Fusing Conventional ECG QRS Detection Algorithms with an Auto-associative Neural Network for the Detection of Ectopic Beats

G. Clifford, L. Tarassenko, N. Townsend
Department of Engineering Science, University of Oxford
Parks Road, Oxford, OX1 3PJ, UK
Fax +44 (0)1865 273905, Email: gari@robots.ox.ac.uk

Abstract

The onset of a potentially fatal arrhythmia is often preceded by abnormal morphologies in the QRS complex, the main feature in the electrocardiogram. However, these ectopic beats are difficult to detect as their shape is very similar to those found in a normal sinus rhythm. We show how an auto-associative multi-layer perceptron can be trained to detect normal beats only, so that the subtle abnormalities in shape of ectopic beats become clearly identifiable. Details of how to train the network for use in a clinical environment are given utilising a new parameter, the *variance ratio*. Results for a study of a combination of algorithms to produce a robust ectopic beat detector are presented. Finally we discuss an on-line implementation for patient-specific adaptability.

Keywords

Auto-associative, Arrhythmia, ECG, Ectopic, Electrocardiogram, Neural Network, QRS Detection.

1 Introduction

The electrocardiogram (ECG) is usually recorded with surface electrodes on the chest. It is a record of the heart's electrical activity, with a number of segments corresponding to the depolarisation and subsequent repolarisation of various regions of the heart muscle. Firstly, there is a rapid positive-going transition from the Q-wave to the R-wave, immediately followed by a negative-going transition from the R-wave to the S-wave making up the high-frequency spike known as the QRS complex (as in Figure 1). The most common diagnostic use of the ECG is the identification of abnormal rhythms and therefore heart rate

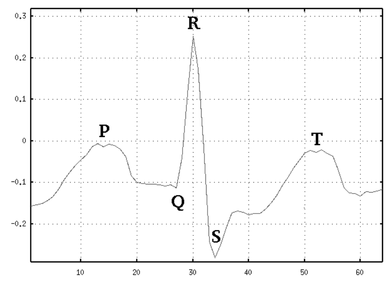

Figure 1: Typical ECG waveform sampled at 256 Hz. The points P, Q, R, S and T are known as the P-wave, Q-wave, R-wave, S-wave and T-wave respectively.

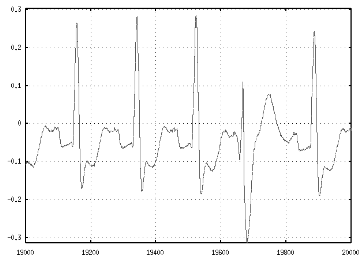

Figure 2: Ventricular Ectopic Beat ($t \approx 19700/256 =$ 77s through file 109) in a normal heart beat sequence.

computation is a fundamental requirement of ECG analysis. The easiest way to calculate the heart rate is to measure the interval between successive QRS complexes. Most QRS detection algorithms are based upon slope and energy information, consisting of a band-pass filter (usually from 10 to 40 Hz) to emphasize the high-frequency spike, followed by differentiation, squaring and time integration to highlight further the upslope and downslope of the QRS complex [4].

Abnormal rhythms (arrhythmias) are often preceded by the occurrence of ectopic beats. These are beats which do not originate from the sino-atrial node, the region of the heart from which normal beats originate. The most common type of ectopic beat is the ventricular ectopic beat (VEB), in which an ectopic beat in a ventricle gives rise to an extra R-wave, indicative of a premature ventricular contraction. It is important to try and detect these as reliably as possibly, as they are often precursors of major arrhythmias. The shape of a VEB is usually that of a distorted QRS complex (see Figure 2), but with an energy content sufficiently close to that of a normal beat that most VEBs trigger QRS detectors. They are therefore mis-identified as normal beats and their presence in an ECG record is missed. Although template matching is commonly used as a method of identifying abnormal beats, this requires the application of a large set of heuristics to the morphology of the ECG. Since inter-patient morphology variability can be quite considerable, a set of patient independent features must be identified (such as QRS broadening). It is difficult to select criteria for generating abnormal templates in an online situation and this type of method often requires several iterations through the data to allow accurate classification.

In this paper, we describe how an auto-associative neural network can quickly learn patient-specific details of an ECG and differentiate between normal and abnormal morphologies with far greater accuracy than conventional methods. We show how its output can be combined with a conventional QRS detection algorithm (mentioned above) and template matching in order to identify VEBs correctly. During learning, the auto-associative network is tuned to detect the normal beats for that subject, so that the subtle abnormalities in shape of ectopic beats become clearly identifiable using the processed output of the trained network. We show that this method is superior to a combination of template matching and conventional filtering and describe a method for real time implementation realisable on standard commercial hardware.

2 The test database

The algorithms described in this paper were developed using the MIT-BIH arrhythmia database, which is a well-known database of ECG records [3]. Each record consists of approximately 30 minutes of two channels of ECG data annotated by an expert clinician who has marked the occurrence of every beat, both normal and abnormal. For the purpose of this paper, we will discuss the performance of our algorithms on file 109 which contains 32 VEBs and 2067 normal beats but the learning procedure described here can be applied to any ECG record. File 109 contains two channels of digitised ECG data which has been down-sampled to 256Hz with an accuracy of 12 bits using the tools provided with the MIT database (http://ecg.mit.edu/).

3 Auto-associative neural network

An auto-associative network is a neural network which learns to reproduce the input data at its output. As will be shown below, a multi-layer perceptron (MLP) configured as an auto-associative network can be trained to be a highly-tuned waveform matching system. In an auto-associative MLP, the hidden layer performs a data compression function and so the number of units, M, in the hidden layer is much lower than the input dimensionality, D, of the data (see Figure 3). If the MLP is trained by minimising a standard sum-of-squares error criterion, then the auto-associative network performs a projection of the data onto the first M principal components [2], whether the hidden units have linear or non-linear (sigmoid) activation functions. It is theoretically possible to perform non-linear dimensionality reduction by adding two extra layers of non-linear units either side of the central hidden layer, but we have found that it is very difficult in practice to train such net-

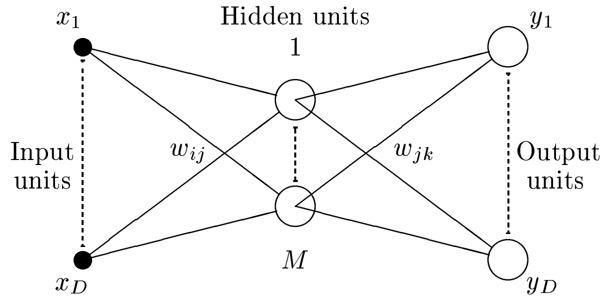

Figure 3: Architecture of D-M-D auto-associative MLP with a single layer of hidden units.

works as they are very susceptible to being trapped in local minima. We therefore use a single hidden layer of linear units in all of the work described in this paper. Of course, it would be possible to use standard techniques to perform principal component analysis (PCA) instead but the on-line learning capabilities of the auto-associative MLP is a key requirement in this application, as will be made clear at the end.

4 The training set

The training set is assembled as follows. The data from the first channel in record 109 is further downsampled from 256Hz by a factor of 4 to limit the input/output dimensionality of the auto-associative MLP, whilst preserving the key features of the ECG waveform. For each of the first N heart beats in the record, we select the one-second segment which is centered on the peak of the R-wave ("R-peak"). With an effective sampling rate of 64 Hz, this means that there are 64 samples in each one-second segment (31 or 32 samples either side of the peak).

A conventional QRS detector, adapted from the design described by Pan and Tompkins [4], is used to identify the R-peak. The aim of the learning procedure is to train the auto-associative network to reproduce as accurately as possible at its output these R-peak centered waveforms. In our initial experiments, one thousand R-peak centered waveforms were used to construct the training set, corresponding to approximately the first half of the ECG record. In the second set of experiments, only three hundred waveforms were used, corresponding (for a heart rate of 60 beats per minute) to five minutes of training data. The second set of experiments correspond to the standard conditions under which adaptive ECG algorithms have to operate within clinical practice [1].

As we want the auto-associative network to be highly tuned to normal ECG waveforms for that subject, waveforms which are corrupted even by a moderate amount of movement artefact or noise must not be allowed into the training set. A simple template matching algorithm is therefore used as a gating procedure to eliminate corrupted waveforms from the training set. Each R-peak centered waveform is compared with the average of the last ten valid waveforms; if the Euclidean distance between the waveform and the current average is below the rejection threshold, then that waveform is accepted into the training set and used to compute an updated average. Note that such a template matching procedure on its own would *not* be suitable to detect ectopic beats. With template matching, if the threshold for rejection is set too high, many normal beats, as well as some of the ectopic beats, are rejected; if it is set too low, none of the ectopic beats are rejected (even though a few of the normal beats are). For the purpose of training the MLP a high rejection threshold is set to guarantee that no abnormal beats are used at the expense of losing some normal beats from the training set.

5 Training the auto-associative MLP

The set of either N=1000 or 300 R-peak centered normal waveforms (for record 109) is used to train a 64-16-64 auto-associative network. The number of hidden units is not critical; a 4:1 compression ratio is sufficient to allow the salient features to be encoded without significant information loss. The auto-associative MLP is trained using gradient descent to minimise the mean squared error (MSE) at the output of the network. The standard error-backpropagation algorithm [2, 5] is used, with a learning rate of 0.05, to calculate the weight updates in each layer of the network. When training first starts, the MSE decreases rapidly before levelling off after several hundred iterations through the training set. The MSE, however, is not an ideal criterion for monitoring the progress of learning in an auto-associative network. A new parameter [6], the variance ratio, V_{ratio}, is used to decide when learning should be stopped.

V_{ratio} is a measure of how closely the output patterns match the input patterns; it is defined as the ratio of the variance of the patterns reconstructed at the output with respect to the variance of the input patterns (i.e. the training set).

The variance of the dth dimension in a set of N training vectors, $\mathbf{x}^n = [x_1^n, x_2^n, ..., x_D^n]^T$, where $n = 1, ..., N$, is given by:

$$v_d = \frac{1}{N-1}\sum_{n=1}^{N}(x_d^n - \mu_d)^2$$

where μ_d is the mean value of x_d over the N patterns. The overall variance V is obtained by summing over all D dimensions ($D = 64$ here):

$$V = \sum_{d=1}^{D} v_d$$

V_{ratio} is simply the ratio of the values of V for the output $\mathbf{y}$ and input $\mathbf{x}$.

Figure 4 shows the change in V_{ratio} as training proceeds with a training set of 1000 patterns from file 109. After 300 iterations (epochs), the variance ratio flattens out at a value of 0.8. This corresponds

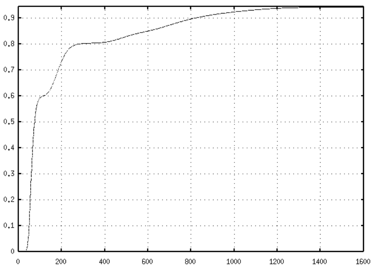

Figure 4: Change in the value of the variance ratio during the training of the auto-associative MLP. The horizontal axis shows the number of iterations through the training set (epochs).

to a local minimum for which each of the input patterns is reconstructed as the *mean of the training set*. It is therefore essential to continue training past this point, a value of V_{ratio} as close as possible to 1.0 being required to capture the *variability associated with normality*. Figure 4 shows this point being reached after 1300 epochs when V_{ratio} has increased to 0.95.

6 Detection of ectopic beats

Once the auto-associative neural network has finished training, it has learnt the details of a normal R-peak centered ECG waveform for that subject, including the allowable inter-beat variability. In test mode, the ECG waveform is slid across the 64 input nodes of the auto-associative MLP, sample by sample (at the 64 Hz rate). When the test (input) waveform is centered on its peak, as in Figure 5, the reconstructed (output) waveform is very similar. When the input waveform is *not* centered on the R-peak, even by a small amount (7 samples in Figure 6), the reconstructed waveform is very different from the input waveform. We could use the mean squared error (MSE) between the output pattern **y** and input pattern **x** to quantify the difference between the two waveforms. When a normal beat is present, centered on the R-peak in the input window, the MSE is close to zero; as soon as the input waveform has moved past this point, the MSE becomes much larger. As we are interested in *shape* rather than amplitude differences, we emphasize these by energy normalisation of both the input and reconstructed patterns, prior to computing the MSE between the two normalised waveforms. Finally, it is easier to identify the normal beats from the background by using the *reciprocal* MSE, rather than the MSE itself. As shown in Figure 7 (upper trace), each QRS complex gives a very sharp peak in the reciprocal MSE (RMSE) plot over time. Thus, the RMSE, after energy normalisation, is used as the in-

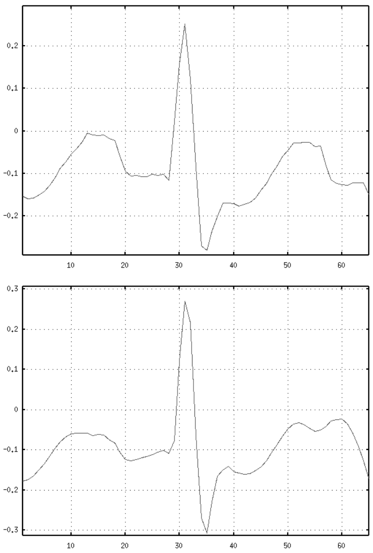

Figure 5: Reconstructed output waveform (lower trace) from the trained auto-associative MLP when the input waveform (upper trace) is centered on its R-peak.

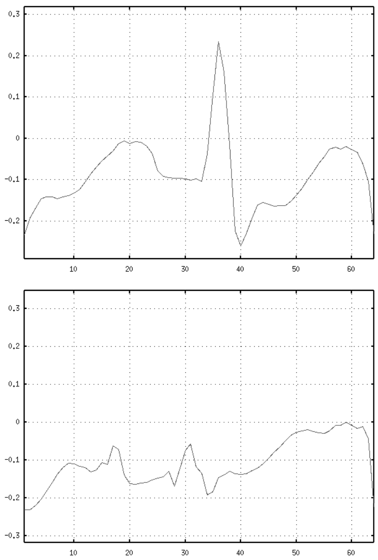

Figure 6: Reconstructed output waveform (lower trace) from the trained auto-associative MLP when the input waveform (upper trace) is off-centre.

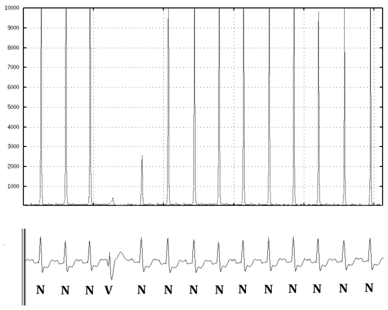

Figure 7: Variation of Reciprocal MSE (RMSE) with time as the ECG waveform shown underneath is p-resented to the trained auto-associative MLP using a one-second sliding window (64 samples).

dex for waveform matching and is defined as follows:

$$\text{RMSE} = \left\{ \sum_{d=1}^{64} (y'_d - x'_d)^2 \right\}^{-1}$$

where

$$x'_d = x_d \,/\, \sum_{d=1}^{64} x_d^2 \quad \text{and} \quad y'_d = y_d \,/\, \sum_{d=1}^{64} y_d^2$$

Figure 7 shows the RMSE plot for a 15-second section of ECG from file 109. The fourth beat is a ventricular ectopic beat (labelled as V by the expert clinician). Very high values of the RMSE (around 10,000) correspond to the R peaks of the normal beats. The ectopic beat generates only a very small rise in the value of the RMSE (well below 1000) with respect to the background. Interestingly, the normal beat which follows the ectopic beat does produce a clear peak in the RMSE but the amplitude of this peak is smaller than that associated with the other normal beats. This shows how well tuned to normality the trained auto-associative MLP is: the end of the ectopic beat runs into the next beat and affects the shape of the P wave which precedes the QRS complex for that beat. Hence it is a slightly distorted waveform (with respect to normality) which is presented to the MLP when that beat's R-peak is at the centre of the input window and so the RMSE peak is reduced to a value of around 2,500 for that beat.

In the first method which we investigated the detection of ectopic beats is achieved by running the conventional QRS detector in parallel with the trained auto-associative MLP and using the truth table shown in Table 1. For the auto-associative neural network, a 1 indicates that the waveform within the input window generates an RMSE value greater than

Neural network	QRS detector	Classification
0	0	No Beat
0	1	VEB
1	0	Normal
1	1	Normal

Table 1: Truth table for an auto-associative network/QRS detector combination algorithm.

No. training patterns	1000	300
Normal Identified as Normal	2055	2052
Normal Identified as VEB	12	15
VEB Identified as VEB	28	28
VEB Identified as Normal	0	0

Table 2: Performance of auto-associative network/QRS detector combination algorithm (VEB a Venticular Ectopic Beat).

1,000. Using the truth table of Table 1, the results given in Table 2 are obtained on file 109.

With 1,000 patterns in the training set, the auto-associative MLP is trained for 2,000 iterations. When the number of training patterns is reduced to 300, the number of iterations required for the variance ratio to reach a value of 0.95 increases to 10,000. In other words, the product of the number of training patterns and the number of training epochs remains approximately constant.

Table 2 shows that there are a few normal beats (12 or 15) which are identified as VEBs. These are normal beats which have slightly distorted shapes because of slight movement artefact or high-frequency noise. The classification could easily be restored to "Normal" by introducing a further test based on timing: these beats occur when expected (according to the heart rate at the time) whereas VEBs have a certain degree of prematurity. There are also four VEBs (out of a total of 32) which are missed because the QRS detector does not fire when those beats occur.

In the second method which we tested, the conventional QRS detector is replaced by the template matching algorithm used to select waveforms for the training set. The templates are selected by running the conventional QRS detector with the same rejection criteria used for selecting the auto-associative MLP training data on the first five minutes of the ECG only. Again, this method is run in parallel with the trained auto-associative MLP.

Two template matching criteria were considered: firstly, the MSE ($\sum_{d=1}^{64} (y'_d - x'_d)^2$) between the normalised 64-D template waveform, $x'_1, x'_2, ..., x'_{64}$ and the observed ECG waveform, $y'_1, y'_2, ..., y'_{64}$ (both cen-

No. training patterns	1000	300
Normal Identified as Normal	2055	2052
Normal Identified as VEB	12	15
VEB Identified as VEB	31	31
VEB Identified as Normal	0	0

Table 3: Performance of auto-associative network/template matching combination algorithm (VEB a Venticular Ectopic Beat).

tered on the R-peak); secondly, C, the correlation coefficient given by

$$C = \sum_{d=1}^{64} (y_d^n - \mu^n) \times (x_d^n - \nu)$$

where $\mu^n = \frac{1}{64}\sum_{d=1}^{64}(y_d^n)$, is the mean value of the n^{th} pattern and ν is the mean value of the template calculated as

$$\nu = \frac{1}{DM}\sum_{d=1}^{D}\sum_{m=1}^{M}(x_d^m),$$

with x_d^m the d^{th} element of the m^{th} valid QRS complex, D=64, the number of points in the template and M=10, the number of R-peak centered ECG waveforms.

The correlation coefficient was chosen as a criterion since it is a standard technique in image processing and it offers an intuitive measure of how two patterns differ: $C << 1$ represents a poor match whilst $C = 1$ represents a one-to-one correspondence between patterns. Note that in this case energy normalisation of the output and the input patterns is not necessary since the mean of the respective vectors is removed from each point.

The combination of the auto-associative MLP with the template matching algorithm for differentiating between normal and abnormal beats was investigated using both the MSE and C criteria. The correlation coefficient was found to give significantly better results [7]. The results of Table 3 are for the combination of the auto-associative MLP and template matching algorithm with the correlation coefficient. With $C = 0.4$, only 12 Normals and 1 VEB were missed by the template matching technique, therefore 31 out of the 32 VEBs were correctly identified for patient 109.

7 Conclusion

This paper has described results from a robust, adaptive algorithm for the detection of ventricular ectopic beats, an important requirement for the early detection of the onset of potentially fatal arrhythmias. These results indicate that template matching techniques can provide an accurate QRS detection technique when the templates are selected using a conventional QRS detector. If these algorithms are then combined with an auto-associative MLP a highly sensitive abnormal beat detector can be implemented. The auto-associative MLP can be trained on the first five minutes of ECG data, which makes it usable in clinical practice.

Often simple statistical techniques and template matching are insufficient to distinguish between normal and abnormal QRS complexes. However, we have shown that the MLP can be trained to learn normal waveform shapes, including the variability associated with normality.

With on-line learning, it is also possible for the MLP to adapt to the gradual changes in waveform morphology which can occur over several days of continuous monitoring. This simply requires the learning procedure and the template matching algorithm described above to run continuously in the background. This algorithm is simple enough to run on cost effective, commercial hardware.

References

[1] AAMI ECAR.: *Recommended Practice for Testing and Reporting Performance Results of Ventricular Arrhythmia Detection Algorithms*, Association for the Advancement of Medical Instrumentation, 1987.

[2] Bishop, C.M.: *Neural Networks for Pattern Recognition*, OUP, Oxford, 1995.

[3] Moody, G.B.: *Harvard-MIT Division of Health Science and Technology's MIT-BIH arrhythmia database*, Cambridge, MA, 1989.

[4] Pan, J. and Tompkins, W.: *A Real-time QRS Detection Algorithm*, IEEE Transactions on Biomedical Engineering, **32**(3) (1985), 230-236.

[5] Tarassenko, L.: *A Guide to Neural Computing Applications*, Arnold, London, 1998.

[6] Tarassenko, L., Clifford, G., Townsend, N.: *Detection of Ectopic Beats in the Electrocardiogram Using an Auto-associative Neural Network*, submitted to Neural Processing Letters, 2000.

[7] Clifford, G.: *Novelty Detection and Classification in Electrocardiograms*, Dept. of Engineering Science, Oxford, 1999.