Proceedings of the 29th Annual International
Conference of the IEEE EMBS
Cité Internationale, Lyon, France
August 23-26, 2007.

Shape-based Matching of ECG Recordings

Tanveer Syeda-Mahmood, David Beymer, and Fei Wang

Abstract— **An electrocardiogram (ECG) is an important and commonly used diagnostic aid in cardiovascular disease diagnosis. Physicians routinely perform diagnosis by a simple visual examination of ECG waveform shapes. In this paper, we address the problem of shape-based retrieval of ECG recordings, both digital and scanned from paper, to infer similarity in diagnosed diseases. Specifically, we use the knowledge of ECG recording structure to segment and extract curves representing various recording channels from ECG images. We then present a method of capturing the perceptual shape similarity of ECG waveforms by combining shape matching with dynamic time warping. The shape similarity of each recording channel is combined to develop an overall shape similarity measure between ECG recordings. Results are presented that demonstrate the method on shape-based matching of various cardiovascular diseases.**

I. INTRODUCTION

An electrocardiogram (ECG) is an electrical recording of the heart that depicts the cardiac cycle. It is routinely used as a first course of choice in diagnosing many cardiovascular diseases. A normal ECG waveform (in lead II) has a characteristic shape indicated in Fig. 1b. The segment labeled P represents the phase of atrial depolarization/contraction when deoxygenated blood enters the heart from the right atrium and oxygenated blood enters the heart from the lungs into the left atrium (Fig. 1a). The QRS segment represents the phase of ventricular depolarization/contraction when blood enters the right and left ventricles for ejecting into the pulmonary artery and Aorta respectively. Finally, the T segment represents ventricular repolarization where the ventricles relax to allow the cycle to begin again. Many disturbances in the heart function show as characteristic variations in the sinus rhythm waveform of Fig. 1b, and can be used as cues to diagnose the disease. Fig. 1c shows such a modification in the ECG due to premature ventricular contraction, where the heart skips a beat only to beat very strongly in the next, causing a missed R segment. Physicians routinely make diagnosis by a simple visual examination of these ECG waveform. It is common knowledge to physicians that patients with the same disease have similar-looking ECG shape in the relevant channels. Examples of such similarity can be seen from Fig. 2 which shows ECG recording of several patients all diagnosed with bundle branch block.

The goal of this paper is to capture this fundamental intuition used by physicians in a computational algorithm for finding shape similarity of ECG recordings. We will handle both digital ECG recordings as well as scanned paper ECGs using novel image processing techniques. Specifically, we use the knowledge of ECG recording structure to segment and extract curves representing various recording channels from ECG images. We then present a method for capturing the shape similarity of ECG waveforms by modeling the morphological variations in ECG shapes for the same disease as a constrained non-rigid translation transform. This non-rigid alignment transform is recovered using a variant of dynamic time warping that explicitly accounts for missing and spurious fiducial features in ECG images. Results are presented that demonstrate the effectiveness of the method in retrieving matching ECGs depicting the same disease from a large database of ECG recordings.

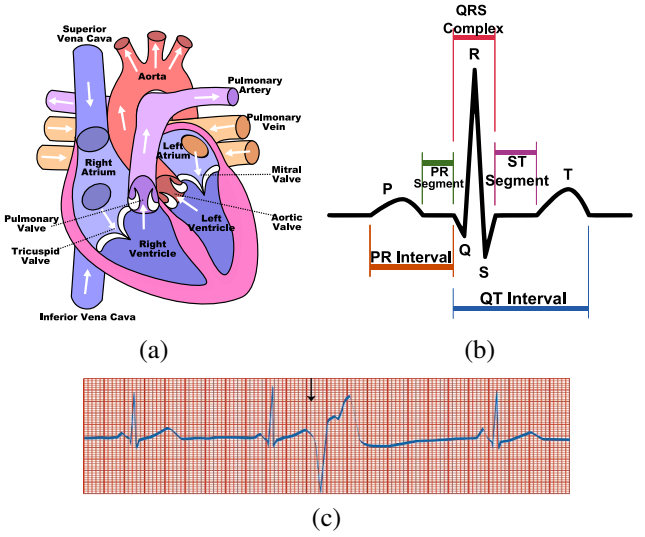

Fig. 1. Illustration of the Electrocardiogram and its variations under a disease. (a) The heart cycle. (b) A normal ECG. (c) An abnormal ECG showing premature ventricular contraction. *Sources: (a) Wikipedia Commons, under GNU Free Documentation License, (b) Wikipedia, in public domain.*

The paper makes several novel contributions. First, while most existing work has addressed the problem of analyzing the ECG of a single patient, we focus on the problem of searching for matching ECGs from a database. As such, it is the first practical application of content-based retrieval techniques for diagnosis validation during decision support. Next, our system caters to both scanned and digital ECGs. Finally, our ECG waveform digitization algorithm is fully automatic and does not need operator intervention to set system parameters.

The rest of the paper describes our approach in detail. In Section II, we discuss related work in ECG analysis. In Section III, we describe extraction of ECG waveforms from

The authors are with the Healthcare Informatics Group, IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120 USA.

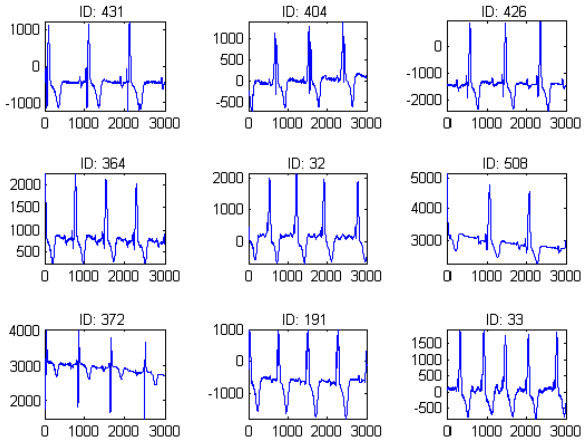

Fig. 2. Illustration of similarity of ECG shapes for similar diseases.

ECG images. Sections IV and V present our shape matching algorithm. In Section VI, we present results demonstrating its use for shape-based retrieval of similar ECG images to infer the similarity of diseases.

II. RELATED WORK

A number of research and commercial systems exist for digitizing ECG waveforms on hardcopy printouts [18], [5], [13], [12], [14], [15], [16], [4]. They typically start by detecting the ECG waveform and separating it from the background grid, usually via color segmentation or a morphological operator [12]. Once detected, the waveform pixels are thinned [5], if necessary, and traced using line tracking [14] or more advanced model fitting such as snakes [4]. The greatest limitation of existing approaches is the need for manual intervention, especially to locate bounding boxes around the lead waveforms. Lobodzinski, et al.'s ECG waveform recognition technique [14] includes templates for parsing the ECG page, although it is unclear from their description if this includes automatic lead extraction.

There are a number of algorithms available for single ECG analysis [3], and for ECG classification based on neural network [11], expert and fuzzy expert systems [9], machine learning methods, wavelet transforms [8] and genetic algorithms. However, very little work has been done to date in automatic search of similar ECGs and using it to infer disease similarity. The rule-based methods rely on the accuracy of the P-Q-R-S-T segment detection [3]. Errors in estimation of these feature values can cause major errors in disease-specific interpretation. Further, in order to distinguish combinations of diseases, a finer shape analysis of the ECG waveform may be required. The parametric modeling methods, on the other hand, are good at spotting major disease differences but can't take into account fine morphological variability due to heart rate (e.g. ventricular vs. supra-ventricular tachycardia) and physiological differences. Related work in the time alignment of ECGs also exists. Dynamic time warping has been a popular technique in ECG frame classification [10], and more recently, in the recognition of heart beat patterns for synthetically generated signals [17]. In all such alignments, however, the amplitude of the signal was used rather than a detailed modeling of the shape. Moreover, the DTW algorithm used did not explicitly model the morphological changes in the signal across patients with similar diseases as it does not take into account missing and spurious fiducial features during alignment.

III. EXTRACTION OF ECG SHAPES FROM IMAGES

Although most hospitals now have digital ECG recorders, much of the legacy ECG data is still in paper form, inaccessible to analysis by computer. Unlocking these ECG records printed on paper and exposing them to digital analysis would be useful. It provides additional data for current ECG analysis techniques, as well as historical data to be used for comparative studies. For instance, the Framingham Heart Study [2] has been collecting ECG data since 1948, with the data collected before 1990 in hard-copy form. Statistical or search-based ECG analysis techniques that require large ECG databases can be populated with such samples of legacy data. Therefore the analysis of paper ECG may still be relevant for disease-specific retrieval.

We now describe the extraction of ECG waveforms from the scanned ECG images. Processing consists of three steps (a) background removal, (b) lead segmentation, and (c) ECG waveform extraction.

A. Preprocessing

ECG images, scanned at a resolution of 600 dpi, capture the waveform trace at a good thickness, usually about 4 pixels. Also, due to the limited number of colors in the image, the foreground can be easily segmented from the background grid using the intensity histogram. Fig. 3b shows the result of this pre-processing of the ECG image shown in Fig. 3a (only a portion of the scanned image is shown for brevity). Notice that this operation retains both the ECG waveforms as well as the lead annotation which will be used for lead segmentation described next.

B. Lead segmentation

The ECG image records all 12 channels or lead recordings by interlacing 3 second intervals from combinations of leads per row as seen from Fig. 4. They always occur in the same order (first row: I, AVR, V1,V4, second row: II,AVL,V2,V5 and third row: III,AVF,V3,V6) all occurring aligned in columns. Since different diseases manifest differently in each of the leads (for example, lead V1 shows hypertrophy more clearly than in other leads), it is important to isolate the different leads from the ECG image. Because the lead labels in ECG recordings are placed in fixed positions relative to the channel recordings, we can locate them using normalized correlation with local templates. As the scanned images are large (5000 x 4000 pixels), and since grey-level correlation is more robust than binary correlation, we smooth the background subtracted image using a Gaussian pyramid [6]. If the scanner dpi is known, the scale can be set to directly

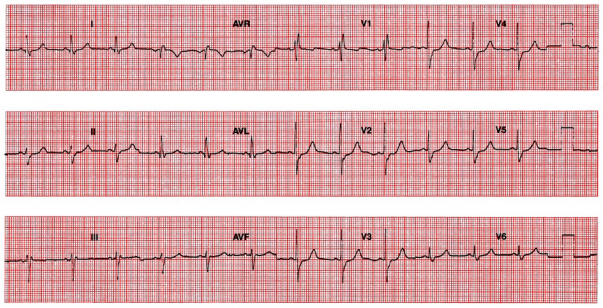

(a)

(b)

Fig. 3. Illustration of the pre-processing for ECG waveform extraction. (a) ECG image showing a 12 channel recording. (b) processed image after background subtraction.

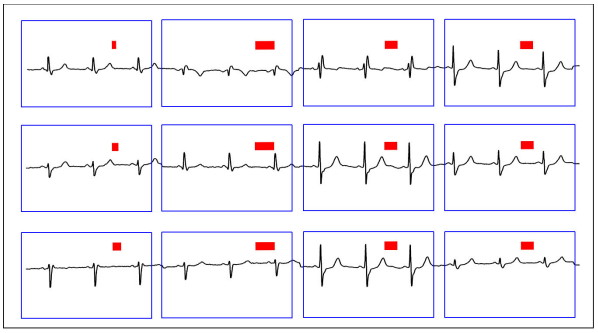

Fig. 4. ECG waveforms extracted from the image of Fig. 3a for all 12 recording leads.

compute the proper pyramid level for correlation. Otherwise, a coarse-to-fine search is performed over the pyramid levels with the lead templates.

To increase the accuracy of template matching, we match the templates in order of their character size starting from the distinctive 3 character labels (AVR, AVL, AVF) to the featureless lead label I. Since the labels appear in a grid pattern, finding one label constrains the location of others in the same row or column. Fig. 4 shows some detected lead labels using red bounding boxes.

Since the labels are fixed relative to the recording channels, the ECG recordings per channel can now be segmented relative to the location of the labels. The method is made invariant to image rotation by measuring skew in the pattern of detected lead labels.

C. ECG waveform extraction

After each lead position in the image is segmented, we extract the ECG waveforms as curves in the respective image segments. Due to noise and dropouts in the recordings, there are often gaps which cause problems in curve extraction. Note that these recordings are actually time series or functions of *lead-voltage vs time*. Thus, a general purpose algorithm such as curve following or skeletonization, may not enforce the constraint that a single $y-$value occurs for each x position in image coordinates. In places where the axis bifurcates or turns back on itself, a post-processing step would be required for choosing between multiple values of y for a particular value of x.

Since the ECG curve may be multiple pixels thick at the scanned resolution, our curve tracing algorithm follows the upper and lower edges of the curve. Fig. 5 depicts these curves, $y_u(x)$ and $y_l(x)$. Our tracing algorithm uses a first order linear filter, modeling both the position $y(x)$ and velocity $y'(x)$ for each curve. To extend the curve trace from $x-1$, we first use the velocity estimate to predict the current location at x

$$y_u^-(x) = y_u(x-1) + y'_u(x-1), \tag{1}$$

where $y_u^-(x)$ is the prediction (y_l is similar). Using this prediction to extend $y_u(x-1)$ and $y_l(x-1)$, we define two search ranges centered around these values:

$$\mathbf{R}_u = \left[y_u^-(x) - W, y_u^-(x) + W\right] \tag{2}$$

$$\mathbf{R}_l = \left[y_l^-(x) - W, y_l^-(x) + W\right] \tag{3}$$

where W is a search window. In practice, W needs to be set large enough to handle the voltage spike at an R wave. Next, we define

$$y_u(x) = \min_{y \in \mathbf{R}_u} y \quad \text{such that} \quad T(x,y) = 0 \tag{4}$$

$$y_l(x) = \max_{y \in \mathbf{R}_l} y \quad \text{such that} \quad T(x,y) = 0 \tag{5}$$

where $T(x,y)$ is the segmented foreground image containing the curves and lead labels. After tracing the upper and lower envelopes of the ECG curve, y_u and y_l, the average $(y_u + y_l)/2$ is used as the traced value $y(x)$.

To provide robustness to noise and missing curve fragments, we use morphological operators and a grouping algorithm across the gaps. A morphological open (erode + dilate) helps fill in holes inside the curve, and a morphological close (dilate + erode) eliminates noise pixels near the curve. To close small gaps, curve tracing is started using a number of seed points along the expected ECG curve location. Gaps are closed between pairs of consecutive fragments if the gap is small enough (4 pixels = 10 msec). For gaps that are larger than the gap threshold, we have found that most larger breaks in the curve occur at the R wave, where the signal spikes up and down. At the R wave, however, the curve is nearly vertical, so the sampling in Equations (2)-(3) is nearly tangent to the curve. Thus, we can handle a large fraction of vertical dropout, as shown in Fig. 5.

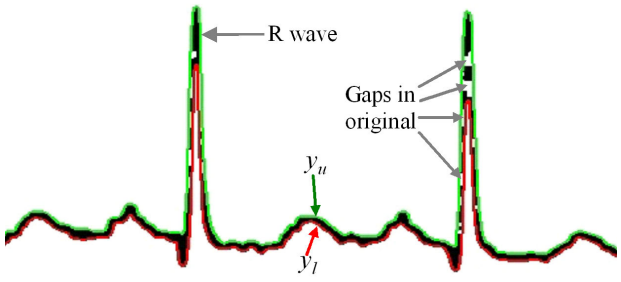

Fig. 5. Illustration of curve tracing across gaps in ECG images. Upper and lower bounding curves of the ECG, y_u and y_l, are shown in green and red. The R wave spikes show a number of gaps in the original curve that are correctly bridged by our tracing technique.

Given the 3x4 table layout located using the lead labels, we trace curves across all columns in each of the 3 rows. The rectangle enclosing each table cell is used to assign curve fragments to the corresponding voltage lead. In addition, the bounding boxes around each lead label are excluded from curve tracing to prevent the alteration in shape due to merger with the lead label pixels. Fig. 4 shows the ECG curves extracted from the scanned image of Fig. 3a corresponding to the various leads. The blue bounding boxes show the tabular structure used to parse the curves by lead. Once the lead ECG waveforms are extracted, the pixels coordinates can now be expressed as a one-dimensional time series. The next section describes the matching of these time series.

IV. MODELING SHAPE VARIATIONS OF ECG

We now turn to the problem of matching the ECG waveform shapes. As described in Section I and as seen from Fig. 2, patients with the same disease label have similarities in their ECG. This observation of similarity, however, is after factoring out a number of morphological variations that can be attributed to heart rate variability, and the measurement variability in ECG recordings that affect the amplitude levels. Further, there seems to be a built-in tolerance to the small relative translation of characteristic segments of the ECG such as the P,Q,R,S,T as long as the shape of these segments is preserved. The first row of Fig. 2 illustrates this perceptual shape similarity in ECG from channel V1 for male and female patients all diagnosed with right Bundle Branch Block.

A. Modeling the ECG shape variation

We begin by modeling the shape variations in ECG taken from patients diagnosed with the same disease. Consider an ECG $g(t)$ corresponding to disease X. For ease of discussion, we assume that the relevant channel for the diagnosis of disease X has already been pre-selected. Further, due to the periodic nature of ECG, we assume that the duration of the signal considered for shape matching can be restricted to a single heart beat. Consider another ECG $f(t)$ that is a potential match to $g(t)$ corresponding to the same channel. The signal $f(t)$ is considered perceptually similar to $g(t)$ if a non-rigid transform characterized by $[a,b,\Gamma]$ can be found such that

$$\left|f'(t)-g(t)\right| \leq \delta \tag{6}$$

where $|.|$ represents the distance metric that measures the difference between $f'(t)$ and $g(t)$, the simplest being the Euclidean norm and

$$f'(t) = af(\Phi(t)) \quad \text{with} \quad \Phi(t) = bt+\Gamma(t) \tag{7}$$

where the $[a,b]$ is the linear component of the transform and Γ is the non-linear translation component. The parameters a and b can be solved by normalizing in amplitude and time. That is, if we transform $f(t)$ and $g(t)$ such that

$$\hat{f}(t) = \frac{f(t)-f_{min}(t)}{f_{max}(t)-f_{min}(t)} \quad \text{and} \quad \hat{g}(t) = \frac{g(t)-g_{min}(t)}{g_{max}(t)-g_{min}(t)}, \tag{8}$$

then $a=1$.

To eliminate solving for b, we can normalize the time axis, so that all time instants lie in the range [0,1]. Since the duration we are considering is a single heart beat long, the time normalization can be easily achieved as:

$$\vec{f}(t) = \hat{f}(t/T_1) \quad \text{and} \quad \vec{g}(t) = \hat{g}(t/T_2) \tag{9}$$

where T_1 and T_2 are the heart beat durations of $f(t)$ and $g(t)$ respectively. With this time normalization, $b=1$. Such amplitude and time normalization automatically makes the shape modeling invariant to voltage variations in ECG recordings, as well as variations in heart rate.

Since the non-uniform translation Γ is a function of t, we can recover it at important fiducial points (features) in the normalized signals, and recover the overall shape approximation by time interpolation. Let there be K features extracted from $\vec{f}(t)$ as $F_K = \{(t_1,\vec{f}_1(t_1)),(t_2,\vec{f}_2(t_2)),\ldots(t_K,\vec{f}_K(t_K))\}$ at time $\{t_1,t_2,\ldots t_K\}$ respectively. Let there be M fiducial points extracted from $\vec{g}(t)$ as $G_M = \{(t'_1,\vec{g}_1(t'_1)),(t'_2,\vec{g}_2(t'_2)),\ldots(t'_M,\vec{g}_M(t'_M))\}$ at time $\{t'_1,t'_2,\ldots t'_M\}$ respectively. If we can find a set of N matching fiducial points $C_\Gamma = \{(t_i,t'_j)\}$, then the non-uniform translation transform Γ can be defined as:

$$\Gamma(t) = \begin{cases} t_i & \text{if } t=t'_j \text{ and } (t_i,t'_j)\in C_\Gamma \\ t_r + \left(\frac{t_s-t_r}{t'_l-t'_k}\right)(t-t'_k) & \text{where } (t_r,t'_k),(t_s,t'_l)\in C_\Gamma \end{cases} \tag{10}$$

and t'_k is the highest of $t'_j \leq t$ and t'_l is the lowest of $t'_j \geq t$ that have a valid mapping in C_Γ. Other interpolation methods besides linear (eg. spline) are also possible.

Using Equations 6 and 10, the shape approximation error between the two time series is then given by:

$$\left|f'(t)-g(t)\right| = \left|\vec{f}(\Gamma(t))-\vec{g}(t)\right| \tag{11}$$

For each $g(t)$, we would like to select Γ such that it minimizes the approximation error in (11) while maximizing the size of match C_Γ.

Finding the best matching ECG based on shape can then be formulated as finding the $g(t)$ such that

$$g_{best} = \arg\min_g \left|\vec{f}(\Gamma(t))-\vec{g}(t)\right| \tag{12}$$

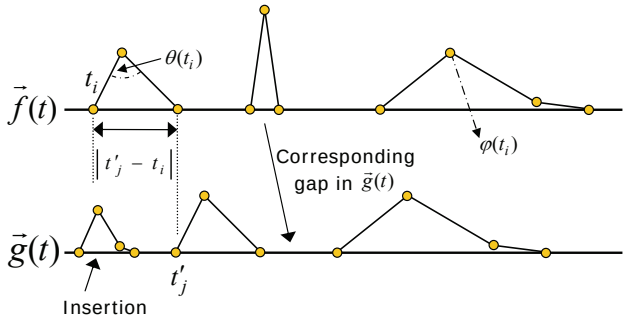

Fig. 6. Illustration of insertion and gaps in ECG alignment.

while choosing the best Γ for each respective candidate match $g(t)$.

B. Solving for Γ

If we consider the feature set F_K, G_M extracted from the respective time series as sequences, the problem of computing the best Γ reduces to finding the best global subsequence alignment using the dynamic programming principle. The best global alignment maximizes the match of the time series fragments while allowing for possible gaps and insertions. Gaps and insertions correspond to signal fragments from feature set F_K that don't find a match in set G_M and vice versa. In fact, the alignment can be computed using a dynamic programming matrix H where the element $H_{i,j}$ is the cost of matching up to the ith and jth element in the respective sequences. As more features find a match, we want the cost to increase as little as possible. The dynamic programming step in our case becomes:

$$H_{i,j} = \min \begin{cases} H_{i-1,j-1} + d(\vec{f}(t_i), \vec{g}(t'_j)) \\ H_{i-1,j} + d(\vec{f}(t_i), 0) \\ H_{i,j-1} + d(0, \vec{g}(t'_j)) \end{cases} \quad (13)$$

with initialization as $H_{0,0} = 0$ and $H_{0,j} = \infty$ and $H_{i,0} = \infty$ for all $0 < i \leq K$, and $0 < j \leq M$. Here $d(.)$ is the cost of matching the individual features described in detail next. Also, the first term represents the cost of matching the feature point $\vec{f}(t_i)$ to feature point $\vec{g}(t'_j)$ which is low if the features are similar. The second term represents the choice where no match is assigned to feature $\vec{f}(t_i)$. Fig. 6 illustrates the accounting of insertions and gaps in the computation of the match between the two time series.

C. Characterizing fiducial point similarity

If we regard the time series as curves, a natural choice of fiducial points are the corners as shown by the circles in Fig. 6. The corners can be easily obtained as the end points of a line segment approximation to curve. The shape information at each corner is modeled using the following parameters

$$S(\vec{f}(t_i)) = \left\langle t_i, \vec{f}(t_i), \theta(t_i), \phi(t_i) \right\rangle \quad (14)$$

where $\theta(t_i)$ is the included angle in the corner at t_i, and $\phi(t_i)$ is the orientation of the bisector at corner t_i. Using the angle of the corner ensures that wider QRS complexes are not matched to narrow QRS complex as these can change the disease interpretation. The angular bisector, on the other hand, ensures that polarity reversals such as inverted T waves or change in ST elevation can be captured. We assume $(\theta(t_i), \phi(t_i))$ are both normalized to lie in the range [0,1] as are t_i and $\vec{f}(t_i)$. The fiducial points in $\vec{g}(t)$ can be defined similarly. The cost function $d(\vec{f}(t_i), \vec{g}(t'_j))$ is then given as the Euclidean distance between the two fiducial points using the 4 parameters as

$$d(\vec{f}(t_i), \vec{g}(t'_j)) = \begin{cases} \sqrt{\begin{array}{l}(t_i - t'_j)^2 + (\vec{f}(t_i) - \vec{g}(t'_j))^2 + \\ (\theta(t_i) - \theta(t'_j))^2 + \\ (\varphi(t_i) - \varphi(t'_j))^2\end{array}} & \begin{array}{l}\text{if } \left|t_i - t'_j\right| \leq \lambda_1 \text{ and} \\ (\vec{f}(t_i) - \vec{g}(t'_j))^2 \leq \lambda_2 \\ \left|\theta(t_i) - \theta(t'_j)\right| \leq \lambda_3 \\ \left|\varphi(t_i) - \varphi(t'_j)\right| \leq \lambda_4\end{array} \\ \infty & \text{otherwise} \end{cases} \quad (15)$$

The thresholds $(\lambda_1, \lambda_2, \lambda_3, \lambda_4)$ are determined through a prior learning phase in which the expected variations per disease class is noted. The cost function $d(\vec{f}(t_i), 0)$ can be computed by substituting $t'_j = 0, \vec{g}(t'_j) = 0$ and $\theta(t'_j) = 0, \phi(t'_j) = 0$ in Equation 15. The cost function $d(0, \vec{g}(t'_j))$ can be similarly computed.

V. SHAPE MATCHING ALGORITHM

We now describe the overall shape matching algorithm to align a pair of ECG signals extracted from the ECG images from the same observation lead. The algorithm consists of several steps as described below.

A. Extraction of single periods for matching

To extract the single heart beat duration, we normalize a signal $f(t)$ in amplitude as given in Equation 8, and compute the autocorrelation function. As shown in Fig. 7b the peaks in the autocorrelation function correspond to the various periodicity patterns found in the signal. We note the most common inter-peak duration as representative of a heart beat duration and extract a segment of recovered duration from the ECG signal. This segment becomes the basis of our shape-based alignment scheme. The single heart-beat containing region extracted using the inter-peak distance in the autocorrelation function of Fig. 7b is shown in Fig. 7c.

The normalization of the time axis for a single heart beat duration is performed as given by Equation 9. This ensures that all signals being compared are one heart beat long and have their time values range from 0 to 1.0.

The fiducial points extracted from time series are corners. A simple line segment approximation that does a recursive partitioning of the time series curve is used. A threshold on minimum length = 5, and amplitude deviation of 0.01 was found sufficient to remove much of the noise while still keeping the main P,Q,R,S,T features.

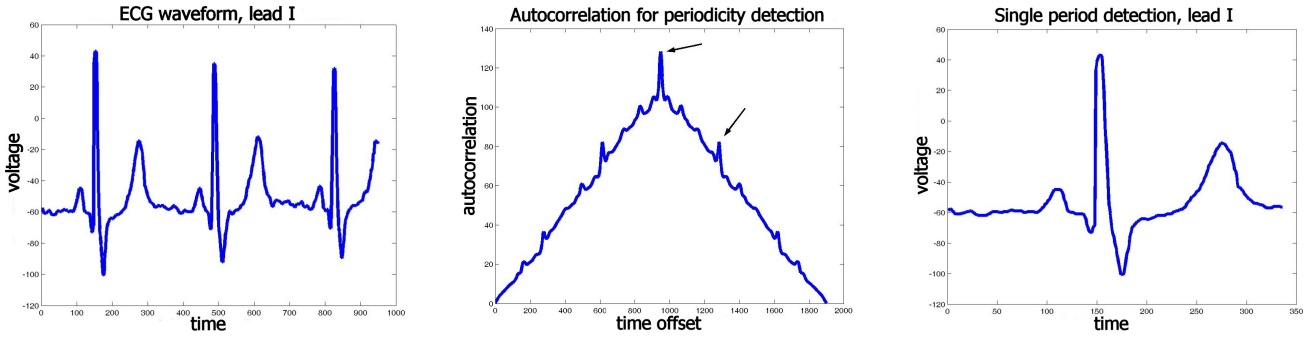

(a) (b) (c)

Fig. 7. Illustration of pre-processing steps in shape matching algorithm. (a)ECG waveform extracted from the ECG image of Fig. 3a for lead I. (b) Autocorrelation-based period detection (the distance between the two arrows. (c) Single period extraction for matching from (b).

B. Shape matching algorithm

By selecting one heart beat interval from the original signal, there is an initial translation bias depending on the starting point for such interval choice. The signals as a result may need to be circularly shifted to perform an initial registration. As the translation required is usually much larger than that is allowed during DTW alignment, it is extracted separately. For this, we cross-correlate the two signals to be matched and note the translation corresponding to the peak as the initial translation. Since the signals are periodic, a circular shift is performed.

Once the pair of signals are initially registered, the DTW alignment is performed as described in Section IV-B. The alignment transform is then used to project one signal onto the other as given in Equation 10, and the residual error is evaluated using Equation 11.

This pair-wise matching of single heart beat intervals is repeated over multiple such heart beat segments over the available data and the average residual error is used to rank the matches. Finally, channel information from each of the leads is combined through a winner take-all mechanism to retain the best matching ECGs as those that match in the highest number of channels with a score exceeding a chosen threshold. Future work will explore other channel fusion methods.

VI. RESULTS

We now present results of finding disease similarity by shape matching of ECGs. We collected scanned images of ECG from several collections including the Harvard data set of 301 ECGs [7], and 168 ECG images scanned from an ECG practice book for physicians [1]. In addition, we combined the digital ECG recordings from PTB benchmark database available from Physionet (http://www.physionet.org) which contains 12 channel 547 ECG recordings sampled at 1000 Hz of 294 patients with diseases including those listed in Table I, giving rise to a total dataset of size 1016 x 12 = 12,192 channels.

A. Examples

We illustrate shape-based disease similarity detection using an example in Fig. 8. Fig. 8a shows the single periods of ECG waveforms extracted as described in Section V-A. As can be seen, the waveforms are similar except for a non-rigid translation transform. This is computed using the DTW algorithm described in Section IV-B to give an alignment as shown in Fig. 8c. The resulting alignment of the waveforms is shown in Fig. 8b. The improvement in shape matching due to non-rigid DTW alignment can be clearly seen by comparing the simple overlaid shapes in Fig. 8b. As can be seen from Fig. 8c, the alignment is close to the diagonal illustrating a good match.

Fig. 2 shows the results of shape-based retrieval of matching ECGs from the dataset described above. The query ECG is shown in Fig. 2, left-top and corresponds to a patient diagnosed with Bundle Branch Block (BBB). The top 8 matches using the DTW-based shape matching algorithm are shown in Fig. 2 ordered left-to-right, top to bottom. Upon disease label verification, 6 of these were found to have right BBB.

Disease	Number of queries	Average # matches returned	Average recall	Average precision
Healthy Control	76	100.7	94.5%	71.3%
Hypertrophy	7	8.7	79.8%	67.3%
Left BBB	16	21.4	87.3%	61.6%
Right BBB	6	8.2	83.3%	62.2%
MI	362	584.2	96.2%	59.6%

TABLE I
ILLUSTRATION OF PRECISION AND RECALL FOR VARIOUS DISEASES BASED ON ECG SHAPE MATCHING.

B. Precision and recall

We evaluated the precision and recall values by using all available samples per disease as queries and retrieving matches above a threshold of 0.4. The number of matches retained averaged across the queries for various diseases is shown in Column 3 of Table I. The precision and recall values given below were averaged over the queries tested

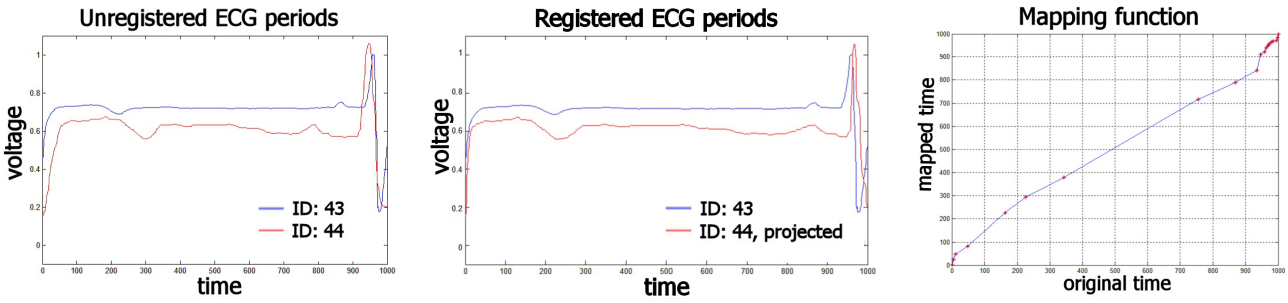

(a) (b) (c)

Fig. 8. Illustration of shape matching by non-rigid alignment. (a) Two ECG waveforms that are shape-wise similar except for a non-rigid translation transform. (b) Result of alignment. (c) DTW-alignment function.

for the respective classes.

$$\text{Recall} = \frac{\text{Number of correct matches selected}}{\text{total number of correct matches present}} \quad (16)$$

$$\text{Precision} = 1 - \frac{\text{Number of incorrect matches selected}}{\text{total number of matches returned}}$$

As can be seen from Table I, the shape matching has high recall even when the precision is not as high. On analyzing the results, we found the precision could be improved as the disease labels in the PTB database were sometimes incomplete, particularly in cases where a patient has multiple diseases and only a few of them are annotated for the corresponding ECG.

VII. CONCLUSIONS

In this paper we have presented an algorithm for non-rigid alignment of ECG shapes extracted from ECG recordings, both digital and scanned paper recordings. Starting from the key idea that patients with similar disease labels have similarities in their ECGs, we demonstrated an algorithm for disease similarity detection, where the precision-recall values for different disease groups indicate the effectiveness of the proposed algorithm.

REFERENCES

[1] Ecg wave-maven: Harvard electrocardiogram datasets. http://ecg.bidmc.harvard.edu/maven/mavenmain.asp.

[2] The Framingham Heart Study, Online Bibliography. http://www.framinghamheartstudy.org/Publications.asp.

[3] C. Alexakis, H. Nyongesa, R. Saatchi, N. Harris, C. Davies, C. Emery, R. Ireland, and S. Heller. Feature extraction and classification of electrocardiogram (ecg) signals related to hypoglycaemia. *Computers in Cardiology*, pages 537–540, 2003.

[4] F. Badilini, T. Erdem, W. Zareba, and A. J. Moss. ECGScan: a method for conversion of paper electrocardiographic printouts to digital electrocardiographic files. *Journal of Electrocardiology*, 38:310–318, 2005.

[5] H. Bhullar, D. deBono, J. Forthergill, and N. Jones. A computer based system for the study of QT intervals. *Computers in Cardiology*, pages 533–536, 1992.

[6] P. J. Burt and E. H. Adelson. The laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, COM-31,4:532–540, 1983.

[7] D. Davis. *Quick and Accurate 12-Lead ECG Interpretation*. Lippincott, Williams, and Wilkins, 2005.

[8] P. de Chazal, B. G. Celler, and R. B. Reilly. Using wavelet coefficients for the classification of the electrocardiogram. *Proc IEEE. Conf. on Engineering in Medicine and Biology Society*, 1:64–67, 2000.

[9] L. Gargasas, R. Ruseckas, and R. Jurkoniene. An expert system for diagnosis of coronary heart disease (chd) with analysis of multicardiosignals. *Medical & Biological Engineering Computing*, pages 734–735, 1999.

[10] B. Huang and W. Kinsner. Ecg frame classification using dynamic time warping. *Proc. IEEE Canadian Conf. on Elec. and Comp. Engineering*, 2:1105 – 1110, 2002.

[11] N. Izeboudjen and A. Farah. A new neural network system for arrhythmia's classification. In *Neural Computation*, pages 208–212, 1998.

[12] T. Kao, L.-J. Hwang, Y.-H. Lin, T.-H. Lin, and C.-H. Hsiao. Computer analysis of the electrocardiograms for ECG paper recordings. *Proc IEEE. Conf. on Engineering in Medicine and Biology Society*, pages 3232–3234, 2001.

[13] W. Lawson, G. Wagner, R. Startt-Selvester, and G. Ybarra. New method for digitization and computerized analysis of paper recordings of standard 12-lead electrocardiograms. *Computers in Cardiology*, pages 41–44, 1995.

[14] S. M. Lobodzinski, U. Teppner, and M. Laks. State of the art techniques for the preservation and reuse of hard copy electrocardiograms. *Journal of Electrocardiology*, 36:151–155, 2003.

[15] S. Mitra, M. Mitra, and B. Chaudhuri. Generation of digital time database from paper ECG records and fourier transform-based analysis for disease identification. *Computers in Biology and Medicine*, 34:551–560, 2004.

[16] E. Morales, D. Sevilla, J. Pierluissi, and H. Nazeran. Digitization and synchronization method for electrocardiogram printouts. *Proc IEEE. Conf. on Engineering in Medicine and Biology Society*, pages 1588–1591, 2005.

[17] V. Tuzcu and S. Nas. Dynamic time warping as a novel tool in pattern recognition of ecg changes in heart rhythm disturbances. *IEEE Intl. Conference on Systems, Man and Cybernetics*, 1:182–186, 2005.

[18] L. E. Widman and L. S. Hines. Digitization of electrocardiograms by desktop optical scanner. *Journal of Electrocardiology*, 24(4):325–338, 1991.