Analysis of Cardiac Motion with Recursive Comb Filtering

John C. McEachen II[†], Arye Nehorai[†], and James S. Duncan[†,‡]

Departments of [†]Electrical Engineering, and [‡]Diagnostic Radiology
P.O. Box 208042
Yale University
New Haven, CT 06520-8042

ABSTRACT

A framework for temporal analysis of left ventricular (LV) endocardial wall motion is presented. This approach uses the technique of two-dimensional comb filtering to model the periodic nature of cardiac motion. A method for flow vector computation is presented which defines a relationship between image-derived, shape-based correspondences and a more desirable, smoothly varying, set of correspondences. A recursive filter is then constructed which takes into consideration this relationship as well as knowledge of temporal trends. Experimental results for contours derived from cycles of actual cardiac magnetic resonance (MR) images are presented. Applications to the analysis of regional LV wall function are discussed.

Keywords: recursive filtering, non-rigid motion, shape-based tracking, harmonic analysis

1 INTRODUCTION

Knowledge of the correspondence between given frames from an image sequence of the cardiac cycle allows precise description of the associated motion of the heart and measurement of physical parameters that aid in analysis of the heart's dynamics. In this paper we develop a novel technique for tracking motion of the Left Ventricular (LV) wall of the heart using two-dimensional harmonic analysis of endocardial deformation. The novelty of this method arises in the ability to incorporate knowledge of temporal trends into the tracking problem. The efforts in this paper build upon previous work presented [1], [2], and [3]. Specifically, a temporal model is defined and a statistical framework structured on previously defined data-to-solution relationships. This provides substantial improvements over *ad hoc* techniques currently in use. The approach emerged from research into the more general problem of understanding the non-rigid motion that might be associated with deformable objects' bounding contours and surfaces in images.

Within computer vision, considerable attention has been given in recent years to tracking and understanding object motion [4]. However, this work is dominated by the study of the motion of rigid (often people-made) objects. Because of its obvious relationship to the study of natural (e.g. anatomical and/or biological objects), we have been interested in developing approaches to study non-rigid object motion. We point out that this appears to be an area ripe for more attention. There are various examples of types of non-rigid motion that researchers find interesting to study, for example cloud movement from visible light images [5] and cell motility from microscopic images [6]. The primary problem that we deal with in this paper is the study of the movement of the LV boundary of the heart obtained from any of several two-dimensional diagnostic imaging modalities in fixed views. Later in this section we refer to some of the previous work in the computer vision literature that has influenced the development of the approach proposed in this paper.

1.1 Background

The problem of quantifying regional LV wall function is not a new one in the field of cardiology, and in fact a number of previous approaches have been attempted. The most common of these attempt to relate regional function to

endocardial motion. These fall into three rough categories: 1) detecting endocardial motion by observing gray level change (e.g. [7]), 2) finding the boundary wall of the ventricle first and tracking it (e.g. - [8]), and 3) attempting to track landmarks (e.g. - implanted markers [9] or using MR signal tags [10].) There are significant problems with each of these previous approaches that affect their ability to accurately estimate wall-motion-derived clinical parameters. The primary objections include the inability to estimate point-wise motion and track specific points on the wall over time in the first two cases and the inherent problems of measuring the motion of a 3-D spatially-deforming object from a 2-D image in all cases.

The ability to track points on object surfaces over time has also been a principal goal of both rigid and non-rigid motion research in the field of computer vision. Snakes, or active contour models, have been the primary focus of many recent efforts (*e.g.* - [11], [12] and [13]). This method uses a physically based description of the contours to constrain the solution of movement to future frames. The physical model requires detailed prior knowledge of stiffness and mass characteristics which is often not available. Generalizations are often made to counter this problem, but this severely compromises the premise of the method in the first place. Additionally, all of these methods are frame-to-frame iterations and do not take advantage of any temporal trends. Finally, current active contour models do not take into account the periodic nature of our problem.

One method using a deformable model, [14], uses vibrational modes to parameterize its object of concern. However, these modes are not to be confused with the type of harmonics used in our method. The vibrational modes do not take into account any periodic nature associated with either the contour or its periodic deformation. Again, the method in [14] is a frame-to-frame iteration that does not incorporate extended temporal trends.

Recent efforts of a related nature to the approach described here have been proposed [15], [16]. These approaches differ in the term used for regularization of the flow vectors. Additionally, [15] uses an interpolation scheme to quantify shape between sub-sampled shape features. The work in [16] also optimizes a cost functional over a contour space as in our approach for a single contour, but maps shape-based contour correspondences back into Euclidean space values to perform smoothing of the flow vector field.

The technique of harmonic analysis and enhancement is well established [17]. The major thrust of research in this area has focused on one-dimensional signals where the fundamental frequency is unknown (*e.g.* - [18]).

1.2 Algorithm/Overview

Our approach to tracking and quantifying the motion of the LV wall is based upon several important assumptions. First, we assume that the endocardial contours must be extracted to provide an estimate of the location of the LV wall in each image frame over an entire cycle. Second, we assume that a set of reliable shape tokens can be derived from each of the extracted boundaries. Finally, we assume that motion is such that the tokens change only a small amount from frame to frame. Given these assumptions we have developed an algorithm for computing non-rigid object motion from a sequence of images.

The remainder of the paper reads as follows: first we explain the frame-to-frame tracking (or flow vector computation) of contours, next we present a model that incorporates the temporal constraints of our problem, and then we develop a statistical filter by combining the attributes of our temporal model with our solution for flow vector computation. We examine the results of this algorithm as applied to a typical MR study of a healthy canine heart and then offer an example of how this information can assist in regional LV function analysis. Finally, we discuss ongoing and planned future initiatives for our algorithm and summarize our findings.

2 FLOW VECTOR COMPUTATION

Using the image[1] sequence-derived contours that define the moving endocardial LV boundary, we now set out to compute flow vector fields that will aid in the tracking of point trajectories of contour pairs. For clarity, we begin by explaining the algorithm in this section using continuous mathematics, and lead up to a final implementation that uses a discrete time version of this equation. It is helpful to state the following definitions at the outset. First, for

[1] we typically use LV short axis slices, but the approach is generalizable to any 2D image sequence of the moving LV

the flow vector computation and for the sake of explanation, we note that we will consider two endocardial contours, $\mathcal{C}_i$ and $\mathcal{C}_{i+1}$, derived from successive image frames representing times t_i and t_{i+1} in an image sequence. These are part of a discrete set of contours representing point sets that exist over a set of T image frames in a two-dimensional Euclidean image space, i.e. $\mathcal{C}_i \subset \mathcal{R}^2$, $\forall i \in [0, ..., T]$. Furthermore, since each is actually a plane curve, we note that there exists a mapping for each curve from arclength in $\mathcal{R}^1$ back to the image space in $\mathcal{R}^2$. This mapping will be helpful in laying our full smooth flow vector algorithm below. We will use $s_i \in \mathcal{C}_i$ and $s_{i+1} \in \mathcal{C}_{i+1}$ to index individual points on the continuous curves in terms of arc length and m_i and m_{i+1} to index equal, unit-length samples along the discrete versions of the same curves. It will be useful later to note that the sets $\mathcal{C}_i$ and $\mathcal{C}_{i+1}$ are ordered and have values that increase monotonically in terms of their arc length. We assume that our temporal sampling is fine enough that the length of a contour is approximately the same between any two frames[2]. Thus, if L_i is the length of the contour found in frame t_i, we assume that $L_i \simeq L_{i+1}$.

2.1 Local Segment Matching

Correspondences between points on the initial contour $\mathcal{C}_i$ and points on a successive contour $\mathcal{C}_{i+1}$ in the sequence are found by matching shape properties of contour segments surrounding each of the points as discussed in [1]. We use a fixed-size portion of the contour at time t_{i+1} to define the search window, $W(s_i)$ for any point on the continuous contour found at time t_i. Rather than estimating a continuous mapping here, we use a subset of discrete samples on both the initial contour, where each sample point is denoted by m_i, and the contour from the following frame, where now each sample point is denoted by m_{i+1}. Due to our assumption of approximately equal-length contours at both times, we can initially index the sample points on both contours as $N+1$ equally-spaced samples of unit length such that $m_i, m_{i+1}, m \in [0, N]$. Thus, these initial contour samples are monotonically increasing, ordered sets. Now, we set up the correspondence problem as one of finding the points on the contour at time t_{i+1} that best match each of the $N+1$ points m_i. Due to the equal length, coincident index assumption above, this problem can be viewed as finding the deviation from a one-to-one mapping of the mapping of m_i into m_{i+1}, $\forall m_i \in [0, N]$. Thus, we will use the notation m_{i+1}^{shape} to denote the point on the contour at time t_{i+1} that best matches each monotonically ordered point m_i $(= m)$. Furthermore, we note that $m_{i+1}^{shape} = m + \delta_m^{shape}$, where δ_m^{shape} is the optimal deviation from the original one-to-one mapping at original sample point m.

A discrete version of the shape-based matching problem defined in [1] reduces to finding δ_m^{shape} for each sample m using the following discrete minimization:

$$\begin{aligned} \delta_m^{shape} &= \arg\min_{\delta_m} e_{BE}(\delta_m) \\ &= \arg\min_{\delta_m} \frac{1}{2} EI \sum_{m' \in M'} \left[\kappa_{i+1}(m' + \delta_m) - \kappa_i(m')\right]^2 \end{aligned} \tag{1}$$

where δ_m is allowed to vary such that the possible match choices for each point $m_i = m$ on the contour at time t_i are limited to points on the contour at time t_{i+1} within a range $(m + \delta_m) \in W(m_i)$. Here, $\kappa_i(m)$ defines the curvature at point m computed from the discrete contour data found at time t_i. We note that the index m' incorporates the overall matching problem index m, and defines the discrete contour segments surrounding any reference point $m_i = m$ or any candidate point $m_{i+1} = m + \delta_m$. Thus, $m' \in M'$, where the set $M' = \{m - \gamma, ..., m, ..., m + \gamma\}$. The parameter γ is defined as half of index length of the discrete version of the contour segment length L' defined earlier, and has been initially fixed to be equal to 10% of the window size $W(m_i)$. Finally, the constants E and I represent Young's modulus and the inertial component, respectively, associated with bending energy. See also [1].

For the method described below, we find it useful to define a vector $\mathbf{m}_{i+1}^{shape}$ that is of dimension $N+1$ and is a stacking of the correspondences found for each original $m_i = m = [0, N]$. In other words $\mathbf{m}_{i+1}^{shape} = \mathbf{m} + \left[\delta_0^{shape}\, \delta_1^{shape}, \ldots, \delta_m^{shape}, \ldots, \delta_N^{shape}\right]^T$ It is also useful to note that the correspondences determined above essentially represent discrete samples of a continuous function $\phi_{shape}(s_i)$ that maps all points s_i into s_{i+1}.

We note that temporal trajectories for a given point m_i can be derived over a set of periodic frames by both starting with frame $i = 1$, matching it with frame $i = 2$ and proceeding to frame $i = T$ or theoretically we could start with frame $i = 1$, match it with frame $i = T$ and proceed to frame $i = 2$. Obviously, since this knowledge

[2] obviously this is not true, but is a reasonable approximation

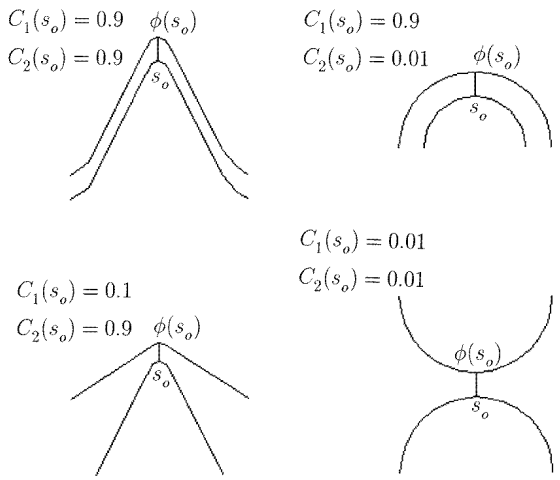

Figure 1: Examples of the four extreme combinations of our two confidence values. The upper left figure represents a close, unique match, hence both C_1 and C_2 are high. The upper right figure depicts a close but non-unique match, hence C_1 is high while C_2 is low. The lower left figure shows a situation where the match is not ideal but is unique among its neighbors, therefore C_1 is low and C_2 is high. Finally, in the lower right we have a situation where the match is poor and not unique, hence both C_1 and C_2 are low.

of periodicity is not incorporated into our initial shape matches the trajectories developed in the former (forward, labeled $\mathbf{m}_i^f$) direction will not exactly match those generated in the latter (backward, labeled $\mathbf{m}_i^b$) direction. To account for this we run our algorithm in both directions and then weight the two matches based on their distance from their initial frame as follows:

$$\mathbf{m}_i = \frac{T+1-i}{T}\mathbf{m}_i^f + \frac{i-1}{T}\mathbf{m}_i^b \tag{2}$$

Corresponding confidence measures are currently weighted in a similar manner although we are developing a method that uses the initial bending energies to recompute the new corresponding confidences. Current efforts of ours are focusing on modeling the number of frames considered in this forward-backward approach (currently T) as a variable parameter of scale, where matches are designated based on their fit in a global to local perspective. This would lead to a more robust formulation than the present weighting method.

2.2 Confidence Measures

The need for confidence measures is brought about by notions of: 1) how matches in any one region are treated in the sense of their relative strength when compared with matches created in other local regions, and 2) how decisions are handled in regions where there are many plausible matches. Thus, the strength and uniqueness of a match is used to help drive the solution of the overall flow field. Guided by the work in [19], we define two confidence measures, both based upon the profile of the error measure within each search region $W(m_i)$. First, an overall confidence at each discrete sample m_i (could also be each continuous point s_i) in the best match vector (m_{i+1}^{shape}) found by searching within search region $W(m_i)$ is made dependent on the inverse of the value of the error measure at that point, $e_{BE}(\delta_m^{shape})$: $C_1(m_i) = \frac{1}{k_1 + k_2 e_{BE}(\delta_m^{shape})}$ where k_1, and k_2 are scaling constants. A second confidence measure is given by $C_2(m_i) = \frac{1}{k_1 + k_2 \sigma(m_i)}$, where $\sigma(m_i)$ is the standard deviation of the error measure within window $W(m_i)$, to weigh the contribution of optimal matches found to each point m_i , such that the match is weighted more heavily if it is a unique match, and less if it was a choice out of many equally-good matches in a search region. Figure 1 provides examples of the four extreme combinations of these confidences.

2.3 Finding Smooth, Shape-Matched Flow Fields

Given pairs of contours representing the boundary of the LV wall at two temporal points in an image sequence, the initial shape-based segment matching and confidence measures defined in the last two subsections, we now seek to estimate a set of flow vectors that describe the motion of the LV wall as best as can be approximated from this two-dimensional data (i.e. in the case of the heart we only have partial information about the problem, since the object is a 3D object that nonlinearly deforms over time). A natural additional model we can now add to the basic shape-based matching approach is that, with sufficient spatial sampling, we can assume that adjacent points on the contour deform smoothly between any two frames[3]. Thus, the final computation is an optimization problem. We seek to find a vector flow field that corresponds points on two contours found from two consecutive temporal frames in an image sequence by compromising between adhering to the points where the local shape of the contours best match and the model that adjacent points on the contour move smoothly.

2.4 Optimization in the Contour space

Given the initial shape-based segment match and confidence measures defined in the last two subsections, we now develop an approach that solves for a smooth, shape- matched flow field between any two contours in arclength space, s. The objective is to find the best functional estimate for mapping points s_i on the initial contour, $\mathcal{C}_i$, into points s_{i+1} on the next contour in the temporal sequence, $\mathcal{C}_{i+1}$. From above, we note that $\phi_{shape}(s_i)$ represents the continuous mapping, found by the piecewise linear interpolation of discretely corresponded points according to the bending energy metric described by equation (1)[4], of any point s_i on contour $\mathcal{C}_i$ into a point s_{i+1} on the contour from the next temporal frame, $\mathcal{C}_{i+1}$. Furthermore, $C_1(s_i)$ and $C_2(s_i)$ represent the confidences computed over fixed size search regions $W(s_i)$ on $\mathcal{C}_{i+1}$ for each and every continuous point s_i. Putting all of this together, we now define an optimization problem that will permit us to solve for a new, smooth mapping, $\phi^*(s_i)$, that is a smoothed version of the shape- based mapping:

$$\begin{aligned}\phi^* &= \arg\min_{\phi} F(\phi|\phi_{shape}) \\ &= \arg\min_{\phi} \int_0^L C_1(s_i)C_2(s_i)\left[\phi(s_i) - \phi_{shape}(s_i)\right]^2 ds_i + \int_0^L \left[\phi''(s_i)\right]^2 ds_i \end{aligned} \tag{3}$$

We note here that the smoothest possible mapping of m_i onto m_{i+1} for all m is a linear one-to-one mapping. We adhere to this by including a term dependent on the second derivative of the mapping into the optimization functional. We express this functional in discrete form by using a sum for the integral and a second order difference for the smoothing term as follows:

$$\begin{aligned}\phi^* &= \arg\min_{\phi} F(\phi|\phi_{shape}) \\ &= \arg\min_{\phi} \sum_{m_i=0}^{N} C_1(m_i)C_2(m_i)\left[\phi(m_i) - \phi_{shape}(m_i)\right]^2 + \left[2\phi(m_i) - \phi(m_i+1) - \phi(m_i-1)\right]^2 \end{aligned} \tag{4}$$

This discrete problem can be reposed conveniently in vector notation. We define a diagonal confidence matrix, $\mathbf{C}$ as

$$\begin{pmatrix} C_1(1)C_2(1) & 0 & \cdots & 0 \\ 0 & \ddots & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \cdots & 0 & C_1(N)C_2(N) \end{pmatrix}$$

where diagonal elements are equal to the product of the two confidences for a given m_i.

[3] note that temporally smooth changes from frame to frame for any one point are implicit in the small- deformation- between- frames assumption in the bending energy model

[4] recall that $\mathbf{m}_{i+1}^{shape} = \mathbf{m} + \left[\delta_0^{shape}\, \delta_1^{shape}, \ldots, \delta_m^{shape}, \ldots, \delta_N^{shape}\right]^T \equiv$ a sampled version of $\phi_{shape}(s_i)$

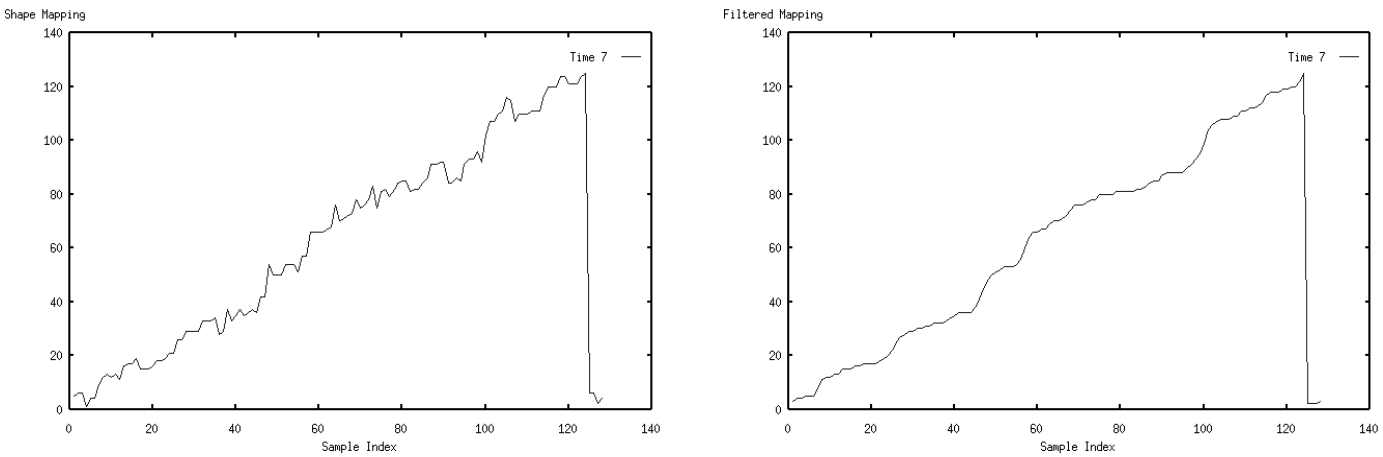

Figure 2: Samples of typical mapping functions for $N = 128$ samples. (Left) The shape-based mapping function. Note it's jagged appearance. Places where the slope is negative indicate vectors that cross. (Right) The filtered mapping function. Note how it maintains general shape qualities of the shape-based match but is monotonically increasing. The sharp drop at the end of the curve indicates where the wrap-around occurs.

If we use $\mathbf{m}_i^{shape}$ to denote the discrete values of $\phi_{shape}(m_i)$ stacked as a vector, and in similar fashion use $\mathbf{m}_i$ to denote $\phi(m_i)$ and $\mathbf{m}_i^*$ to denote the optimal estimate $\phi^*(m_i)$, we can express (4) as follows:

$$F = (\mathbf{m}_i - \mathbf{m}_i^{shape})^T \mathbf{C} (\mathbf{m}_i - \mathbf{m}_i^{shape}) + (\mathbf{D}\mathbf{D}\mathbf{m}_i)^T (\mathbf{D}\mathbf{D}\mathbf{m}_i) \tag{5}$$

where $\mathbf{D}$ is a first order differential operator given by

$$\begin{pmatrix} 1 & 0 & \cdots & \cdots & -1 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & 0 & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & -1 & 1 \end{pmatrix}$$

and $\mathbf{m}_i^{shape}$ is the initial shape-based mapping vector.

It can be shown that the matrix $(\mathbf{C} + \mathbf{D}^T\mathbf{D}^T\mathbf{D}\mathbf{D})$ is positive definite. By multiplying, combining, and eliminating non-related terms (see [3]) we find the optimal estimate for $\mathbf{m}_i$ that minimizes F is given as:

$$\mathbf{m}_i^* = (\mathbf{C} + \mathbf{D}^T\mathbf{D}^T\mathbf{D}\mathbf{D})^{-1}\mathbf{C}\mathbf{m}_i^{shape} \tag{6}$$

The vector $\mathbf{m}_i^*$ can be viewed as the nearest local minimum on this optimization surface which qualitatively represents a vector flow field that contains a compromise between 1) the best local matches for all points on the first contour to the second contour, with each match being weighted according to uniqueness of the shape in the match region and, 2) requiring each of these vectors to have a magnitude and direction that smoothly agrees with its spatially neighboring vectors. Additionally, posing the optimization problem in the contour space implicitly constrains the displacement vectors to the contours themselves. Figure 2 provides a comparison of the shape-derived $(\mathbf{m}_i^{shape})$ and filtered $(\hat{\mathbf{m}}_i^*)$ mapping functions for a given frame.

3 THE TEMPORAL MODEL

Our model for analyzing the temporal dynamics of endocardial wall motion is motivated by the desire to incorporate prior knowledge of periodicity. We note that our representation of endocardial motion is periodic in both contour space ($m_i \in [0, N]$) and time ($i \in [0, T]$). This is a unique aspect of the problem of tracking LV wall motion which has

not been exploited. Additionally, our model provides a stable state-space system which facilitates iterative filtering as will be discussed later.

The periodic mapping function described above can be viewed as a sum of sine waves whose frequencies are integral multiples of the lowest (or fundamental) frequency. Such a signal is said to be harmonic. If we consolidate all our continuous mapping functions into one two-dimensional function, $\phi(s,t)$, we can express the entire manifold as follows:

$$\phi(s,t) = \sum_{p=1}^{P}\sum_{r=1}^{R} C_{p,r}\sin(p\omega_o s + r\tau_o t + \psi_{p,r}) \tag{7}$$

where $\omega_o = \frac{2\pi}{N}$ is the fundamental frequency over the space of the contour and $\tau_o = \frac{2\pi}{T}$ is the fundamental frequency of the cardiac cycle over time. The index p represents harmonics of ω_o, and similarly r indexes τ_o. P and R represent the highest significant harmonics of ω_o and τ_o respectively. $C_{p,r}$ and $\psi_{p,r}$ are the amplitude and phase of the (m,n)th harmonic component of $\phi(s,t)$ respectively.

Through standard trigonometric manipulation this can be expressed as:

$$\phi(s,t) = \sum_{p=1}^{P}\sum_{r=1}^{R} g_{p,r}\cos(p\omega_o s + r\tau_o t) + h_{p,r}\sin(p\omega_o s + r\tau_o t) \tag{8}$$

where

$$g_{p,r} = C_{p,r}\sin(\psi_{p,r}),\quad h_{p,r} = C_{p,r}\cos(\psi_{p,r}) \tag{9}$$

For the purpose of developing a system equation, (8) can be converted into a vector product in the following manner for a given frame i:

$$\mathbf{m}_i = \Xi_i \eta \tag{10}$$

where $\mathbf{m}_i$ is the mapping function to frame i,

$$\eta = [g_{1,1}...g_{1,R}, g_{2,1}...g_{2,R}...g_{P,R}, h_{1,1}...h_{P,R}]^T \tag{11}$$

and Ξ_i is

$$\begin{pmatrix} \cos(\omega_o + \tau_o i) & \cos(\omega_o + 2\tau_o i) & \cdots & \sin(P\omega_o + R\tau_o i) \\ \vdots & \vdots & \vdots & \vdots \\ \cos(\omega_o m + \tau_o i) & \cos(\omega_o m + 2\tau_o i) & \cdots & \sin(P\omega_o m + R\tau_o i) \\ \vdots & \vdots & \ddots & \vdots \\ \cos(\omega_o N + \tau_o i) & \cos(\omega_o N + 2\tau_o i) & \cdots & \sin(P\omega_o N + R\tau_o i) \end{pmatrix}$$

Ξ is of dimension $N \times 2PR$ and η of dimension $2PR \times 1$. Since η theoretically represents a parameterization of the entire manifold it is suitable for use as a state vector as will be discussed in the next section.

4 RECURSIVE FILTERING

Having defined a relationship between the "observed" shape-based matches and our desired mapping solution as well as an extended temporal model for this system, we now derive an observation equation and subsequent iterative filter to exploit these properties. Because we are dealing exclusively with deterministic data at the present time we will use the deterministic equivalent of the Kalman filter, the Recursive Least Squares (RLS) filter, to derive our final temporally developed solution. We are currently investigating a temporal confidence term that would be used analogous to the observation noise covariance matrix. Additionally, we intend to develop a statistical distribution of the state vector based on empirical measurements. We anticipate inclusion of these last two quantities will allow us to evolve the RLS filter into a Bayesian Kalman filter [20].

The harmonic state vector, η, is particularly well-suited for its assignment due to the fact that $\eta_{i+1} = \eta_i$ for any given frame i. This relationship implies non-existence of any system noise. η will evolve as additional measurements are included in the estimation process.

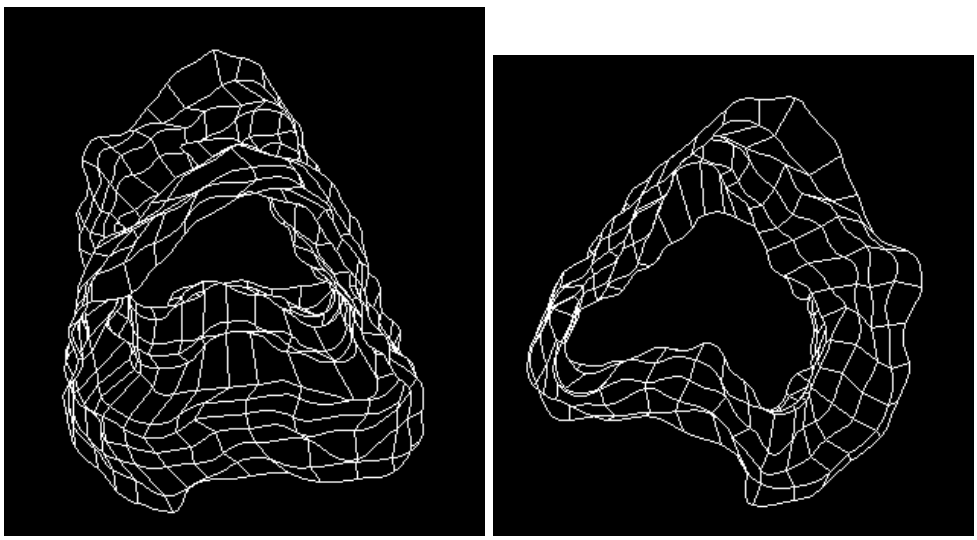

Figure 3: Two views of a complete cycle of sixteen contours segmented from MR images of a LV slice of a healthy canine heart. The contours are stacked as a surface with time as the third dimension to aid in visualizing shape trends.

First, we establish a relationship that will define our observation equation. If we assign $\mathbf{H} = (\mathbf{C}+\mathbf{D}^T\mathbf{D}^T\mathbf{D}\mathbf{D})^{-1}\mathbf{C}$, and note that $\mathbf{H}$ is positive definite, we can combine (6) and (10) as follows:

$$\mathbf{m}_i^{shape} = \mathbf{H}_i^{-1}\Xi_i\eta + \nu_i \tag{12}$$

ν_i represents observation noise at frame i due to truncation of harmonics and the accompanying ill-determined relationship ($\dim\mathbf{m_i} \neq \dim\eta$). (12) is considered the observation equation because it relates the data-derived shape-based mapping to the state vector.

The development of the RLS filter follows from [21]. We assign $\mathbf{G}_i = \mathbf{H}_i^{-1}\Xi_i$. Determination of the filter gain matrix is expressed as:

$$\mathbf{K}_i = \mathbf{P}_{i-1}\mathbf{G}_i^T(\mathbf{I}_N + \mathbf{G}_i\mathbf{P}_{i-1}\mathbf{G}_i^T)^{-1} \tag{13}$$

This is followed by the state covariance matrix update:

$$\mathbf{P}_i = \mathbf{P}_{i-1} - \mathbf{K}_i\mathbf{G}_i\mathbf{P}_{i-1} \tag{14}$$

Finally, the new estimate of the state vector is formed:

$$\hat{\eta}_i = \hat{\eta}_{i-1} + \mathbf{K}_i(\mathbf{m}_i - \mathbf{G}_i\hat{\eta}_i) \tag{15}$$

The solution mapping is retrieved after T iterations of the RLS filter have been computed by carrying out the following assignment $\forall i$:

$$\hat{\mathbf{m}}_i^* = \Xi_i\hat{\eta} \tag{16}$$

This final mapping represents a shape-driven, spatially and temporally smooth, and periodically constrained solution to the posed problem.

5 RESULTS

To illustrate our algorithm, we employ it in analyzing a set of sixteen contours from MR images of the complete cycle (end-diastole (ED) to end-systole (ES) and back to ED) of a healthy canine LV (figure 3). These contours were derived automatically using the method described in [22]. Each frame consisted of 128 (= N) samples. These sample spacings were chosen through repeated experimental trials which determined them to be sufficient to capture

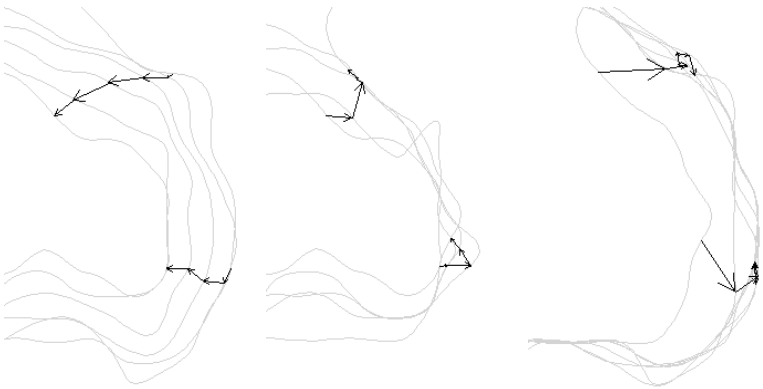

Figure 4: Trajectories of markers implanted in the LV wall. Trajectories are divided into three sections to avoid cluttering the graphic. (Left) Trajectories for frames 1-6. (Middle) Frames 6-10. (Right) Frames 10-16.

Figure 5: Corresponding trajectories estimated with a frame-to-frame shape-based tracking algorithm without any periodicity constraints (shown in black) compared with marker trajectories (shown in grey). (Top) Frames 1-6. (Middle) Frames 6-10. (Bottom) Frames 10-16.

Figure 6: Corresponding trajectories estimated with the recursive filter presented in this paper (in black) compared with marker trajectories (in grey). (Top) Frames 1-6. (Middle) Frames 6-10. (Bottom) Frames 10-16.

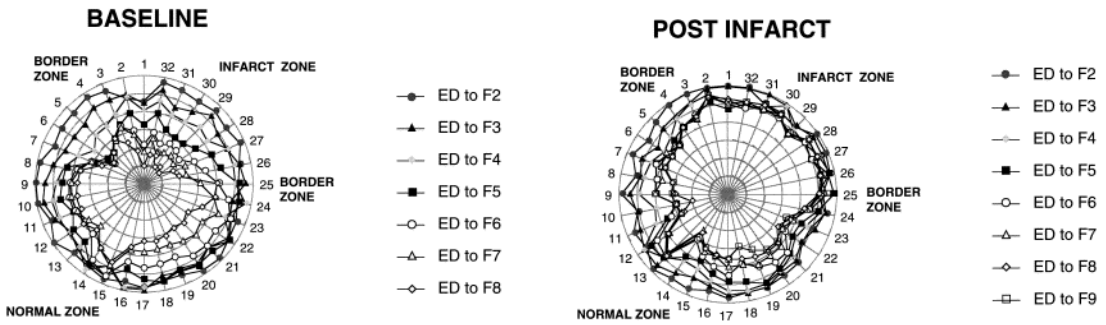

Figure 7: Radial plots of the extent of motion of equally spaced points around the endocardial wall. The points were tracked over 8 frames of a 16 frame cycle from ED to ED. (Left) Distance from starting point (ED) of a healthy canine heart. (Right) Distance from starting point of the same heart after infarction. Note the difference in wall motion around the infarct zone (upper right) before and after infarction. Of particular importance is the condition of the surrounding regions denoted as the border zones.

cardiac features and change. A token length of 5 ($= M'$) samples was used to determine each shape match over a window of 9 ($= W(m_i)$) samples (token ends extending past the limit of the window). $N/4$ harmonics were used in the m space ($P = 32$) and T harmonics were used for the t space ($R = 16$). When filtering we used a sample subset of N (typically 32 samples). Thus, Ξ was of size 32×1024 and η of size 1024×1. These numbers were chosen to limit computation time while still avoiding the distortion introduced by exclusively lower harmonics.

We immediately observe the large number of temporal trends occurring over the contour set in figure 3. We also note how well our algorithm does in following these trends while maintaining reasonable separation between the various trajectories.

In figure 4, we highlight the track of two markers that were implanted on the LV wall of this same canine heart. Figure 5 compares the corresponding pair of trajectories (i.e. - those with the same starting point) obtained using our old shape-based tracking algorithm without any temporal filtering. Finally, figure 6 compares the results of our new recursive filtering algorithm for the same trajectories.

We observe that the old shape-based tracking algorithm diverges from the actual track, especially in the case of the upper trajectory which completes the cycle far from it's initial starting position. The lower trajectory is somewhat better but still strays near ES.

On the other hand, when our periodicity constraint is incoporated the estimated trajectory is much closer to the actual track and in many instances exactly matches it. Even in situations where the tracks have diverged we note that filtered estimate still follows the general flow of the actual track. This is especially remarkable in the final few frames where significant jitter in the actual track is noted. The most obvious reason for any divergence is that we are trying to track the two-dimensional projection of what is actually a three-dimensional motion.

5.1 Analysis of LV Wall Motion

As an example of how our motion tracking algorithm is used in related analysis, we include figure 7. The frame-to-frame tracking procedure is applied to contours of a canine MR study before and after onset of an infarct. The graphs of figure 7 are radial plots for each trajectory and illustrate absolute distance from the starting point of the trajectory's movement. Comparing the two plots, the region of infarction is easily identified by the lack of relative motion in the area. More important though, are regions surrounding the infarct which contain hibernating tissue. Identification of these regions through their associated motion can assist in determining if surgical repair will be successful.

6 SUMMARY

We have a presented a novel tracking algorithm for analyzing the motion of periodically deforming contours such as those from images of the endocardial LV wall. This algorithm uses a technique that has not seen prior use in computer vision literature - that of multi-dimensional harmonic analysis. Using this technique we developed a method of filtering our data-derived shape-based mapping function to obtain a more desirable, smoothly varying mapping function. This final mapping function is the result of a process that combines feature adherence, spatial smoothness and a knowledge of temporal trends into one iterative filter.

We presented results based on actual MR data that show considerable promise when compared to the trajectories of implanted markers and previous trajectory estimates. Of course, we cannot expect to achieve precisely the same motion correspondeces since the true trajectory extends in three spatial dimensions, however we are encouraged these results and intend to integrate them with concurrent work of our group in matching three-dimensional surfaces.

Our near-term efforts are focused on developing a temporal confidence matrix based on the amount of correlation existing in the shape-based matches of our forward-backward approach. This development will be all the more powerful if we can also apply the idea using the number of forward-backward frames as a parameter of scale. This temporal confidence matrix can be used as a observation noise covariance matrix allowing us to convert our RLS filter into a deterministic Kalman filter.

We are also examining the possibility of a scale-space approach in determining the best shape match. This is done by using token length and/or search window size as parameters of scale. This will also have a significant effect on the accuracy of confidence measures.

On a larger scale, we hope to obtain empirical statistical distributions for our system state vector, η, based on data from implanted markers or other means of physical measurement. This would allow us to take our filter to a further level of stochastic inference, a Bayesian Kalman filter. Empirical investigation will also allow us to do more accurate verification of actual wall trajectories. Prior attempts have been difficult due to the sparse placement of physically implanted markers.

7 ACKNOWLEDGMENTS

The authors would like to thank Dr. Albert Sinusas and Don Dione of the Department of Cardiology, Yale School of Medicine, for supplying the radial plots used in figure 7.

This research is supported in part by NIH Grant T15 LM07056 from the National Library of Medicine and by NIH Grant R01HL44803 from the National Heart, Lung and Blood Institute. The work of A. Nehorai was supported by the Air Force Office of Scientific Research under Grant F49620-93-1-0096, the Office of Naval Research under Grant N00014-91-J-1298 and the National Science Foundation under Grant MIP-9122753.

References

[1] J. S. Duncan, R. L. Owen, L. H. Staib, and P. Anandan. Measurement of non-rigid motion using contour shape descriptors. In *Computer Vision and Pattern Recognition*, pages 318–324, Lahanai, Maui, HI, June 1991. IEEE.

[2] John C. McEachen and James S. Duncan. A constrained analytic solution for tracking non-rigid motion of the left ventricle. In *IEEE 18th Annual Northeast Bioengineering Conference*, pages 137–138, Univ. of Rhode Island, March 1992.

[3] John C. McEachen and James S. Duncan. Shape-based tracking of naturally occurring annuli in image sequences. In *Computer Vision and Pattern Recognition*, pages 613–614, New York City, June 1993. IEEE Computer Society.

[4] Thomas S. Huang, editor. *IEEE Workshop on Visual Motion*. IEEE Computer Society, Princeton, NJ, 1991.

[5] J. K. Aggarwal and R. O. Duda. Computer analysis of moving polygonal images. *IEEE Trans. on Computers*, 24(10):966–976, 1975.

[6] G. Oster and A. Perlson. The physics of cell motility. *Suppl. to J. Cell Science*, pages 35–54, 1988.

[7] S. Song and R. Leahy. Computation of 3d velocity fields from 3d cine ct images. *IEEE Trans. on Med. Imag.*, 10:295–306, Sept 1991.

[8] H. Gelberg, B. Brundage, S. Glantz, and W. Parmley. Quantitative left ventricular wall motion analysis: A comparison of area, chord and radial methods. *Circulation*, 59(5):991–1000, 1979.

[9] C. Slager, T. Hooghoudt, P. Serruys, J. Schuurbiers, and J. Reiber et al. Quantitative assessment of regional left ventricular motion using endocardial landmarks. *JACC*, 7(2):317–326, 1986.

[10] E. Zerhouni and et. al. Tagging of the human heart by multiplanar selective RF saturation for the analysis of myocardial contraction. In *Abstracts of the Ann. Meeting of the Soc. of MR in Imaging*, page 10, San Francisco, 1988.

[11] Demetri Terzopoulos and Richard Szeliski. Tracking with kalman snakes. In Andrew Blake and Alan Yuille, editors, *Active Vision*, pages 3–20, Cambridge, MA, 1992. The MIT Press.

[12] Chahab Nastar and Nicholas Ayache. Non-rigid motion analysis in medical images: a physically based approach. In H. H. Barrett and A. F. Gmitro, editors, *Information Processing in Medical Imaging*, pages 17–32, Flagstaff, AZ, June 1993. Springer-Verlag.

[13] Frederic Leymarie and Martin D. Levine. Tracking deformable objects in the plane using an active contour model. *IEEE Trans. on PAMI*, 15(6):617–634, June 1993.

[14] Alex Pentland, Bradley Horowitz, and Stan Sclaroff. Non-rigid motion and structure from contour. In *IEEE Workshop on Visual Motion*, pages 288–293, Princeton, NJ, October 1991. IEEE Computer Society.

[15] M. Demi, R. Calamai, G. Coppini, and G. Valli. A visual framework for the study of cardiac motion. *Computers in Cardiology*, 1990.

[16] Isaac Cohen, Nicholas Ayache, and Patrick Sulger. *Tracking Points on Deformable Objects Using Curvature Information*, volume 588. Lecture Notes in Computer Science: ECCV '92. Springer-Verlag, Berlin, 1992.

[17] S. M. Kay and S. L. Marple Jr. Spectrum analysis - a modern perspective. *Proceedings IEEE*, 69:1380–1419, Nov. 1981.

[18] Arye Nehorai and Boaz Porat. Adaptive comb filtering for harmonic signal enhancement. *IEEE Trans. on ASSP*, 34(5):1124–1138, Oct. 1986.

[19] P. Anandan. A computational framework and an algorithm for the measurement of visual motion. *International Journal of Computer Vision*, 2:283–310, 1989.

[20] Harry L. VanTrees. *Detection, Estimation, and Modulation Theory (Part I)*. John Wiley and Sons, New York, 1968.

[21] T. Soderstrom and P. Stoica. *System Identification*. Prentice-Hall International, Englewood Cliffs, NJ, 1989.

[22] Lawrence H. Staib and James S. Duncan. Boundary finding with parametrically deformable models. *IEEE Trans. on PAMI*, 14(11):1061–1074, Nov. 1992.