A Recursive Filter for Phase Velocity Assisted Shape-based Tracking of Cardiac Non-rigid Motion*

John C. McEachen II[†], François G. Meyer[◇,†], R. Todd Constable[‡],
Arye Nehorai[†], and James S. Duncan[†,‡]

Departments of [†]Electrical Engineering, [◇]Mathematics, and [‡]Diagnostic Radiology
333 Cedar St.
P.O. Box 208042
Yale University School of Medicine
New Haven, CT 06520-8042

E-mail: mceachen@noodle.med.yale.edu

Categories:
Medical Computer Vision
Motion Analysis
Systems and Applications

Abstract

A framework for tracking pointwise periodic non-rigid motion of the heart's left ventricular (LV) wall is presented which incoporates information from two different magnetic resonance imaging (MRI) techniques. New developments in phase-contrast cine MR imaging have produced spatial maps of instantaneous velocity that have proven accuracy within the the myocardium, or wall, of the heart. This information is combined with shape-based matching techniques to provide improved estimates of trajectories, especially in regions where shape information is limited. These raw trajectories act as input to a recursive least squares (RLS)

*This research is supported in part by NIH Grant R01HL44803 from the National Heart, Lung and Blood Institute.

filter which applies the constraints of temporal periodicity and spatial smoothness for the final estimate. The results of the RLS filter are compared with the motion of actual implanted markers. Comparisons are also made between exclusively shape-based filtered and phase-contrast enhanced trajectory estimates using both phantom and actual canine heart MR images.

Introduction

Tracking non-rigid motion, particularly in medical image sequences, has received considerable interest within the computer vision community in recent years [1], [2]. Cardiac image analysis has seen a large share of this effort specifically in identifying myocardial damage due to heart disease. The requirement for non-invasive diagnosis of regional cardiac function has confounded the problem, however, leaving potential researchers with no concrete landmarks to assist them in their analysis. Cardiac MRI analysis in particular does afford the researcher some useful constraints which we explore in this paper. Specifically, the periodic nature of the cardiac cycle is well suited for the tools of harmonic analysis and is an area ripe for further study. We present a temporal model based on a sum of sinusoids that forms the basis for our (RLS) tracking filter. Additionally, new developments in phase-contrast cine MRI encoding have led to accurate spatial maps of velocity within the myocardium, or wall, of the LV. We employ a finite-element based mesh across the myocardium to capture the global or regional motion of the LV, then use this information to assist our traditional shape-based correspondence method ([3]) in refining our point-wise estimate of the frame-to-frame trajectories. We show that this method is particularly useful in regions with limited shape variation. The work in this paper is based on previous efforts in [4], [5] and [6].

Background

The ability to track points on object surfaces over time has 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.* - [7], [8] and [9]). 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, the above methods require a set of corresponding landmarks that known beforehand. Finally, all of these methods are frame-to-frame iterations and do not take advantage of any temporal trends. Current active contour models do not take into account the periodic nature of our problem.

One method using a deformable model, [10], uses vibrational modes to parametrize 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 [10] assumes known landmarks. It is a frame-to-frame iteration that does not incorporate extended temporal trends and is not pursuing point correspondence maps between image frames.

Recent efforts of a related nature to the approach described here have been proposed [11], [12]. These approaches differ in the term used for regularization of the flow vectors. Additionally, [11] uses an interpolation scheme to quantify shape between sub-sampled shape features. The work in [12] 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 [13]. The major thrust of research in this area has focused on estimating parameters of one-dimensional signals where the fundamental frequency is unknown (*e.g.* - [14]). Little has been explored in the area of enforcing periodic constraints which adds to the significance of the method presented below.

The problem of quantifying regional LV wall function is not new to the field of cardiology either. 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. [15]), 2) finding the boundary wall of the ventricle first and measuring segment changes (e.g. - [16]), and 3) attempting to track landmarks (e.g. - implanted markers [17] or using MR signal tags [18].) 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. In the case of MR signal tags, the spatial resolution of the measured motion depends on the pixel size of the tags and the methods have difficulty imaging the entire cardiac cycle.

Phase contrast approaches to MRI motion analysis have shown promising gains over previous MR signal encoding techniques such as spin tagging. Van Dijk ([19]) first suggested use of the NMR phase to record cardiac wall velocity measurements using a spin echo gated MR imaging sequence. Since then others ([20], [21]) have shown that velocity information can be encoded into the phase of the NMR signal and incorporated into a conventional cine sequence. Additionally, methods have been developed for tracking cardiac wall motion based solely on these velocity maps ([6]).

In phase contrast cine MRI the spectrometer cycles at a constant rate (asynchronous to the cardiac cycle) with timing of the phase-encoding increments driven by the cardiac cycle trigger.

Two sequences that interrogate the same section but contain different flow encoding are interleaved during each cardiac cycle. The differential flow encoding consists of different first moments of the magnetic field gradient in a selected direction. After Fourier transformation, pixel-by-pixel phase differences are proportional to velocity in the encoded direction. Gradient waveforms in the other directions are unchanged. These proportional velocities are taken as a pixel-by-pixel velocity map over the entire image. In cardiac MRI analysis, this velocity map has been shown to be highly accurate for regions embedded in the myocardium. Unfortunately, pixels near the boundaries of the LV become influenced by the spectral content of surrounding non-homogeneous material resulting in a considerable number of outliers. Thus the phase-contrast method in itself is not suitable for tracking point-wise trajectories along the LV wall. It does, however, provide us with considerable general information regarding global LV motion. It is in this manner that we incorporate the information into our method below.

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 is organized into six sections. In the first three we are concerned with the frame-to-frame tracking (or flow vector computation) of contours. This begins by considering our phase contrast velocity map as an physical flow field and representing the myocardium with a triangulation mesh that organized flow information into distinct cells. We draw on the insight of continuum mechanics and finite element analysis to extract parameters of motion for vertices lying on the boundary. The motion estimate for these points is then refined using our traditional shape-based matching method. We establish constraints for smoothness between trajectories and use this in defining a frame-to-frame regularized solution. Next we present a model that incorporates the temporal constraints of our problem, and then develop a statistical filter by combining the attributes of our temporal model with our solution for flow vector computation. Finally, we discuss ongoing and planned future initiatives for our algorithm and summarize our findings.

Analyzing Velocity Flow

The aim of this paper is to describe a method to track the deformation of material points inside the myocardium. The approach makes it possible to track and analyze basic temporal variations

of stretching and twisting of regions inside the myocardium. For each frame, we have access to the instantaneous velocity field. We emphasize the case where these measurements are provided by phase-contrast MRI, although we note that velocities can be directly estimated from the spatiotemporal changes of the intensity function [15]. In this method, the myocardium is divided into small elementary cells. Inside each cell the rate of deformation is assumed to be homogeneous. As we justify below, this makes it possible to describe locally the velocity field using an affine model. Accordingly, the internal deformation of the body is uniquely described by a deforming mesh. The numerical integration of the velocity at each node of the mesh yields new estimates of position at the next instant. With these new estimates of position we will be prepared to employ our shape-based matching method for the purposes of fine tuning our trajectories in the next section.

Continuous Piecewise Linear Model of the Velocity Field

We subdivide the myocardium into triangles. The boundary of the subdivision approximates as well as possible the real boundaries of the myocardium. The deformation in each triangle characterizes the continuum deformation of the body in this neighborhood. We shall now describe the velocity field inside each triangle, referring to a few concepts from continuum mechanics [22].

In our analysis of the deformable motion we assume that the rate of deformation tensor, $\mathbf{D}$, and the vorticity tensor, $\mathbf{W}$, are locally homogeneous. However we note that $\mathbf{D}$ and $\mathbf{W}$ will change from one neighborhood of the myocardium to another. Similar assumptions in the biomechanical literature have assumed that the deformation gradient tensor in the myocardium was locally homogeneous [23, 24, 25]. The physical assumption of local homogeneity of rate of deformation results in a local linear (or affine) model. Let $\mathbf{x}$ be a point in the deformed body. For any point $\mathbf{x}+\delta\mathbf{x}$ in a small neighborhood around $\mathbf{x}$ we have

$$\mathbf{u}(\mathbf{x}+\delta\mathbf{x},t) = \mathbf{u}(\mathbf{x},t) + \mathbf{L}(\delta\mathbf{x}) \tag{1}$$

The physical assumption of local homogeneity of rate of deformation results in the locally linear (or affine) model (1). This model appears to be a good tradeoff between the exhaustive description of the entire set of noisy observed velocities, and a crude model of the velocity, which would assume a constant motion inside each cell. Furthermore this model allows us to estimate the rate of deformation, an essential parameter in the analysis of the myocardial deformation.

Figure 1 is an example of a typical triangulation of the myocardium. Let Δ be the mesh associated with the triangulation. The collection of all the nodes defined by the partition Δ is indexed from 1 to N. We consider the space P^0_Δ of piecewise polynomial functions in $C^0(\mathbb{R}^2)$ in two arguments and of degree one. Any function in P^0_Δ is uniquely determined by its values on the nodes of the triangulation. This property follows from the fact that along the boundary of two cells a piecewise

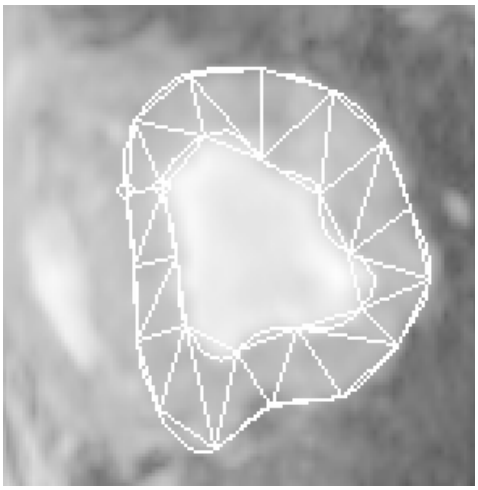

Figure 1: Triangulation mesh overlaid on the short-axial slice of a canine left ventricle (LV).

linear function reduces to a straight line. Formally, we can define a local basis of P_Δ^0. Let ξ_k be the linear function which takes the value 1 at the vertex k and vanishes on all other nodes. ξ_k has the shape of a pyramid. This basis was first suggested in [26], and has been extensively used since then for approximating the solutions of differential equations in finite elements methods [27]. Any function p in P_Δ^0 can then be expressed as

$$p = \sum_{k=1}^{N} \alpha_k \xi_k$$

The basis $(\xi_k)_{k=1\ldots N}$ is not orthogonal but the product $\int_\Delta \xi_k \xi_j$ is equal to zero each time that k and j are not in the same triangle. We can now express our model of the entire velocity field inside the myocardium. Inside each triangle the velocity field is described by the linear model (1). Across the boundaries the velocity is continuous. Each component of the velocity field is thus a function of P_Δ^0. We can expand the Eulerian model of the velocity field in the local basis $(\xi_k)_{k=1\ldots N}$. In 2-D we have

$$(u, v)(t) = \left(\sum_{k=1}^{N} u_k \xi_k, \sum_{k=1}^{N} v_k \xi_k \right)(t) \tag{2}$$

The advantage of this local basis is that the entire velocity field is completely described by the value of the velocity at each vertex of the triangulation. This important result sets up a one-to-one correspondence between the dense motion of the material points inside the medium, and the motion of the nodes of the mesh. We note that we are not working with a fixed grid, but rather with a deforming mesh that follows the myocardium. Therefore the position and the support of the function ξ_k will change in each frame.

Approximation of the Velocity Field

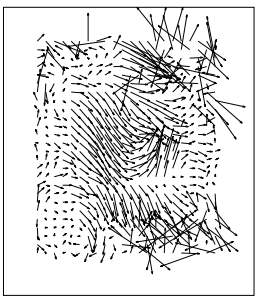

Figure 2: An example of the unprocessed velocity flow map generated from the phase-contrast cine MR image of the short-axial slice of a canine LV.

Next we describe the approximation of an observed vector field by a vector field of piecewise continuous linear functions in two variables. In fact this approximation plays two roles. First, it yields a local description of the velocities with a small number of crucial parameters. Second, it provides a regularization of the observed velocities. Indeed, velocity fields obtained from optical-flow method tend to be noisy. The quality of the estimate of the flow is usually not homogeneous. Velocity fields obtained from phase-contrast measurements are visually smoother but yield significant outliers on regions near the boundary of the myocardium, as shown in figure 2. A general framework for locally regularizing the velocity field has been developed. The global vector field is approximated locally by continuous linear piecewise vector fields. In each triangle the flow is described by (1). Furthermore the approximation will be continuous across the common boundary of two cells.

Given an observation of the velocity field $\mathbf{u} = (u, v)$, our goal is to estimate the best approximation of $\mathbf{u}$ in a space of smooth vector fields. We seek an approximation $\mathbf{p} = (p, q)$ of $\mathbf{u}$ in $P^0_\Delta \times P^0_\Delta$ that minimizes the error

$$< \mathbf{p} - \mathbf{u}, \mathbf{p} - \mathbf{u} >= \int_\Delta (\mathbf{p} - \mathbf{u})^2 = \int_\Delta (p - u)^2 + \int_\Delta (q - v)^2 \tag{3}$$

The two terms $\int_\Delta (p-u)^2$ and $\int_\Delta (q-v)^2$ are independent and can be minimized independently. We now have two similar minimization problems for each component of $\mathbf{u}$, that can be solved independently. In other words (3) is separable and the two components (u, v) can be approximated separately. Therefore we need only to explain how to solve the problem of minimizing

$$\int_\Delta (p - u)^2 \tag{4}$$

where u is a one scalar function, and $p \in P^0_\Delta$. We expand p in the local basis $(\xi_k)_{k=1\ldots N}$

$$p = \sum_{i=1}^{N} \alpha_i \xi_i \tag{5}$$

The problem of minimizing (4) can then be expressed using (5) as seeking the coordinates $(\alpha_i), i = 1, \ldots, N$ so that the following error will be minimized

$$\min_{\alpha_i \in \mathbb{R}^N} \int_{\Delta} \left(\sum_{i=1}^{N} \alpha_i \xi_i - u \right)^2 \tag{6}$$

This is a classic least-squares problem. Taking the partial derivative with respect to α_i, the solution of (6) is given by the normal equations

$$\mathbf{A} \quad = \mathbf{r} \tag{7}$$

where

$$= [\alpha_1, \ldots, \alpha_N]^T$$

$$A_{ij} =< \xi_i, \xi_j >, i, j = 1, \ldots, N \text{ and } r_i =< \xi_i, u > i = 1, \ldots, N$$

As mentioned before, (ξ_i) is a local basis, and most of the inner products, A_{ij} vanishes to zero. Thus the matrix $\mathbf{A}$ has a banded structure.There exist a very effective way to compute the matrix $\mathbf{A}$ and the vector r. The approach is directly derived from the theory of finite elements [27]. $\mathbf{A}$ needs to be calculated and inverted only once for both components of the velocity.

Estimation of displacements The velocity estimated at each node of the mesh is integrated in time, and yields an approximation of the position of the given node at the next instant. The numerical integration is based on a forward Euler method. Our analysis exploits the predicted position of the mesh. Provided the velocity estimates are accurate, the position of a given material point on $\mathcal{C}_{i+1}$ should lie in a small neighborhood around the predicted position $(x, y)(t_{i+1})|t_i)$. At this juncture, it is logical to look at similarities in shape to guide our final correspondence. Thus we discuss our shape matching algorithm below.

Shape-based Refinement

In the next two sections, we refer to the the image[1] sequence-derived contours that define the moving LV boundary with the intent of computing 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 the flow vector computation and for the sake of explanation, we note that we will consider two 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

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

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 $L_i \simeq L_{i+1}$.

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 in a manner similar to [3]. In the method of [3], a fixed-size portion of the contour at time t_{i+1} to define was used the search window, $W(s_i)$, for any point on the continuous contour found at time t_i. This portion of $\mathcal{C}_{i+1}$ chosen for $W(s_i)$ was that centered on the point of nearest Euclidean distance to the initial point of $\mathcal{C}_i$, as discussed in [12]. In our current method, the displacement estimate given by the phase contrast velocity flow map is used to guide the positioning of $W(s_i)$ on $\mathcal{C}_{i+1}$. This is done simply by taking the point on $\mathcal{C}_{i+1}$ of nearest Euclidean distance to *the Euclidean location* of the displaced estimate, $(x, y)(t_{i+1})|t_i)$, given by the integration of the phase velocity map for the initial point on $\mathcal{C}_i$. The search window, $W(s_i)$, is then centered on this point as shown in figure 3. We hypothesize that this approach allows us to incorporate global torque and translation given by phase velocity information with local shape continuity correspondence.

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

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

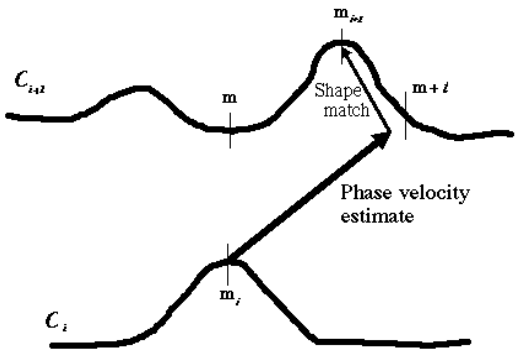

Figure 3: Combining the displaced estimate of the phase velocity map with shape-based matching. The thick arrow indicates the displacement of estimated by processing phase velocity. The thin arrow indicates the resulting shape match extended from the phase velocity estimate. $\mathbf{m}+t$ indicates where the search window, $W(s_i)$, is centered.

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 [3] reduces to finding δ_m^{shape} for each sample m using the following discrete minimization:

$$\begin{aligned}\delta_m^{shape} &= \arg\min_{\delta_m} e_{BE}(\delta_m) \\ &= \arg\min_{\delta_m} \frac{1}{2} EI \sum_{m'\in M'} \left[\kappa_{i+1}(m'+\delta_m) - \kappa_i(m')\right]^2 \end{aligned} \tag{8}$$

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 $(m+\delta_m) \in W(m_i)$. 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 $M' = \{m-\gamma+\tau, ..., m+\tau, ..., m+\gamma+\tau\}$. 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 $W(m_i)$. τ represents the offset in contour space due to the displacement estimate of the velocity flow map. The constants E and I represent Young's modulus and the inertial component, respectively, associated with bending energy. Finally, $\kappa_i(m)$ defines the curvature at point m computed from the discrete contour data found at time t_i. See also [3].

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 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:

$$\mathbf{m}_i = \frac{T+1-i}{T}\mathbf{m}_i^f + \frac{i-1}{T}\mathbf{m}_i^b \tag{9}$$

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.

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 [28], 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 4 provides examples of the four extreme combinations of these confidences.

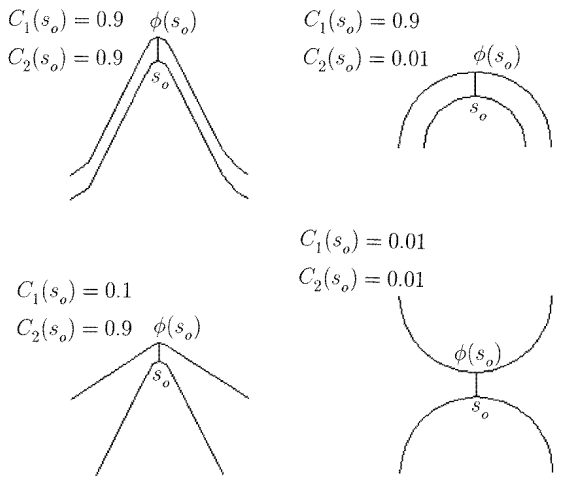

Figure 4: Examples of the four extreme combinations of our two confidence values. The upper left figure represents a close, unique match, hence both C_1 and C_2 are high. The upper right figure depicts a close but non-unique match, hence C_1 is high while C_2 is low. The lower left figure shows a situation 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.

Flow Vector Computation

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 guided by velocity flow and confidence measures defined in the last two sections, 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.

Optimization in the Contour space.

Given the initial shape-based segment match and confidence measures defined in the last two sub-

[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

sections, 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 (8)[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:

$$\begin{aligned}\phi^* &= \arg\min_{\phi} F(\phi|\phi_{shape}) \\ &= \arg\min_{\phi} \int_0^L C_1(s_i)C_2(s_i)\left[\phi(s_i) - \phi_{shape}(s_i)\right]^2 ds_i + \int_0^L \left[\phi''(s_i)\right]^2 ds_i \end{aligned} \tag{10}$$

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:

$$\begin{aligned}\phi^* &= \arg\min_{\phi} F(\phi|\phi_{shape}) \\ &= \arg\min_{\phi} \sum_{m_i=0}^{N} C_1(m_i)C_2(m_i)\left[\phi(m_i) - \phi_{shape}(m_i)\right]^2 + \left[2\phi(m_i) - \phi(m_i+1) - \phi(m_i-1)\right]^2 \end{aligned} \tag{11}$$

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.

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 (11) as follows:

$$F = (\mathbf{m}_i - \mathbf{m}_i^{shape})^T \mathbf{C}_{i+1} (\mathbf{m}_i - \mathbf{m}_i^{shape}) + (\mathbf{DDm}_i)^T (\mathbf{DDm}_i) \tag{12}$$

[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 $\phi_{shape}(s_i)$

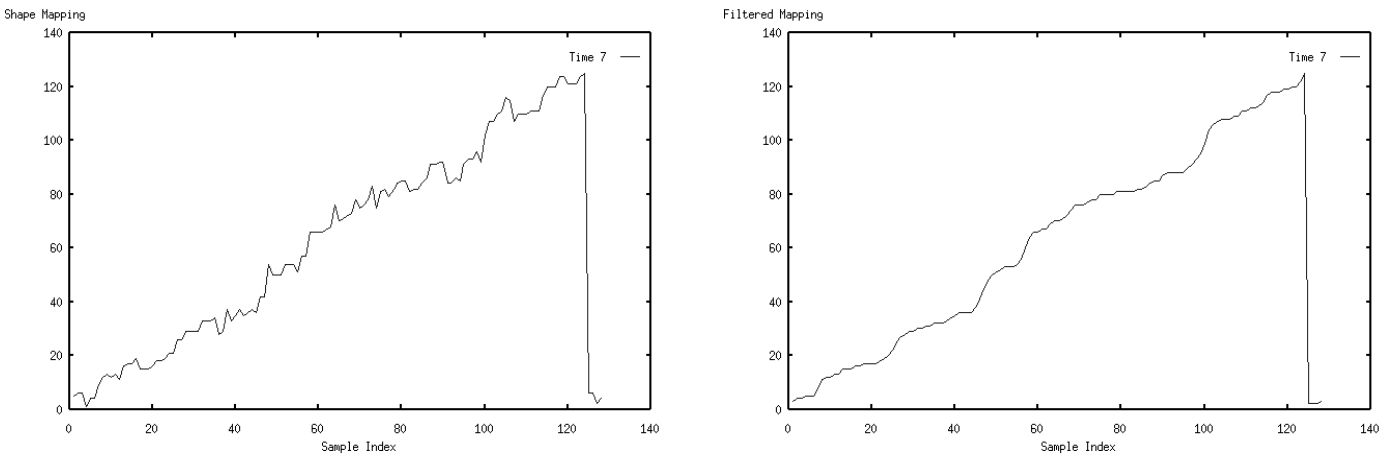

Figure 5: Samples of typical mapping functions for $N = 128$ samples. (Left) The shape-based mapping function. Note it's jagged appearance. Places where the slope is negative indicate vectors that cross. (Right) The filtered mapping function. Note how it maintains general shape qualities of the shape-based match but is monotonically increasing. The sharp drop at the end of the curve indicates where the wrap-around occurs.

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 [29]) we find the optimal estimate for $\mathbf{m}_{i+1}$ that minimizes F is given as:

$$\mathbf{m}_{i+1}^* = (\mathbf{C}_{i+1} + \mathbf{D}^T\mathbf{D}^T\mathbf{D}\mathbf{D})^{-1}\mathbf{C}_{i+1}\mathbf{m}_{i+1}^{shape} \tag{13}$$

The vector $\mathbf{m}_{i+1}^*$ 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 5 provides a comparison of the shape-derived ($\mathbf{m}_i^{shape}$) and filtered ($\hat{\mathbf{m}}_i^*$) mapping functions for a given frame.

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:

$$\phi(s,t) = \sum_{p=1}^{P} \sum_{r=1}^{R} C_{p,r} \sin(p\omega_o s + r\theta_o t + \psi_{p,r}) \tag{14}$$

where $\omega_o = \frac{2\pi}{N}$ is the fundamental frequency over the space of the contour and $\theta_o = \frac{2\pi}{T}$ is the fundamental frequency of the cardiac cycle over time. The index p represents harmonics of ω_o, and similarly r indexes θ_o. P and R represent the highest significant harmonics of ω_o and θ_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:

$$\phi(s,t) = \sum_{p=1}^{P} \sum_{r=1}^{R} g_{p,r} \cos(p\omega_o s + r\theta_o t) + h_{p,r} \sin(p\omega_o s + r\theta_o t) \tag{15}$$

where

$$g_{p,r} = C_{p,r} \sin(\psi_{p,r}), \quad h_{p,r} = C_{p,r} \cos(\psi_{p,r}) \tag{16}$$

For the purpose of developing a system equation, (15) can be converted into a vector product in the following manner for a given frame i:

$$\mathbf{m}_i = \Xi_i \eta \tag{17}$$

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

$$\eta = [g_{1,1} \ldots g_{1,R}, g_{2,1} \ldots g_{2,R} \ldots g_{P,R}, h_{1,1} \ldots h_{P,R}]^T \tag{18}$$

and Ξ_i is

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

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.

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 [30].

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.

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 (13) and (17) as follows:

$$\mathbf{m}_i^{shape} = \mathbf{H}_i^{-1}\Xi_i\eta + \nu_i \tag{19}$$

ν_i represents observation noise at frame i due to truncation of harmonics and the accompanying ill-determined relationship ($\dim \mathbf{m_i} \neq \dim \eta$). (19) 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 [31]. We assign $\mathbf{G}_i = \mathbf{H}_i^{-1}\Xi_i$. Determination of the filter gain matrix is expressed as:

$$\mathbf{K}_i = \mathbf{P}_{i-1}\mathbf{G}_i^T(\mathbf{I}_N + \mathbf{G}_i\mathbf{P}_{i-1}\mathbf{G}_i^T)^{-1} \tag{20}$$

This is followed by the state covariance matrix update:

$$\mathbf{P}_i = \mathbf{P}_{i-1} - \mathbf{K}_i\mathbf{G}_i\mathbf{P}_{i-1} \tag{21}$$

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

$$\hat{\eta}_i = \hat{\eta}_{i-1} + \mathbf{K}_i(\mathbf{m}_i - \mathbf{G}_i\hat{\eta}_i) \tag{22}$$

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

$$\hat{\mathbf{m}}_i = \Xi_i \hat{\eta} \tag{23}$$

This final mapping represents a shape-driven, spatially and temporally smooth, and periodically constrained solution to the posed problem. Figure 6 reviews the steps of this filtering process in flowchart form.

Results

To illustrate our algorithm, we employ it in analyzing a set of eight 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 7). These contours were derived automatically using the method described in [32]. Each frame consisted of 128 ($= N$) samples. These sample spacings were chosen through repeated experimental trials which determined them to be sufficient to capture 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 ($R = 8$). 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 7. We also note how well our algorithm does in following these trends while maintaining reasonable separation between the various trajectories.

In figure 8, we highlight the track of two markers that were implanted on the LV wall of this same canine heart. Phase contrast velocity information was not available in this experiment so we will concentrate on verifying the usefullness of our filtering approach. Figure 9 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 10 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 incorporated 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

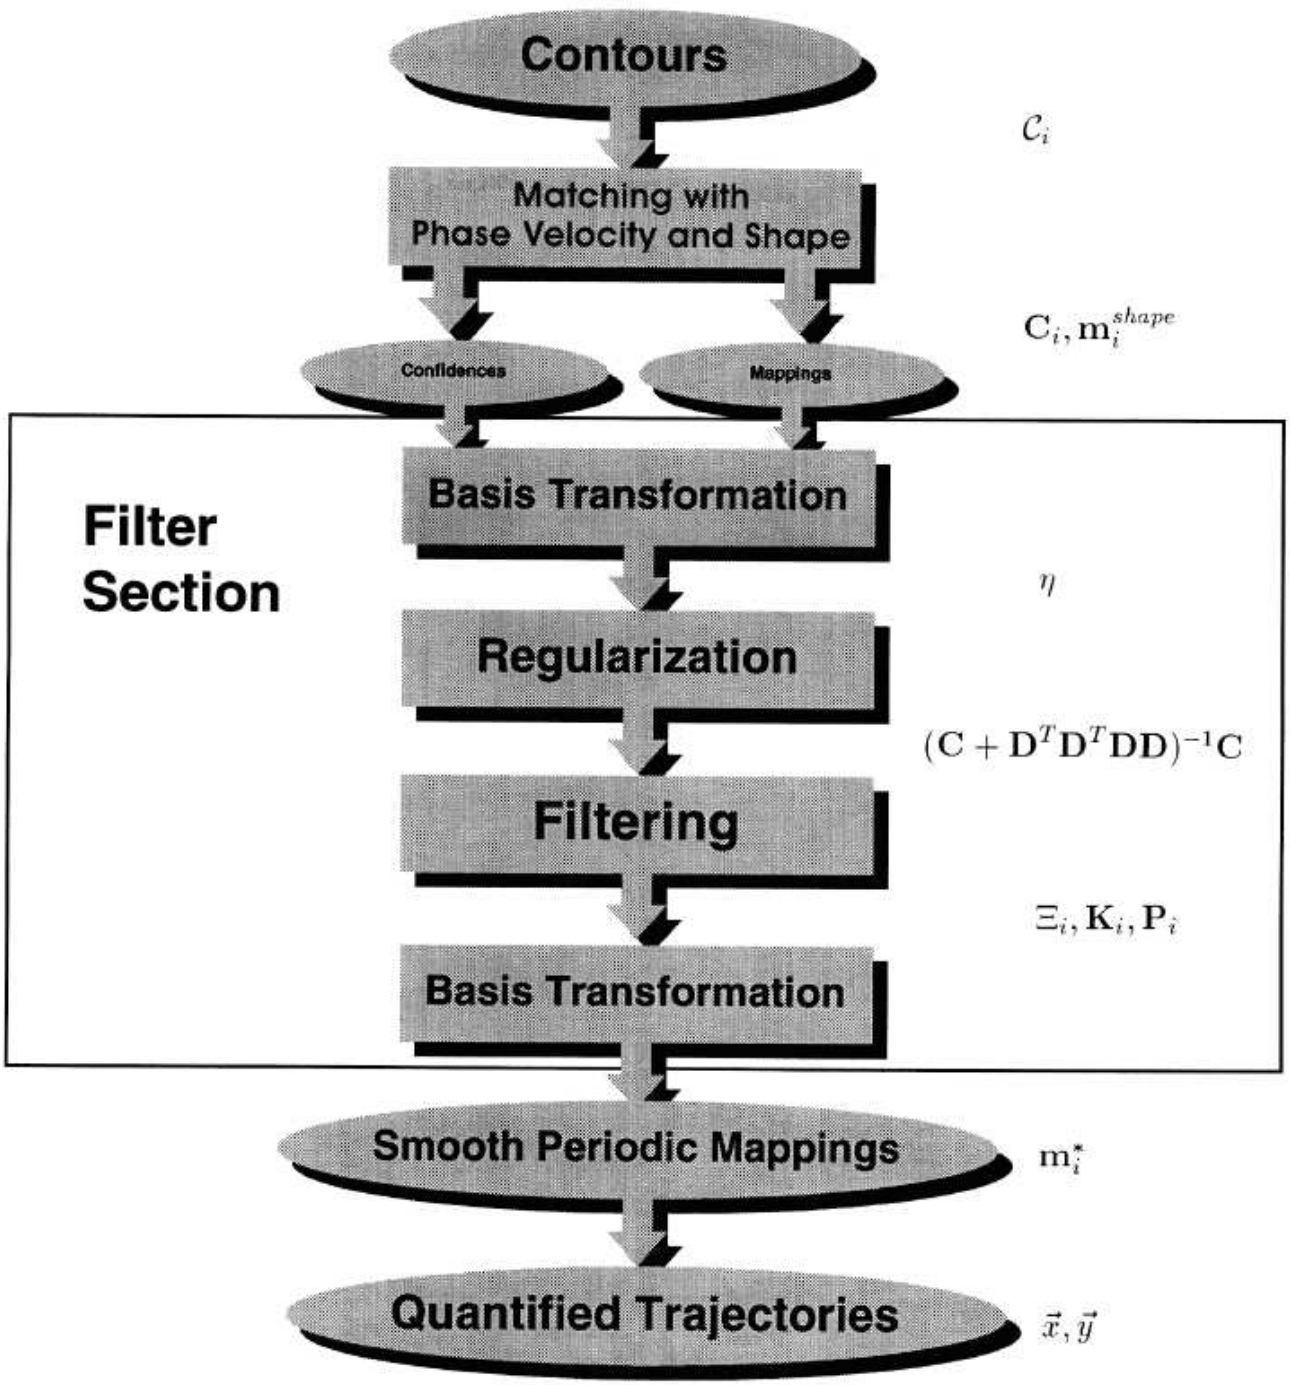

Figure 6: The RLS filtering process.

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.

Figure 11 shows our analysis using MR phase velocity map alone on an imaged cylindrical phantom undergoing rigid-body motion. We see that in areas near the landmarks the velocity map becomes extremely noisy. If the landmarks are ignored and a mesh is developed that accounts for flow strictly internal to the material of the phantom, we see that the mesh deformation itself does a suitable job in tracking the motion of the phantom. We stress that this image sequence is a low-noise circumstance with well defined boundaries and is not typical of cardiac acquisitions.

The effect of a complete absence of shape information is readily seen when our shape-based algorithm alone is applied to the inner boundary of the phantom in figure 12. When the phase velocity information is incorporated, we see a marked improvement. A similar situation occurs when we examine the results of our methods on the outer phantom boundaries in figure 13. In this situation we do have four landmarks and we note that the shape-based method alone performs remarkably well in those areas compared to the actual tracks. As we examine trajectories further away from the landmarks less shape influence from smoothing is propagated and the trajectories go awry as we might expect. Again, incorporation of phase-velocity data considerably improves our results.

When contours from an actual canine heart acquisition are analyzed we see some very interesting results. In figure 14 we examine the boundary of the endocardium. The shape-based method alone does a very plausible job of tracking the actual contours. We then provide the results from our combined shape-based, phase velocity approach in figure 15 for comparison and note the improvement in areas where little shape variation occurs. Finally, in figure 16, we compare composite graphics from tracking the phase velocity information alone using smoothed contours to those trajectories obtained from our combined method. We note that even when we use a smooth restricted set of contours, the mesh generation alone does not do any better. In fact in areas of high motion, the accumulated error becomes significant, leaving the ending trajectory far from its starting point. This is partly because the velocity map in this instance is much noisier as shown in figure 2 above. When we combine the two methods though, we see a marked improvement over both estimates. Examining the contours illustrates that in instances where strong features exist, the shape match refinement helps maintain adherence to the feature, but in areas where little shape information exists the phase velocity maps force the trajectories in the correct direction.

Integration of phase velocity information into our already successful shape-based filtering method (see figure 10) shows considerable promise. Unfortunately, at the time of this paper's deadline, phase

contrast velocity data was not available for any studies on canines with implanted LV markers thereby allowing us to completely verify our phase velocity assisted algorithm. We have plans to conduct these experiments in the January 1995 time.

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 several techniques that have not seen prior use in computer vision literature, specifically velocity maps from phase-contrast cine MRI and multi-dimensional harmonic analysis of deforming contours. Using these combined techniques we developed a method of filtering our phase velocity-driven, 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 knowledge of physical flow, feature adherence, spatial smoothness and 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 correspondences since the true trajectory extends in three spatial dimensions, however we are encouraged by these results and intend to integrate them with concurrent work of our group in matching three-dimensional surfaces.

We showed how incorporation of information from phase velocity analysis can be used to guide the shape correspondence process, especially in areas where shape variation is limited. This method of integrating different modes of MR data shows considerable potential and will be pursued extensively in the future. We have already begun research on a stochastic estimation framework that weighs the confidence of a given shape match with the quality of its corresponding phase velocity estimate.

Other 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.

Extension of the algorithm to three spatial dimensions is the ultimate goal of our research and will no doubt receive the majority of emphasis in the future.

Acknowledgment

The authors would like to thank Dr. Albert Sinusas and Don Dione of the Department of Cardiology, Yale School of Medicine, for performing the experiments used to extract the phase-contrast cine MR data.

References

[1] Dmitry Goldgof and Raj Acharya, editors. *IEEE Workshop on Biomedical Image Analysis*, Seattle, WA, June 1994. IEEE Computer Society Press.

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

[3] 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.

[4] John C. McEachen II, Arye Nehorai, and James S. Duncan. A recursive filter for temporal analysis of cardiac motion. In *IEEE Workshop on Biomedical Image Analysis*, pages 124–133, Seattle, WA, June 1994. IEEE Computer Society Press.

[5] John C. McEachen II, Arye Nehorai, and James S. Duncan. Analysis of cardiac motion with recursive comb filtering. In *Mathematical Methods in Medical Imaging III*, pages 46–57, San Diego, CA, July 1994. SPIE.

[6] R. Todd Constable, Kristina M. Rath, Albert J. Sinusas, and John C. Gore. Development and evaluation of tracking algorithms for cardiac wall motion analysis using phase velocity MR imaging. *Magnetic Resonance in Medicine*, 32:33–42, 1994.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

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

[12] 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.

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

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

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] P. van Dijk. Direct cardiac NMR imaging of heart wall and blood flow velocity. *J. Comput. Assist. Tomogr.*, 8:429–436, 1984.

[20] G. L. Nayler, D. N. Firmin, and D. B. Longmore. Blood flow imaging by cine magnetic resonance. *J. Comput. Assist. Tomogr.*, 10:715–722, 1986.

[21] N. J. Pelc. Mycardial motion analysis with phase contrast cine MRI. In *SMRM, 10th Annual Meeting*, page 17, San Francisco, CA, August 1991. SMRM.

[22] A. J. M. Spencer. *Continuum Mechanics*. Longman Group Ltd., London, 1980.

[23] George D. Meier, Marvin C. Ziskin, William P. Santamore, and Alfred A. Bove. Kinematics of the beating heart. *IEEE Trans. on Biomed. Eng.*, 27(6):319–329, June 1980.

[24] Lewis K. Waldman, Y. C. Fung, and James W. Covell. Transmural myocardial deformation in the canine left ventricle. *Circulation Research*, 57(1):152–163, July 1985.

[25] A. D. McCulloch and J. H. Omens. Non-homogeneous analysis of three-dimensional transmural finite deformation in canine ventricular myocardium. *J. Biomechanics*, 24(7):539–548, 1991.

[26] G. Strang. Piecewise polynomials and the finite element method. *Bulletin of the American Mathematical Society*, 57(1):152–163, July 1985.

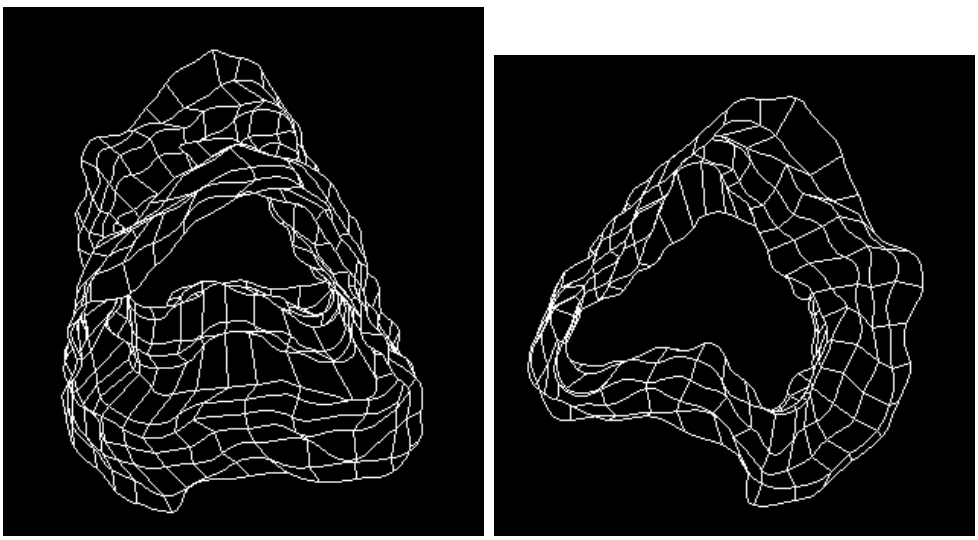

Figure 7: Two views of a complete cycle of sixteen contours segmented from MR images of a LV slice of a healthy canine heart. The contours are stacked as a surface with time as the third dimension to aid in visualizing shape trends.

[27] E. B. Becker, G. F. Carey, and J. T. Oden. *Finite Elements, An Introduction, Volume I.* Prentice Hall, 1981.

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

[29] 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.

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

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

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

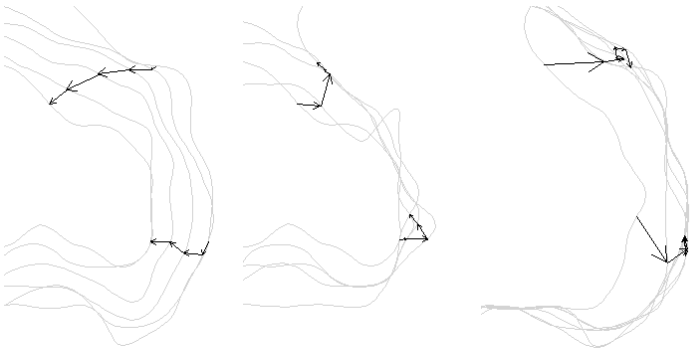

Figure 8: Trajectories of markers implanted in the LV wall. Trajectories are divided into three sections to avoid cluttering the graphic. Phase contrast velocity data was not available for this experiment, its intent was solely to campare accuracy of our RLS filter. (Left) Trajectories for frames 1-6. (Middle) Frames 6-10. (Right) Frames 10-16.

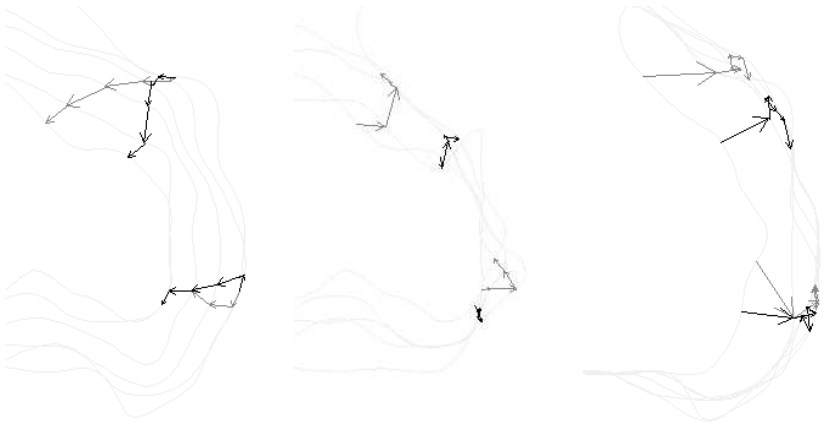

Figure 9: Corresponding trajectories estimated with a frame-to-frame shape-based tracking algorithm without any periodicity constraints (shown in black) compared with marker trajectories (shown in grey). (Top) Frames 1-6. (Middle) Frames 6-10. (Bottom) Frames 10-16.

Figure 10: 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.

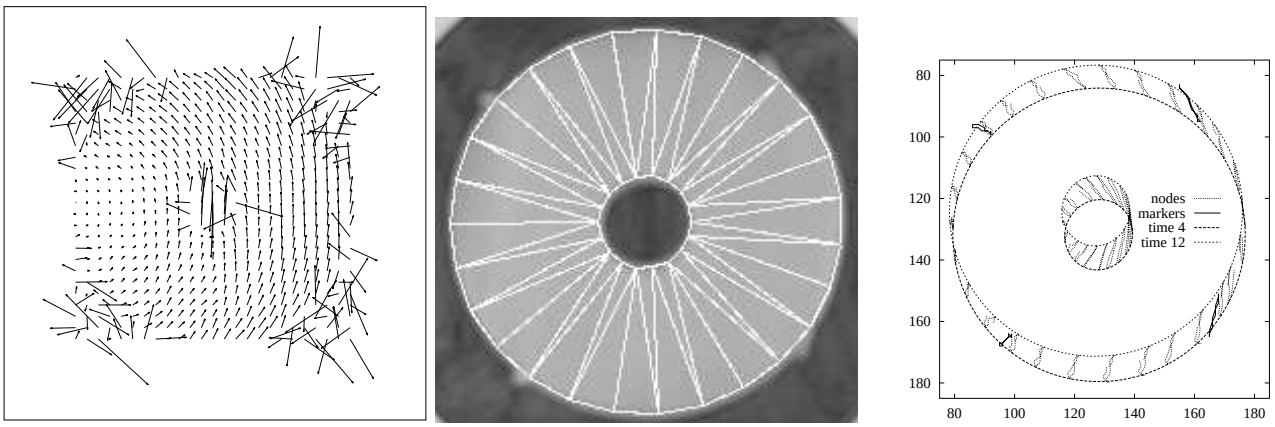

Figure 11: Analysis results of the imaged phantom in rigid-body motion. (Left) The raw velocity map of the image. Note how vectors near the landmarks become very erratic. (Center) A mesh fit to the image that filters out the noisy vectors associated with the landmarks. (Right) The resulting estimated motion. The true motion of the landmarks is overlaid in solid black.

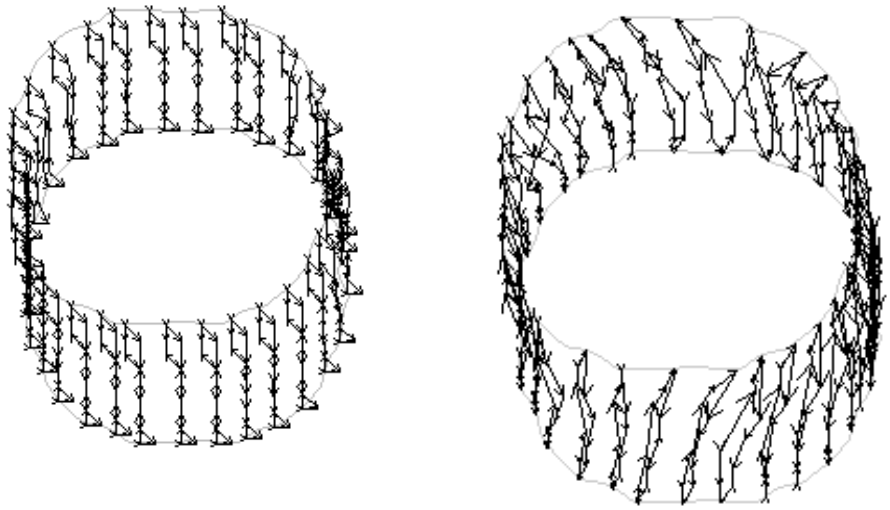

Figure 12: Results of applying the RLS filter to the inner contour of the phantom. (Left) The shape-based filtered estimate without any phase velocity information. Note how the absence of shape information leads the algorithm to find the nearest point in euclidean distance in most instances. (Right) The filtered estimate of shape-based tracking guided by phase velocity information.

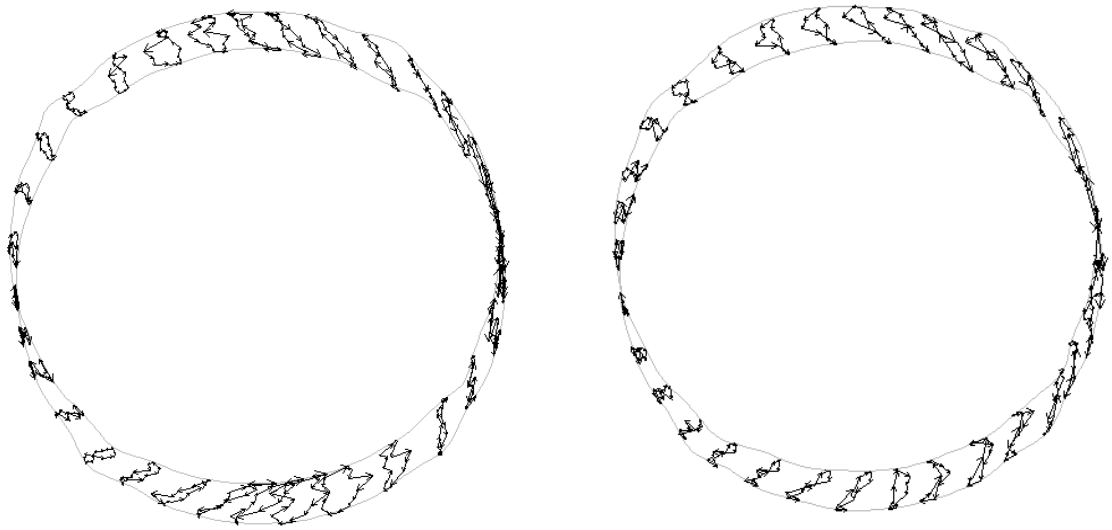

Figure 13: Results of applying the RLS filter to the outer contour of the phantom. (Left) The shape-based filtered estimate without phase velocity information. Note that in areas near the landmarks, the algorithm does a suitable job, however, trajectories get progressively worse as they move away from the smoothing influence of the landmarks. (Right) The filtered estimate of the combined approach.

Figure 14: Tracking the endocardial wall (short axial view) in another canine LV study using our shape-based RLS filter with no phase contrast velocity information. (Left) Frames 1-6. (Middle) Frames 6-11. (Right) Frames 11-16.

Figure 15: Tracking the same study as in figure 14 but this time incorporating the phase velocity information shown in 2 into the filtering process. (Left) Frames 1-6. (Middle) Frames 6-11. (Right) Frames 11-16.

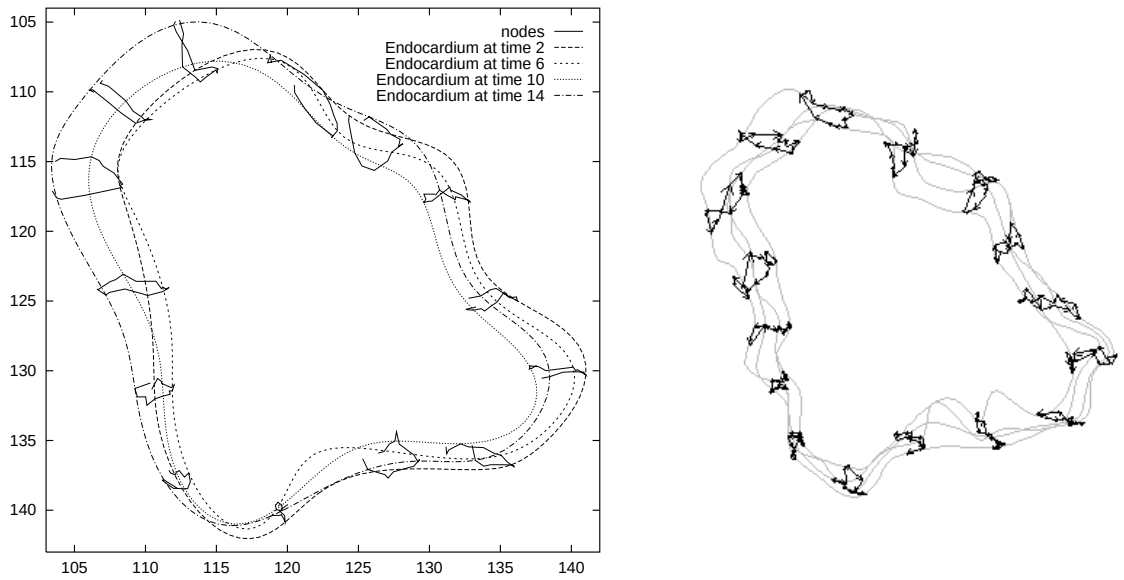

Figure 16: (Left) Results from tracking the same canine study as in 15 using strictly phase velocity information from the deforming mesh. Note for this example, very smooth versions of the contours are used to avoid integrating outliers into the solution. (Right) Composite graphic of results from 15 combining phase velocity and shape information using actual contours. In both graphics only contours from frames 2, 6, 10 and 14 of the 16 frame sequence are shown to avoid cluttering the graphic.