = (\theta - \theta_i) \bmod 2\pi$, $\theta = \mathrm{atan2}(y, x)$ and ω is the angular velocity of the trajectory as it moves around a limit cycle. Baseline wander can be
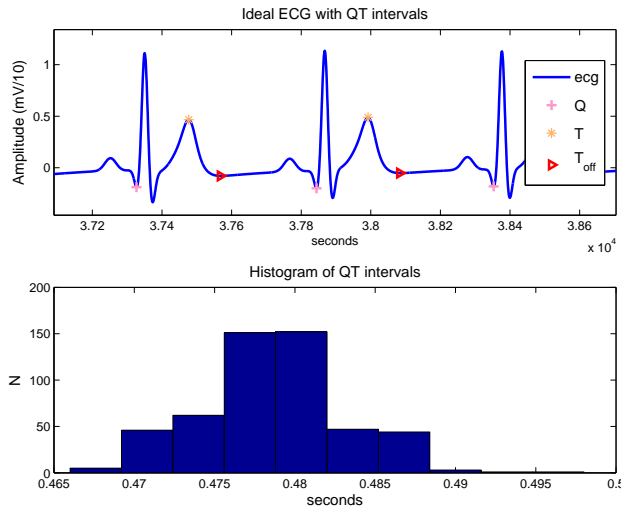
Fig. 6. ECG with QT intervals marked. The lower plot illustrates the distribution of a set of QT intervals for 256 consecutive beats generated from the model.

Table 1. Morphological parameters of the ECG model with heart-rate dilation factor $\alpha = \sqrt{h_{\mathrm{mean}}/60}$.

Index (i)	P	Q	R	S	T$^-$	T$^+$
θ_i	$-\alpha^{\frac{1}{2}}70^o$	$-\alpha 15^o$	0^o	$\alpha 15^o$	$\alpha^{\frac{1}{2}}83^o$	$\alpha^{\frac{1}{2}}90^o$
a_i	0.8	-5.0	30.0	-7.5	$\alpha^{3/2}/2$	$3\alpha^{3/2}/2$
b_i	0.2α	0.1α	0.1α	0.1α	$0.4\alpha^{-1}$	0.2α

introduced by coupling the baseline value z_0 in (3.2) to the respiratory frequency f_2 using

$$z_0(t) = A\sin(2\pi f_2 t), \qquad (3.3)$$

where $A = 0.15$ mV.

These equations of motion given by (3.2) can be integrated numerically using a fourth order Runge-Kutta method[21] with a time step $\Delta t = 1/f_{int}$ where f_{int} is the internal sampling frequency and must be an integer multiple of the output frequency f_s. The value of f_s or f_{int} can often be important and a resampling step is sometimes required. For the ECG, f_s or f_{int} should be at least 500Hz[12]. Fig. 7 illustrates the result of this process using the parameters detailed in table 1. Note that the heart rate *compression* factor α is used on all the parameters (except a_P through a_S). Note in particular that the T wave is adjusted to increase asymmetry and amplitude as the heart rate rises.
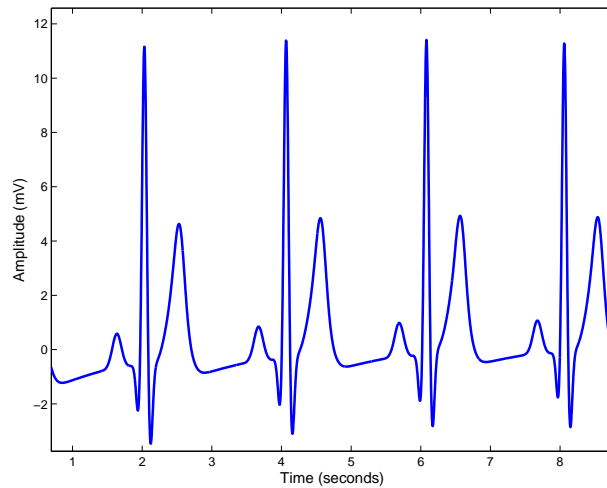
Fig. 7. New artificial ECG after fitting to real data. Note the asymmetric T wave.
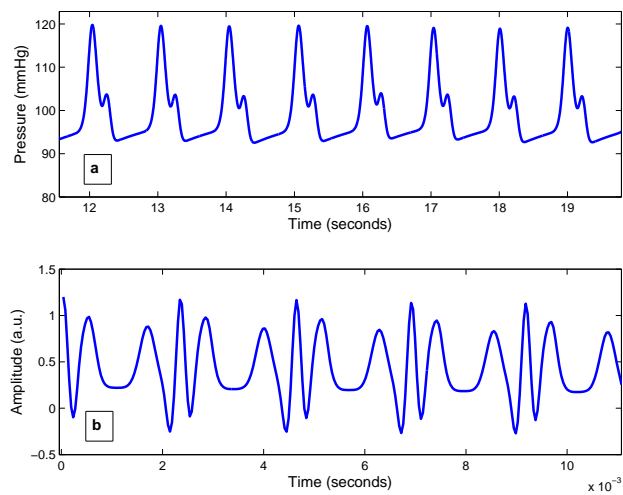

Fig. 8. Examples of possible morphologies from the fitted dynamical model (Eq. 3.2). a) A blood pressure waveform. b) An asymmetric sinc-like function with highly damped tails.

It should be noted of course, that the model is not confined to the representation and reproduction of ECG signals, and can be used for any signal that can be reasonably fitted to a sum of Gaussians (or any other specified set functions). Fig. 8a illustrates an example of a blood pressure waveform, and 8b illustrates a more general pulse train, (a series of asymmetric damped sinc-like functions).

4. Conclusions

The application of a nonlinear gradient descent optimization of a generic signal model to biomedical signals has been shown to provide a framework for providing a noise-free representation of a signal's important components, resulting in an in-band filtering process. The choice of Gaussian kernels allows the identification of statistically meaningful changes in the behaviour of a signal. By simply locating the tails of a Gaussian wave-packet, the onset and offset of distinct features is possible. In particular, the application is demonstrated for the stable location of the QT-interval in the ECG under noisy conditions.

The important difference in this approach in comparison to previous approaches is that the adaptive filter uses an accurate signal- or subject-specific model of the underlying morphology, leading to a virtually distortion-free filter. In fact, the use of the model-fitting procedure can be viewed as a robust filter that rejects in-band noise, yet does not suffer from many of the problems encountered in other filtering procedures (such as ringing).

Since the technique in this paper provides an accurate method for signal representation and noise extraction, it therefore provides an accurate method for analysing the data in the signal which is conventionally considered noise. In particular, the removal of the P, QRS and T waves from an ECG provide an excellent method for measuring high frequency QRS data (such as late potentials). Furthermore, the number of iterations required to fit the model and the error in the final fit, provide an alternative quantification of the noise level in the signal.

Two other applications of this technique are also immediately apparent. Firstly, the representation of the signal in terms of only 18 parameters provides a novel method of compression. Secondly, the distance of the 18 parameters (or a subset of them) from a population cluster for each possible feature, provides a means of classification.

Of course, the use of Gaussian kernels is not a unique choice, and any mixture of mathematical functions can be used to fit a sum of functions to the observation. In fact, other choices of function for certain parts of a feature may yield a better fit (i.e. one that is more robust to noise, with a lower residual error, faster convergence times in fewer operations, using fewer model parameters). Comparisons with Gaussian wavelets are obvious, but distinct differences do exist. For example, the Gaussian wavelet is actually the derivative of a Gaussian and hence the probabilistic interpretation is less clear. Furthermore, rather than linearly convolving the function with the entire time series over many scales, a restricted nonlinear fit

14   *Gari D. Clifford*

is performed at specific locations. It should also be noted that the basis functions necessary for the approach presented in this paper have only one restriction; they must be everywhere continuous (and hence require infinite tails) to avoid discontinuities in the resultant fit. However, the description of the features in terms of Gaussian curves, which have exact (probabilistic) interpretations for starting and end points, allows for a smooth and accurate noise-free representation of a signal with precise boundaries (such as the QT-interval on the ECG). This represents a significant improvement on existing methods for localizing such boundaries, which are generally disturbed by the inherent noise or artefacts in the observation due to filtering processes or exogenous influences.

References

1. N. Nikolaev, Z. Nkolov, A. Gotchev, and K. Egiazarian. Wavelet domain Wiener filtering for ECG denoising using improved signal estimate. In *Proc. of ICASSP '00, IEEE Int. Conf. on Acoustics, Speech, and Sig. Proc.*, volume 6, pages 3578–3581, 5-9 June 2000.
2. G.B. Moody and R.G. Mark. QRS morphology representation and noise estimation using the karhunen-loeve transform. *Computers in Cardiology*, pages 269–272, 1989.
3. AK Barros, A. Mansour, and N. Ohnishi. Removing artifacts from ECG signals using independent components analysis. *Neurocomputing*, (22):173–18, 1998.
4. T. He, G. D. Clifford, and L. Tarassenko. Application of ICA in removing artefacts from the ECG. *Neural Comput. & Applic.*, 15(2):105–116, 2006.
5. *Biomedical Signal Analysis; A case study approach*, chapter 3.5, pages 137–176. IEEE Press, John Wiley & Sons, 2002.
6. P.D. Wirawan and H. Maitre. Multi-channel high resolution blind image restoration. In *International Conference on Acoustics, Speech, and Signal Processing*, volume 6, pages 3229–3232. IEEE, Mar 1999.
7. R.M. Rangayyan. *Biomedical Signal Analysis : A Case-Study Approach*. IEEElsevier Press, 2002. Series on Biomedical Engineering.
8. P. Laguna, R. Jane, O. Meste, P.W. Poon, P. Caminal, H. Rix, and N.V. Thakor. Adaptive filter for event-related bioelectric signals using an impulse correlated reference input. *IEEE Trans. on Biomedical Engineering*, 39(10):1032–44, 1992.
9. P. Lavie. *The echanted world of sleep*. Yale University Press, New Haven, 1996.
10. P. E. McSharry, G. D. Clifford, and L. Tarassenko. A dynamical model for generating synthetic electrocardiogram signals. *IEEE Trans. Biomed. Eng.*, 50(3):289–294, 2003.
11. J. Healey, G. D. Clifford, L. Kontothanassis, and P. E. McSharry. An open-source method for simulating atrial fibrillation using ECGSYN. *Computers in Cardiology*, 31:425–427, 2004.
12. G. D. Clifford and P.E. McSharry. Method to filter ECGs and evaluate clinical parameter distortion using realistic ECG model parameter fitting. *Computers in Cardiology*, 32, 2005.
13. G. D. Clifford and P. E. McSharry. A realistic coupled nonlinear artificial ECG, BP, and respiratory signal generator for assessing noise performance of biomedical signal processing algorithms. *Proc of SPIE International Symposium on Fluctuations and Noise*, 5467(34):290–301, 2004.
14. G. D. Clifford, L. Tarassenko, and N. Townsend. 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*, pages 1623–1628, Beijing, China, August 2000. IFIP, World Computer Congress.

15. J.J. Moré. *Numerical Analysis*, chapter The Levenberg-Marquardt Algorithm. Implementation and Theory, pages 105–132. Springer-Verlag, New York, 1978.

16. G. D. Clifford, A. Shoeb, P. E. McSharry, and B. A. Janz. Model-based filtering, compression and classification of the ECG. *International Journal of Bioelectromagnetism*, 7(1):158–161, 2005.

17. B.-U. Khöler, C. Hennig, and R. Orglmeister. The principles of software QRS detection. *Engineering in Medicine and Biology Magazine, IEEE*, 21(1):42–57, Jan/Feb 2002.

18. Anthony A. Fossa, Todd Wisialowski, Anthony Magnano, Eric Wolfgang, Roxanne Winslow, William Gorczyca, Kimberly Crimin, and David L. Raunig. Dynamic Beat-to-Beat Modeling of the QT-RR Interval Relationship: Analysis of QT Prolongation during Alterations of Autonomic State versus Human Ether a-go-go-Related Gene Inhibition. *J Pharmacol Exp Ther*, 312(1):1–11, 2005.

19. A. Koçer, Z. Aktürk, E Maden, and A. Taşçi. Orthostatic hypotension and heart rate variability as indicators of cardiac autonomic neuropathy in diabetes mellitus. *European Journal of General Medicine*, 2(1):5–9, 2005.

20. R. Shouldice, C. Heneghan, P. Nolan, P.G. Nolan, and W. McNicholas. Modulating effect of respiration on atrioventricular conduction time assessed using PR interval variation. *Med Biol Eng Comput.*, 40(6):609–617, Nov 2002.

21. W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. *Numerical Recipes in C*, chapter 13. Cambridge University Press, 2^{nd} edition, 1992.