Automated QT Interval Analysis with Confidence Measures

NP Hughes, L Tarassenko

Department of Engineering Science, University of Oxford, UK

Abstract

Current automated methods for QT interval analysis suffer from the absence of a confidence value in the resulting measurements. In this paper we present a new approach to QT interval analysis, which produces both a segmentation of the ECG together with an associated measure of confidence in the segmentation. The method is based on a novel hidden Markov model architecture, which is designed to prevent unrealistic segmentations. We utilise the probabilistic nature of the model to derive a confidence measure based upon the log likelihood of the ECG waveform under consideration. The method is demonstrated on an ECG signal containing an ectopic beat, and an ECG contaminated by muscle artefact noise.

1. Introduction

The accurate measurement and assessment of the QT interval is one of the most important problems in ECG analysis. QT interval measurements are particularly significant since a prolonged QT interval is a good indicator of long QT syndrome (LQTS). This is a potentially fatal condition which can give rise to a very fast, abnormal heart rhythm known as *torsade de pointes* [1]. When this occurs the heart cannot beat effectively and the result is generally a sudden loss of consciousness followed by cardiac death.

In the case of clinical drugs trials, the accurate measurement of the QT interval is especially important. In particular, changes in the QT interval are currently the gold standard for evaluating the effects of drugs on ventricular repolarization. In practice, such measurements are carried out manually by specially trained ECG analysts. This is an expensive and time consuming process, which is susceptible to mistakes by the analysts and provides no associated degree of confidence in the measurements. For this reason, much effort has been put into developing automated methods that can accurately and effectively measure the QT interval in ECG waveforms [2].

Currently however, no automated system can achieve the same level of accuracy as an expert ECG analyst. In particular, unusual waveform morphologies (such as those caused by ectopic beats) coupled with the various noise processes which affect the ECG (such as muscle artefact and baseline wander), often result in unreliable QT interval measurements by automated techniques.

In this paper we propose a new approach to automated QT interval analysis, which produces *both* a segmentation of the ECG *and* an associated degree of confidence in the segmentation. Such confidence measures can be used to assess the novelty of the waveform under consideration, and thus to determine a suitable threshold for rejecting QT interval measurements which are deemed to be unreliable.

2. Probabilistic modelling of ECG

Probabilistic modelling offers an attractive framework for the problem of ECG segmentation. In particular, probabilistic models utilise the statistical characteristics of the ECG waveform features during the segmentation process. In addition, the models can take advantage of the sequential nature of the waveform features (i.e. QRS complex follows Baseline which follows P wave etc.) to improve the robustness of the segmentation. Finally, probabilistic models provide "confidence measures" which can be used to assess the quality of the segmentation.

Figure 1 shows a hidden Markov model (HMM) for ECG segmentation. The model architecture is comprised of a "hidden" state sequence (indicated by the clear nodes), which is stochastically related to an observed signal (indicated by the shaded nodes). For ECG segmentation, the hidden state s_t corresponds to the particular waveform feature which is active at time t, and the observed signal sample O_t corresponds to the associated ECG sample.

An HMM is parameterised by an initial state distribution π, a state transition matrix A, and a set of observation

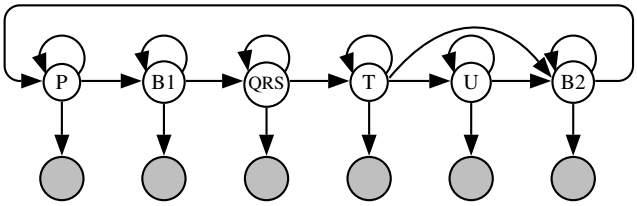

Figure 1. A hidden Markov model for ECG segmentation

densities b_i (for each state i). Once the parameters of the model have been learnt, the most probable state sequence (i.e. the waveform segmentation) for a given ECG signal can be inferred through the use of the Viterbi algorithm [3].

A significant limitation of the standard HMM is the manner in which it models state durations. For a given state i with self-transition coefficient a_{ii}, the probability mass function for the state duration d is a geometric distribution, given by:

$$p_i(d) = (a_{ii})^{d-1}(1 - a_{ii}) \quad (1)$$

For the waveform features of the ECG signal, this geometric distribution is inappropriate. In particular, the distribution naturally favours state sequences of a very short duration. Conversely, real-world ECG waveform features do not occur for arbitrarily short durations, and there is typically a minimum duration for each of the ECG features. In practice this "mismatch" between the statistical properties of the model and those of the ECG results in segmentations that are often unrealistic. These segmentations are characterised by the model incorrectly inferring QRS complexes and T waves of a very short duration [4].

Unfortunately, such segmentations significantly impact upon the reliability of the automated QT interval measurements produced by the model. Thus, in order to make use of the model for automated QT interval analysis, the robustness of the segmentation process must be improved. This can be achieved by incorporating *duration constraints* into the HMM architecture.

More precisely, each duration constraint takes the form of a number specifying the minimum duration for a particular state in the model. For example, the duration constraint for the T wave state is simply the minimum possible duration (in samples) for a T wave. Such values can be estimated in practice by examining the durations of the waveform features for a large number of labelled ECG waveforms.

Once the duration constraints have been chosen, they are incorporated into the model in the following manner: for each state k with a minimum duration of $d_{min}(k)$, we augment the model with $d_{min}(k) - 1$ additional states directly preceding the original state k. Each additional state has a self-transition probability of zero, and a probability of one of transitioning to the state to its right. Thus taken together these states form a simple *left-right* Markov chain.

The most important feature of this chain is that the parameters of the observation density for each state are identical to the corresponding parameters of the original state k (this is known as "tying"). Thus the observations associated with the d_{min} states identified with a particular waveform feature are governed by a single set of parameters (which is shared by all d_{min} states).

Figure 2 shows part of a hidden Markov model with built-in duration constraints for ECG segmentation. In practice we have found that the segmentations of ECG waveforms and the corresponding QT interval measurements are significantly more robust than those of a standard HMM.

3. Confidence measures

Given the segmentation of an ECG waveform by the model, we would like to assess the degree of *confidence* we should have in the segmentation (and thus the corresponding QT and PR interval measurements). We can view such a confidence measure as quantifying the "normality" of the waveform under consideration. The intuition behind this view is that we should have more confidence in the segmentations of ECG waveforms which are similar to those the model was trained on, compared with the segmentations of waveforms which are markedly different from those in the training data.

More formally, we can define the confidence measure for a given ECG waveform as the *log likelihood* of the waveform (under the model). Thus, if an ECG waveform begins at time sample t_1 and ends at time sample t_2, then the corresponding confidence measure θ is given by:

$$\theta = \log\, p(O_{t_1} \cdots O_{t_2} \mid \lambda) \quad (2)$$

where λ are the parameters governing the model. This log likelihood value can be computed efficiently for hidden

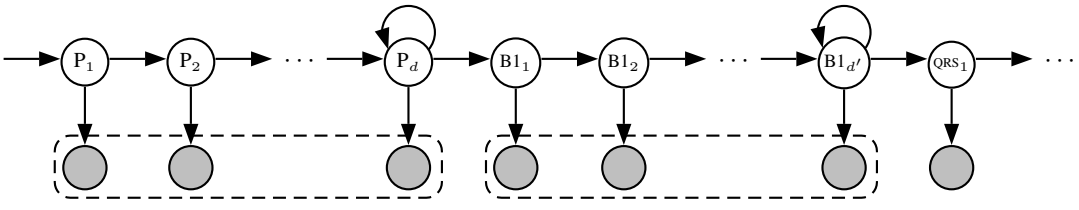

Figure 2. A hidden Markov model with built-in duration constraints for ECG segmentation. The dashed boxes indicate *tied* observation densities.

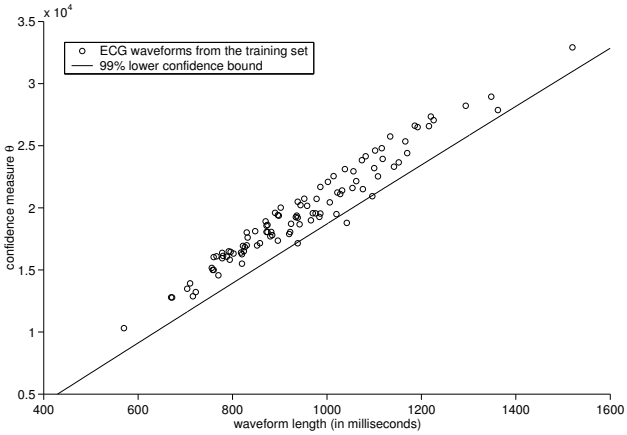

Figure 3. Confidence measure against waveform length for each of the ECG waveforms in the training set.

Markov models using the standard forwards-backwards recursions [3].

When assessing the confidence measure for an ECG waveform in practice, we must consider *both* the log likelihood value for the waveform (as defined by equation 2) and the length of the waveform (i.e. $t_2 - t_1$). This is necessary because the log likelihood of the signal is proportional to both the signal morphology *and* the length of the signal itself. Hence longer signals naturally have a greater log likelihood value than shorter signals. Thus when assessing the confidence measure for a particular ECG waveform, we must determine if the log likelihood value is "normal" for the given waveform length.

In order to determine the range of confidence measures for "normal" ECG waveforms, we first trained a hidden Markov model with built-in duration constraints using a data set of 100 clean ECG waveforms. The data set consisted of a variety of different ECG waveforms (in terms of QRS complex and T wave shape) measured over a range of different heart rates (from 39bpm to 107bpm). Prior to training, each ECG signal was first *encoded* using the undecimated wavelet transform [5]. Once the model had been trained, we then evaluated the confidence measure for each waveform in the data set.

Figure 3 shows a plot of the confidence measures for the waveforms against the corresponding waveform lengths[1]. As noted previously, the confidence measure varies in proportion to the length of the waveform under consideration. In addition, for a given waveform length, there is generally a spread of possible values for the confidence measure.

Using standard regression techniques [6], we fitted a linear regression to the data points and evaluated the lower confidence bound at the 99% level of significance (indicated by the solid line in figure 3). This lower bound thus provides a suitable threshold which can be used to determine if ECG waveforms segmented by the model are sufficiently *novel* as to cast doubt on the reliability of the associated QT and PR interval measurements.

4. Results

To evaluate the effectiveness of the proposed scheme for detecting unreliable segmentations, we used the trained model to segment an ECG signal containing an ectopic beat and an ECG contaminated with muscle artefact noise. Both signals were not present in the original training data.

Figure 4(a) shows the ectopic ECG. The signal consists of an ectopic beat which affects the two central waveforms (beats 2 and 3), and two "normal" waveforms (beats 1 and 4) which lie either side of the ectopic beat. The QT intervals (as determined from the model segmentations) for beats 1 and 4 are 398ms and 404ms respectively, whereas the QT interval for beat 2 is 644ms. Figure 4(b) shows the confidence measures for the four waveforms, together with the 99% lower confidence bound (for the training data). Both waveforms affected by the ectopic beat lie significatly below the confidence threshold, and the normal waveforms lie above this threshold. Thus, by utilising the confidence measures we can automatically detect the unreliable QT interval measurements in the signal.

Figure 4(c) shows the ECG corrupted by muscle artefact noise[2]. The signal consists of two ECG waveforms which are affected by the noise (beats 1 and 2), followed by one waveform which is relatively clean (beat 3). The QT interval for beat 3 is 360ms, whereas the QT intervals for beats 1 and 2 are 514ms and 432ms respectively. Figure 4(d) shows the confidence measures for the three waveforms, together with the 99% lower confidence bound. Both waveforms affected by the muscle artefact noise lie below the confidence threshold, thus enabling the corresponding QT interval measurements to be automatically "flagged" as unreliable. The confidence measure for beat 3 lies on the threshold border, due to the poor location of the P wave onset by the model.

5. Conclusions

The philosophy behind this work is that automated methods for QT interval analysis should not be viewed as potential replacements for human ECG analysts, but rather as *tools* to be used in conjunction with standard human analysis. Thus, by taking advantage of the confidence measure for each segmented ECG waveform, we can automatically highlight those waveforms which are least suitable to analysis by machine. In the future we anticipate

[1] We defined the waveform length as the time from the onset of the P wave to the onset of the P wave for the *following* beat

[2] This type of noise has a bandwidth similar to that of the ECG, and therefore cannot be removed by simple filtering techniques

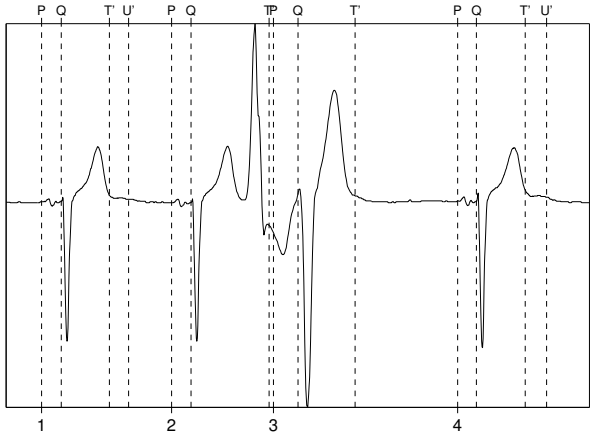

(a) ECG containing an ectopic beat

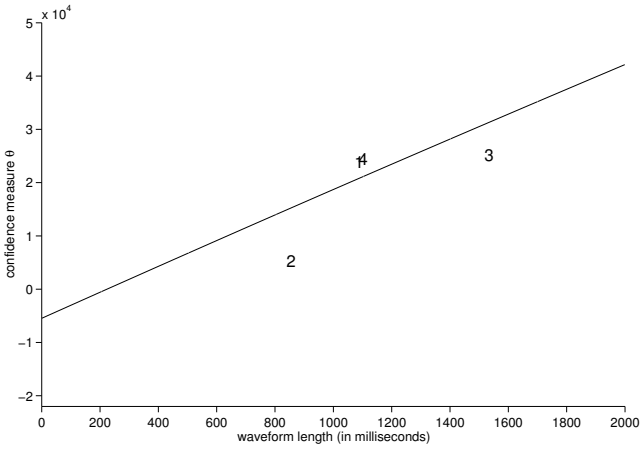

(b) Confidence measures for ectopic ECG

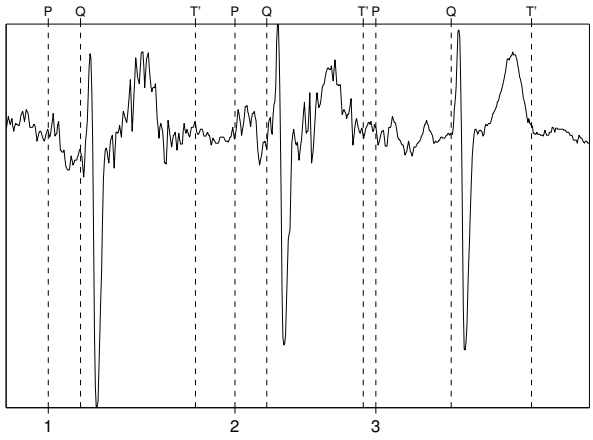

(c) ECG contaminated by muscle artefact noise

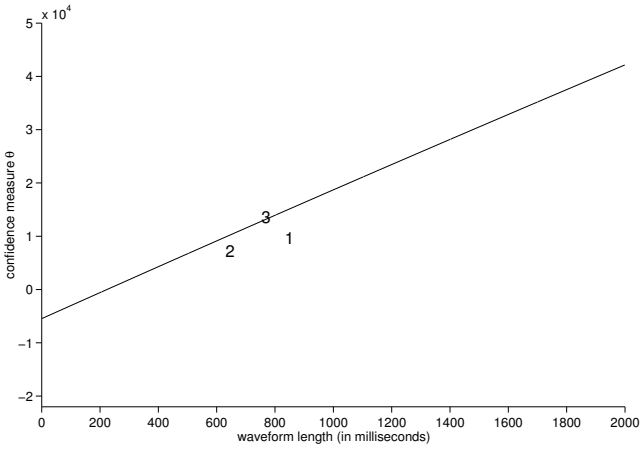

(d) Confidence measures for muscle artefact ECG

Figure 4. Segmentation results and confidence measures for ECG signals containing a.) an ectopic beat and c.) muscle artefact noise. In plots a.) and c.) the lower axis shows the numeric identifier for the start of each waveform, and the upper axis shows the segmentations. Plots b.) and d.) show the confidence measures together with the lower confidence bound.

that such methods will provide a "first-pass" of QT interval analysis by processing large quantities of ECG data and selecting only the most "difficult" waveforms for further analysis by human experts.

References

[1] Houghton AR, Gray D. Making Sense of the ECG. Arnold, 1997.

[2] McLaughlin NB, Campbell RW, Murray A. Accuracy of automatic QT measurement techniques. In Computers in Cardiology. IEEE Press, 1993; 863–866.

[3] Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. Proceedings of the IEEE 1989;77(2):257–286.

[4] Hughes NP, Tarassenko L, Roberts SJ. Markov models for automated ECG interval analysis. In Advances in Neural Information Processing Systems 16. 2004; .

[5] Percival DB, Walden AT. Wavelet Methods for Time Series Analysis. Cambridge University Press, 2000.

[6] Freund RJ, Wilson WJ. Regression Analysis: Statistical Modelling of a Response Variable. Academic Press, 1998.

Address for correspondence:

Nick Hughes
Department of Engineering Science
Parks Road, Oxford, OX1 3PJ, UK
nph@robots.ox.ac.uk