

Point–Tracked Quantitative Analysis
of Left Ventricular Motion from 3D Image Sequences

Pengcheng Shi^{*,†}, Albert J. Sinusas^{†,‡},
R. Todd Constable[†], Erik Ritman^{*}, and James S. Duncan^{*,†}

Departments of *Electrical Engineering,
†Diagnostic Radiology, and ‡Medicine
Yale University
New Haven, CT 06520-8042

Department of Physiology and Biophysics^{*}
Mayo Clinic
Rochester, MN 55905

This work was supported by NIH–NHLBI Grant No. R01HL44803.

Abstract

We propose and validate the hypothesis that we can use differential shape properties of the myocardial surfaces to recover dense field motion from standard three-dimensional image data (MRI and CT). Quantitative measures of left ventricular regional function can be further inferred from the point correspondence maps. The non-invasive, algorithm-derived results are validated in two levels: the motion trajectories are compared to those of implanted imaging-opaque markers of a canine model in two imaging modalities, where sub-pixel accuracy is achieved, and the validity of using motion parameters (path length and thickness changes) in detecting myocardial injury is tested by comparison to post-mortem TTC staining of myocardial tissue, where the Pearson product-moment correlation value is 0.968.

Keywords: *Motion Analysis, Cardiac Function, Shape Modeling, Algorithm Validation*

1 Introduction

It has been the fundamental goal of many efforts in cardiac imaging and image analysis community to assess the regional function of the left ventricle (LV) of the heart. The general consensus is that the analysis of heart wall dynamics (motion, thickening, strain, etc) provides quantitative estimates of the location and extent of ischemic myocardial injury. It is also felt that image-based serial analysis of regional function is helpful in assessing the efficacy of therapeutic agents or mechanical interventions (*i.e.* angioplasty) directed at reducing myocardial injury [14, 22]. However, imaging-based clinical analysis of regional LV function, usually from SPECT and Echocardiography, often results in overestimation of true myocardial infarction. More sophisticated quantitative analysis of myocardial dynamics may improve the accuracy in the estimation of myocardial injury, and add to the understanding of the remodeling process following myocardial infarction as it relates to the extent of myocardial injury.

There have been considerable efforts within the medical image analysis and computer vision communities aimed at establishing sparse correspondences between different time instants of a moving object, as well as mapping the entire object from one time instant to next, both with applications in cardiac motion analysis. In addition, a significant level of activity has been performed within the magnetic resonance (MR) imaging community regarding the measurement of LV motion and deformation using MR tagging, and to a lesser extent, MR phase contrast velocity images. However, because of the complexity of the heart motion, the lack of unambiguous reference landmarks, and the limitations of the existing imaging techniques (*i.e.* out of plane motion problem for two-dimensional (2D) image slices), accurate assessment of LV motion remains a difficult problem.

In this paper, we propose and validate the hypothesis that we can use differential shape properties of the myocardial surfaces to track dense field three-dimensional (3D) trajectory of the surface points over the entire cardiac cycle. The matching and tracking criteria are based upon locating and matching differential geometric surface features, as well as a mathematical optimization reasoning strategy to combine the data-driven matching with a local coherent smoothness model. While it may not be true for non-rigid motion analysis in general, we show that geometrical token tracking to be a viable strategy in myocardial surface motion analysis, as suggested by an earlier visual study [55].

Very few of previous efforts in cardiac image analysis attempted to validate the accuracy of their results against recognized gold standards. We believe that carefully designed ground truth

is essential in validating the results of any image analysis based approach. Towards this goal, we have developed an experimental setup using a canine model where imaging-opaque markers are implanted on the surfaces of the left ventricular wall, and the motion of these markers are used as the gold standard against which the validity of our non-invasive motion analysis algorithms are evaluated. The results of this validation process in two image modalities (standard cine MRI and computer tomography (CT)) are presented. Also, the usefulness of using image-derived motion measures in predicting regional myocardial viability is also validated using post-mortem TTC staining of myocardial tissue, which is the current clinical standard.

We believe that our proposed method overcomes many of the problems associated with the 2D image-based approaches, and complements some of the existing 3D efforts, such as MRI tagging and MR phase contrast imaging. It is non-invasive in nature and can follow a dense field of points. It offers the high accuracy of the gold standard implanted-markers, with sub-pixel average errors for both imaging modalities. And it accurately detects the location and extent of myocardial infarction, which is evident from its high correlation with the TTC staining. Also very importantly, this approach does not rely on special imaging methods such as MR tagging or phase contrast, but rather is applicable to several standard imaging protocols (cine MRI and CT, and we are experimenting with ultrasound images). Our efforts in integrating the shape-based approach with MR tagging and phase contrast information are reported in other papers [53, 51, 34, 32].

2 Previous Efforts

2.1 Medical Image Analysis

Earlier image-based approaches to measuring cardiac motion largely fall into two groups. The first one attempts to measure wall motion by detecting gray level changes over time for each pixel, the optical flow strategy [31, 56, 60]. Besides the issue that optical flow is usually not the same as the motion field of the physical material points [26], it is quite difficult to post-process the results to separate the wall motion from other motion in the images.

The second group attempts to track the motion of pre-extracted boundaries of the cardiac wall, and recent works often use geometric measures as tracking tokens. It includes the earlier works such as the chord shortening method [20], and the centerline method [8], which treats the end-diastole (ED) and end-systole (ES) contour difference as a thickness measure but does not track each contour point through the cardiac cycle. The important work of Slager [55] established

the experimental foundation for salient feature tracking. He used visually extracted contour shape landmarks to track point motion over time from 2D image sequences, and validated the results using implanted physical markers. The correlation between the two motion trajectories was good ($r = 0.86$), which seemed to support the concept of tracking identifiable shape cues over time in cardiac images. More recently, several boundary differential shape matching ideas for LV motion tracking are reported, using Gaussian curvature under conformal stretching surface models [30], hybrid bending and stretching surface model [2], and 2D elastic rod bending model [32, 40].

One key issue in cardiac motion arises from the fact that the heart is a non-rigid moving object that rotates, translates, and deforms in 3D space. Hence, motion measurements from 2D image sequences most likely are not derived from the same points of the heart wall at different time instants, thus do not measure the true movement of the heart. Only 3D based analysis has the potential to accurately describe the true dynamics of the heart. We also want to point out, without much elaboration, that 3D approaches are technically very differently from corresponding 2D approaches, even though conceptually they could be very similar. Extension of 2D approaches to 3D are neither technically trivial nor always possible, because of the added topological complexity of the third dimension.

2.2 Imaging Physics-Based Methods

2.2.1 MR Tagging

In magnetic resonance tagging, a spatially varying pattern of magnetization is encoded into the myocardium prior to each phase-encoded data acquisition cycle of a spin-echo imaging sequence, forming a grid pattern of image voids generated at each temporal instant in the cardiac cycle [5, 66]. Motion occurring between the tag pulse sequence and the image pulse sequence results in distortions of the tag pattern, and the grid deformation can be tracked over a portion of the cardiac cycle. Sparse sets of grid motion measurements of several orthogonal tagging grid acquisitions, where each acquisition provides the 2D motion components of tags in that particular imaging plane, are assembled. 3D motion descriptions are then derived from the tracking of the sparse set of tag-tag or tag-boundary crossings.

MR tagging certainly shows promises but still has some major technical limitations. It remains a fundamental problem to track the tags over the complete cardiac cycle, due to decay of the tags over time. While it is often remedied by performing a second acquisition (a re-tagging)

somewhere later in the cardiac cycle, the same tissue elements are generally not tagged in each of the two tagging processes. Hence, the motion trajectory of the tag crossing point in each tagging sequence is not acquired for the same tissue or point. The further derived motion fields from two tagging processes may be incoherent if the myocardium is heterogeneous and the deformation is finite.

Also, only sparse in-plane motion, the approximated 2D projection of the 3D motion of the material points being imaged, can be obtained from tagging images directly [41]. In order to obtain data pertaining to motion in three dimensions, 2D tag data must be acquired in two orthogonal views, typically short axis and long axis [6]. Once again, the same tissue elements are not tagged in each of the two views, and thus the motion in each view must be seen as partial 2D data at different points that contributes to an overall scheme aimed at estimating the complete 3D motion and deformation.

There are very active researches in the MR tag image analysis. Some of the most interesting ideas include the use of finite element models with spring-like interconnections and nonlinear basis functions [65], the use of locally deformable superquadrics Park94,Park95, the use of stochastic models and statistical estimation techniques [18], the high order polynomial fitting of sparse displacements [39], and the use of B-snake grid [1].

While there is no doubt that MR tagging can provide unique and interesting data regarding LV myocardial movement, there are quite a few processing steps required to assemble the data into meaningful measures of 3D motion. Just having the MR tag data alone does not necessarily mean that physiologically and clinically useful analysis is forthcoming. The proper choice of image analysis and processing algorithms is essential and remains a significant open question.

2.2.2 MR Phase Contrast

Another imaging-physics based technique for cardiac motion analysis employs MR phase contrast, which relies on the fact that a uniform motion of tissue in the presence of a magnetic field gradient produces a change in the MR signal phase that is proportional to the velocity of the tissue [46]. The velocity in any particular spatial direction can be estimated by measuring the difference in phase shift between two acquisitions with different first gradient moments. Hence, velocity maps encoded for motion in all three spatial dimensions may be obtained at multiple time instances throughout the cardiac cycle using a phase contrast cine-MR imaging sequence, which may then be used to provide information on tissue displacement, strain, strain rate, and other quantitative measures of

motion and deformation [15, 38, 46, 43, 45, 44]. In principle, the instantaneous velocity can be derived for each pixel in an image acquisition. However, because of the size of the acquisition region of interest (ROI) for each pixel, accurate velocity information is only available at the middle of the myocardial wall, but not at the wall boundaries.

Since the velocity maps themselves only provide instantaneous motion information, the important first step in using these maps to estimate Lagrange motion parameters is to devise methods that can accurately track each segment of myocardium as it deforms through the heart cycle. Methods have been proposed to use forward integration, backward integration, or combinations of the two, of the velocity information to estimate the displacement vectors [15, 23, 44, 64]. Meyer’s approach combines a spatially regularizing velocity field with temporal Kalman filtering, and produces more robust results in 2D [34]. All these methods assume some kind of constant velocity conditions between time frames, which generally is not true even for small time intervals. Errors resulting from the constant velocity assumptions can be significant if the velocity change rate is large, and small error accumulates as velocity is integrated through the cardiac cycle.

2.3 Computer Vision Based Motion Analysis

Motion measurement from image sequences has long been studied by many computer vision researchers. The primary emphasis has been on the determination of optical flow [4, 25, 27, 36], and the determination of point correspondences between successive frames of rigid objects [7, 21]. They are generally not very successful in estimating the true motion field of deformable objects because non-rigid motion makes it more difficult to enforce appropriate smoothness constraints on the optical flow field which are necessary to deal with the aperture problem.

In computer vision, quantifying the boundary motion of non-rigid objects has often been seen as a two step process: establishing correspondence between certain sub-sampled points of the boundary between time t and time $t+1$, then using these correspondences as a guide to solve for a complete mapping of the object boundary between two time frames. While early approaches used simple global distance measures to find correspondence, more recent matching methods have been mainly based on tracking salient image features over time, from simple tokens such as points or line segments to complex features such as eigenmodes of finite element representations [2, 30, 32, 35, 40, 50, 52].

The task of establishing a complete non-linear mapping between object frames has received more attention. The basic assumption is that estimates of correspondence between sparse indi-

vidual points on objects are either specifically assumed to be known or can be established based on some global distance measures. Physically-based finite element models are used to provide a framework for the mappings [59, 29], and modal analysis is adopted to reduce computational cost [28, 37]. Bookstein used deformable thin-plate splines to interpolate dense correspondences from very sparse initial matches [9]. Locally deformable superquadrics approaches are also used for non-rigid problems [33, 58], including several using MR tagging data to establish initial correspondence [42, 41].

3 Surface Characterization

Since the proposed algorithm is based upon tracking surface shape tokens over time, the endocardial and epicardial surfaces of the LV, as well as their differential properties, over the cardiac cycle are derived from the 4D image data at the start of the process.

3.1 LV Wall Boundary Finding

The boundary finding problem is solved twice for each 2D image slice of a 3D image frame, once for the endocardial border and once for the epicardial one. To specifically address robustness of the segmentation results, we use an integrated approach which combines gradient based deformable boundary finding with region based segmentation [12]. This approach uses Green’s theorem to derive the boundary of a homogeneous region-classified area in the image, and integrates this with a gray level gradient based boundary finder. It combines the perceptual notion of edge information with gray level homogeneity, and is more robust to noise and poor initialization.

For either endocardial or epicardial boundary segmentation, a contour is manually initialized within a 2D image frame representing a middle-ventricle slice of a 3D image, halfway between ED and ES. The above described boundary finding algorithm is run, and the final result is used to initialize similar 2D boundary finding problems in image frames spatially and temporally adjacent to the starting image frame. The results are propagated and the process is repeated until all of the contours that make up the LV surfaces in all 3D frames are detected. During the process, human operator decides the top and bottom slices to be segmented beforehand, and may re-initialize the algorithm if needed. The reproducibility of the algorithm under different initialization and noise conditions has been reported [13]. Figure 1 shows the segmentation results of one of the 3D MR image frame.

Some of the 3D image data, such as MR, may have higher in-plane resolution than inter-plane resolution because of the limitations in the imaging protocol. However, it is often more desirable to have roughly equally sampled representations in all three dimensions to achieve more accurate quantitative analysis and display [61]. Since we are only interested in the boundary points of the LV, a shape-based contour interpolation method [24], using the chamfer distance transformation [10, 11], is implemented. Here, interpolated contours are computed from linear combination of the distance maps of the two spatially adjacent data contours. After the segmentation and the contour-interpolation process, endocardial and epicardial boundaries are represented by contour stacks at each time instant.

3.2 Delaunay Triangulation

We want to tessellate each contour stack to infer a surface representation of the LV boundary, which will then allow the calculation of the geometrical parameters of the surface that form the basis of the motion analysis. Delaunay triangulation is chosen as the surface representation because it is most suitable for volumetric representation of a set of spatial sample points without *a priori* connectivity information. In the case of *from contour stack to surface* that we are facing, even though we do know the explicit relationship for boundary points within any given slice which will be enforced, we do not have any prior knowledge of the inter-slice spatial relationship. The two-dimensional Delaunay facets which are on the *boundary* of a bounded and constrained Delaunay triangulation constitute the corresponding LV surface [54], which defines a symmetrical and isotropic neighborhood point relationship.

For a set \mathbf{V} of N points in 3D, Delaunay triangulation performs a simplex decomposition of the convex hull of the point set, where the vertices of the tetrahedra are contained in the point set [62]. It is a globally optimal triangulation in the sense that the circumsphere of every tetrahedron in the triangulation contains no other point from the point set, and every tetrahedron is as *regular* as possible. With the exception of the degenerate case, the Delaunay triangulation is a unique minimal shape representation for a given set of points, and can be computed efficiently.

A mortified incremental point insertion algorithm is implemented [62]. For LV surface representation, the distribution of the points is limited to certain subspace of \mathfrak{R}^3 , and a subset of the points has *a priori* connectivity relationship. A bounded and constrained Delaunay triangulation is developed [54], and a volumetric tessellation of both endocardial and epicardial points is constructed. The Delaunay facets which are on the *outside boundary* of the triangulation constitute the

epicardial surface, while the facets which are on the *inside boundary* of the triangulation constitute the endocardial surface. The surfaces are left open at the top and bottom slices. Figure 2a shows a Delaunay tessellated MR endocardial surface.

The circumspheres of the Delaunay triangles express a special adjacency relationship among spatial points [63]. For any point Q , its natural neighbor points are those points which share the same Delaunay edges with Q . For example, points F_i in Figure 2b are the natural neighbor points of Q . We can even further define multiple orders of natural neighbor relationships of spatial points from the Delaunay triangulation, hence measure the natural *closeness* of arbitrary points [54]. Even though Euclidean distance measure performs a similar function, the natural neighbor relationship takes into account the geometrical relations between points, thus provides a more accurate description of the relationship. In the remaining of the paper, we use this natural neighbor relationship extensively for defining surface patches and search windows.

3.3 Surface Curvature Estimation

Since our motion tracking algorithm is shape based, it is thus closely related to the characteristics, (and implicitly their reliability, reproducibility, and accuracy), of the LV surface shape and their variations along the entire cardiac cycle. We have developed a method which results in intrinsic, locally smooth curvature maps.

3.3.1 Triangulated Surface Smoothing

As described previously, the myocardial surfaces are represented by the aggregate of the two-dimensional Delaunay facets. These polyhedral surfaces, however, often appear faceted. One reason is that each 2D boundary contour from the segmentation is digitized at the resolution of the image sampling lattice which may not be fine enough for smooth reconstruction of continuous contour. A more fundamental reason arises from the fact that contours of different planes are segmented independently, which means that there is no spatial constraints between contour points of different planes during segmentation, even though Delaunay triangulation does establish an *a posteriori* connectivity relationship

For shape characterization and visualization purposes, the jagged boundary surfaces are desired to be smoothed to achieve locally coherent surface representations. Most of the existing smoothing methods suffer from the *shrinkage* problem, however, which means that when the smoothing process is applied iteratively a large number of times, a surface will collapse into a point

[57]. A two-stage Gaussian smoothing algorithm has been proposed to produce non-shrinking result [57]:

$$\mathbf{x}_{new1} = (1 - \lambda_1)\mathbf{x}_{current} + \lambda_1 \sum_{i=0}^{N-1} \omega_i \mathbf{x}_i \quad (1)$$

$$\mathbf{x}_{new2} = (1 - \lambda_2)\mathbf{x}_{new1} + \lambda_2 \sum_{i=0}^{N-1} \omega_i \mathbf{x}_i \quad (2)$$

Here, $\mathbf{x}_{current}$ is a point in the Delaunay triangulation, $\{\mathbf{x}_i\}$ is its natural neighbor point set, $\{\omega_i\}$ is the weighting coefficient set associated with the neighbor point set, \mathbf{x}_{new1} is the new location of $\mathbf{x}_{current}$ after one iteration of the first stage smoothing, and \mathbf{x}_{new2} is the new location of $\mathbf{x}_{current}$ after one iteration of the second stage smoothing. For each iterative smoothing step, after a first Gaussian smoothing stage with a positive smoothing factor λ_1 which produces a slightly shrunk surface, a second Gaussian smoothing stage is applied with a negative smoothing factor λ_2 which produces a slightly expanded surface. λ_2 should have slightly greater magnitude than λ_1 ($0 < \lambda_1 < -\lambda_2$), with the difference determined by a bounded band which makes sure that the surface does not shrink or expand after each two-stage smoothing. This method is in fact a linear low-pass filter that removes high curvature variations (usually caused by noise), with the two smoothing factors determining the passing band.

Figure 3 shows the effect of the two-stage non-shrinkage smoothing on an endocardial surface. On the left is the original triangulated surface (rendered), and on the right is the smoothed surface of the two-stage non-shrinkage Gaussian smoothing process ($\lambda_1 = 0.33$, $\lambda_2 = -0.34$, iteration number $n = 10$), where the z -coordinates of the points in top and bottom contours are fixed. It is quite obvious, at least visually, that the smoothing process keeps the overall size and shape-feature of the surface. We have applied this non-shrinkage smoothing on all the myocardial surfaces to achieve coherent surface representations.

3.3.2 Local Surface Fitting and Curvature Estimation

To analyze the local shape properties of a point p in a neighborhood of a triangulated surface, polynomials are used to approximate this neighborhood, in the form of the graph of a differentiable function.

We begin the process by defining a local surface patch S at point p as a discrete surface

triangulation consisting of point p itself, *the fiducial point*, and a set of neighboring points of p . Two questions arise for constructing a smooth local surface for these points:

1. How do we define the neighborhood of the fiducial point p ?
2. How do we define the parameterization of the surface patch S ?

The second question is easier to answer: this surface patch is represented by the graph of a differentiable function. A point-dependent local coordinate system is chosen, with its origin at the fiducial point p , and its z axis in the normal direction of the surface patch S at p . The normal at p is determined as the weighted average of the outward normal vectors of all the Delaunay facets which share p as one of their Delaunay vertices. The weighting coefficient for each facet is its normalized area. Since the xy plane agrees with the tangent plane $T_p(S)$ of S at p , any pair of orthogonal directions in $T_p(S)$ can be chosen as the x and y axes. After this local coordinate system is established, surface patch S can be represented in the graph form $z = h(x, y)$, which has a one-to-one $\mathbb{R}^2 \rightarrow \mathbb{R}$ mapping if the neighborhood of p is well chosen.

Traditionally, the neighborhood of point p is defined as a sphere centered at p with a radius R in \mathbb{R}^3 , and all the points that fall inside this sphere constitute the neighboring point set [3, 49]. However, depending on the value of R and the roughness and folding of the local surface, $z = h(x, y)$ may not be a unique function. Figure 4a shows the possible multiple mapping problem in a 2D representation. In this example, the surface patch defined from the neighborhood folds around p , and causes the graph function to have multiple mappings in \mathbb{R}^2 for one value in \mathbb{R} . And since this nearest Euclidean distance based method does not take account of the known connectivity from triangulation, it may also include points which are disconnected from the valid surface points and should not be in the neighborhood of p at all, such as the case shown in Figure 4b. Of course, these two problems may be avoided by choosing a smaller radius R . However, the determination of the value for R is a difficult task without much of the knowledge about the geometry of the underlying surface, which is exactly what we are trying to estimate. An adaptive window for curvature estimation, controlled by the error of the least square fit, is another possible solution [17].

In order to choose a neighborhood of p which is unique (no folding) and local (no non-neighboring points), the natural neighbor relationship between surface points is used to define a more precise and flexible neighborhood. First of all, if we only include the 1-order natural neighbors of p in the neighborhood, the graph function will always have a unique mapping. If the surface

patch is well-behaved, which means that the orientation of one facet does not change dramatically from those of its adjacent facets, higher orders of natural neighbor points can also be included in the neighborhood of p . Depending on how many orders of natural neighbor points are included in the neighborhood, the local surface can be considered to have *multi-scale* details. Secondly, the natural connectivity ensures that only valid neighbor points will be chosen, since each of these points is connected to at least one other chosen point. This is true even if higher order natural neighbor points are included in the neighborhood.

Since the Weingarten mapping matrix only depends on the first and second derivatives of the graph function $h(x, y)$, biquadratic functions are all we need to estimate the curvatures of the surface patch. The estimate of the biquadratic graph function involves finding the five coefficients of the polynomial [3, 49]:

$$z = h(x, y) = a_1x^2 + a_2xy + a_3y^2 + a_4x + a_5y \quad (3)$$

If there are at least four neighboring points, say $n-1$ ($n \geq 5$) points, the least-square estimate of the five coefficients is evident. Since natural neighbor points are used for p , three neighboring points are the worst case (two 1-order neighbor in the same slice, one 1-order neighbor in the adjacent slice). We can always include higher order natural neighbors to have at least four neighbors.

In matrix form, the coefficient vector \vec{a} which minimizes the fit error is:

$$\vec{a} = (A^T A)^{-1} A^T \vec{Z} \quad (4)$$

where A is a $n \times 5$ matrix with row i being $[x_i^2 x_i y_i y_i^2 x_i y_i]$, \vec{Z} is a column vector of the z_i values. Once the coefficient vectors are found, computing the actual differential properties on the fiducial point is straightforward [19].

The graph function and curvature estimation may follow a multi-scale procedure, depending on how many orders of natural neighbor points are included in the neighborhood, to examine the consistence of the shape information across different extent of the chosen neighborhood. Let's denote the k -scale graph function as the one estimated from the neighborhood which includes 0-order, ..., k -order natural neighbor points. If certain prior information is available, we can use lower scale (small k) graph functions to characterize highly curved surface patches to avoid the

smoothing effect, and use higher scale (large k) graph functions to characterize flat surface patches to alleviate the noise effect. The consistency of the estimated shape characteristics through different scales, such as from higher to lower, may also indicate the goodness and reliability of predominant shape features (corners, ridge lines, etc.) in the surface. From our experiments, it is clear that the significant shape patterns are generally consistent across different scales. The curvature maps from lower scale graph functions (small neighborhood) may appear noisier than curvature maps from higher scale graph functions (large neighborhood), while the latter ones may over-smooth the surface shape features. Figure 5 shows the Gaussian curvature, mean curvature, and the two principal curvature maps of an endocardial surface derived from 1-scale, 2-scale and 3-scale graph functions. The motion tracking examples in the remaining of this paper are achieved only using 2-scale estimated curvatures.

4 Non-Rigid Motion Tracking

The movement of a set of LV surface points over cardiac cycle is determined by following local surface shape between successive temporal frames. In the first step of the tracking algorithm, endocardial and epicardial surfaces at one given time instant are used as reference surfaces and are sub-sampled, creating a sparser set of points (from five to ten percent of the total points). Significant shape cues, such as cup and cap points, are always chosen. Other points are selected with random spatial distribution. The reason to use this sparse set of points for initial match is to avoid cross-over between nearby matches. Under the assumption that the LV surface deforms as little as possible between small time intervals, the sub-sampled surface points are matched and followed over the LV surface sequences under a minimum bending energy model. These initially-calculated 3D displacement vectors can be noisy due to inherent data noise and matching uncertainty. Match-confidence-weighted regularization process is utilized to help limiting these problems and to compute the dense field motion for all the surface points.

4.1 Thin Plate Bending Model

Local surface patches on the LV are modeled as thin flexible plates, loosely constrained to deform in a predetermined fashion which will be defined in the next section. The potential energy of an ideal thin flexible plate of elastic material is a measure of the strain energy of the deformation from its equilibrium state, a plane. It is a function of the two principal curvatures of the plate, and is

invariant to 3D rotation and translation [16]:

$$\epsilon = A \left(\frac{\kappa_1^2 + \kappa_2^2}{2} \right) \quad (5)$$

where A is a material related constant for the particular plate, κ_1 and κ_2 are the two principal curvatures of the surface patch of interest. Note that even though the principal curvatures have associated directions, only their magnitudes are relevant in the potential energy metric. The potential energy maps of an endocardial surface over the entire cardiac cycle (sixteen temporal sampling frames) are shown in Figure 6. In these maps, most of the high potential energy features run through several temporal frames or even the entire sequence. This consistency of shape features over time provides the visual basis for matching surface points based upon geometrical tokens.

We modify the idea of surface potential energy slightly to define the bending energy required to bend a curved plate or surface patch to a newly deformed state as:

$$\epsilon_{be} = A \left(\frac{(\kappa_1 - \bar{\kappa}_1)^2 + (\kappa_2 - \bar{\kappa}_2)^2}{2} \right) \quad (6)$$

The principal curvatures of the initial surface patch are κ_1 and κ_2 , while the same parameters of the deformed surface patch are $\bar{\kappa}_1$ and $\bar{\kappa}_2$. This equation assumes that the corresponding points on the surface patches are known, and arrives at a numerical value measuring the energy required to deform the initial surface patch to achieve a deformed shape. To use the square of curvature differences instead of the differences of the curvature squares is because the principal curvatures are signed quantities. Obviously, if the surface patch undergoes rigid motion, there is no bending energy required for the movement since the shape of the patch does not change. Only non-rigid deformation requires bending energy which measures a form of strain energy.

4.2 Shape-Based Surface Point Match

The bending energy of Equation 7 measures the required energy to deform a surface patch from one state to another. However, it is the goal of the motion tracking algorithm to find the point correspondences between temporally successive surfaces in the sequence. We use the bending energy as the matching criterion to obtain the point correspondences, where the essence of the algorithm is the hypothesis that a surface patch deforms as little as possible between successive temporal

frames, and the deformation can be fully characterized by the bending energy.

The motion of the surface patch surrounding each sample point consists of two distinctive parts: the rigid motion (global and local) and the non-rigid deformation. Since rigid motion does not change the shape properties of a surface patch, point correspondences with no deformation can be found by matching surface patches requiring no bending energy. Hence, the minimum bending energy criteria can be used for rigid motion. For the non-rigid deformation, we make the assumption that the surface patch is constrained to deform in a fashion such that the extent of the deformation is determined solely by the required bending energy. This assumption makes it reasonable to use the bending energy measure as the matching criteria for the non-rigid deformation.

Under the strong assumption that the surface patch surrounding each sample point deforms only slightly and locally within a small time interval, we construct a physically-plausible search region on the second surface at time t_{i+1} for each sub-sampled point on the first surface at time t_i . For any point at time t_i which is a cup, trough, rut, ridge, dome, or cap point, a search region W is defined as the area surrounding the intersection of normal vector of the point $\mathbf{x}(\mathbf{u})$ at t_i with the surface at t_{i+1} . For any point at time t_i which is a saddle rut, saddle, or saddle ridge point, a search region W is defined as the area surrounding the nearest surface point at t_{i+1} to the point $\mathbf{x}(\mathbf{u})$ at t_i . W consists of the natural neighbor points of the intersected point, often of the same or slightly higher order as the surface patch fitting process. These points inside W is denoted as point set $\{\mathbf{x}_{W,i}, i = 0, \dots, N - 1\}$. Bending energy measures between the surface patch surrounding point $\mathbf{x}(\mathbf{u})$ and surface patches surrounding points $\mathbf{x}_{W,i}$ are computed, and the point $\bar{\mathbf{x}}$ within the search window W that has the *minimum corresponding bending energy* is chosen as the point corresponding to point \mathbf{x} :

$$\begin{aligned} \bar{\mathbf{x}} &= \arg \min_{\mathbf{x}_{W,i}} \epsilon_{be}(\mathbf{x}, \mathbf{x}_{W,i}) \\ &= \arg \min_{\mathbf{x}_{W,i}} A \left(\frac{(\kappa_1(\mathbf{x}) - \kappa_1(\mathbf{x}_{W,i}))^2 + (\kappa_2(\mathbf{x}) - \kappa_2(\mathbf{x}_{W,i}))^2}{2} \right) \end{aligned} \quad (7)$$

Here, we assume the material constant A is the same for all surface patches surrounding points $\mathbf{x}_{W,i}$ of the search window, and it does not change from t_i to t_{i+1} .

Performing the matching process for all the sub-sampled points of the surface at t_i , the results yield a set of sparse, shape-based, best-matched motion vectors $\mathbf{d}_{init}(\mathbf{u}, t_i) = \bar{\mathbf{x}} - \mathbf{x}$ for pairs of surfaces derived from 3D surface sequence. An example of bending energy-based matching

of an endocardial surface point is shown in Figure 7. The complete trajectory of the point over the cardiac cycle (sixteen temporal frames) superimposed onto the endocardial surface at ES is shown in Figure 7a as an overview, and in Figure 7b as a more detailed blowup. Evidently, the point moves upwards during the contraction phase (ED to ES) and it moves downwards during the expansion phase (ES to ED). This out-of-plane motion demonstrates one important reason to use 3D images for LV motion analysis. It should be noted that the trajectory of this point is almost a complete loop, meaning the point moves back to where it begins. Since there are no periodic motion constraints imposed on the point correspondence tracking over the cardiac cycle, this result cannot be guaranteed. However, we have begun to work on the temporal and periodic motion modeling for two-dimensional image sequences [32]. For our current results in 3D, surface points usually move back near where they begin but not to the exact beginning positions in most cases, with average differences in the range of 1 to 1.5 pixel.

The bending energy measures for all the points inside each search region are recorded as the basis to measure the *goodness* and *uniqueness* of the matching choices. The value of the minimum bending energy in the search region between the matched points indicates the goodness of the match. This is because the matching criteria is based on the notion that the smaller the deformation between matched surface patches, the better the match. Denote this value as m_g , we have the following measure for matching goodness:

$$m_g(\mathbf{x}) = \epsilon_{be}(\mathbf{x}, \bar{\mathbf{x}}) \quad (8)$$

On the other hand, it is desirable that the chosen matching point is a unique choice among the candidate points within the search window. Ideally, the bending energy value of the chosen point should be an outlier (much smaller value) compared to the values of the rest of the points (much larger values). If we denote the mean values of the bending energy measures of all the points inside window W except the chosen point as $\bar{\epsilon}_{be}$ and the standard deviation as σ_{be} , we define the uniqueness measure as:

$$m_u(\mathbf{x}) = \frac{\epsilon_{be}(\mathbf{x}, \bar{\mathbf{x}})}{\bar{\epsilon}_{be} - \sigma_{be}} \quad (9)$$

This metric means that if the bending energy of the chosen point is small compared to some smaller

value (mean minus standard deviation) of the remaining bending energy measures, it is quite unique and reliable, otherwise it is not. Obviously for both goodness and unique measures, the smaller the values are, the more reliable the match. Combining these two measures together, we arrive at one *confidence measure* for the matched point $\bar{\mathbf{x}}$ of point \mathbf{x} :

$$c(\mathbf{x}) = \frac{1}{k_{1,g} + k_{2,g}m_g(\mathbf{x})} \frac{1}{k_{1,u} + k_{2,u}m_u(\mathbf{x})} \quad (10)$$

where $k_{1,g}$, $k_{2,g}$, $k_{1,u}$, and $k_{2,u}$ are scaling constants for normalization purposes. The confidence measures for all the surface matches are normalized to the range of 0 to 1. The matches with the highest confidence have the confidence measure value of 1, and the matches with the lowest confidence have the confidence measure value of 0. In the shape-based work on 2D contour matching [32, 40], a similar goodness measure is used. However, the uniqueness measure in that case is based upon the assumption that the matching metric (the bending energy) is spatially distributed as Gaussian, centered at the matched point. We do not believe this is the situation in the 3D case.

4.3 Optimized Dense Motion Field

A smoothing-interpolation process is needed for accurate estimation of the dense motion field of the surfaces. One reason is that the initial surface matching results are obtained by purely local processes, without any constraint on the motion of neighboring points, even though surface motion of biological origin often varies smoothly in direction and magnitude over space. Another reason is that some sort of interpolation is needed to compute the dense motion field for every point on the surface, while temporal correspondences of only a set of sub-sampled points are found. Under the assumption that the displacement fields of the LV surfaces are sufficiently smooth and that any changes of direction and magnitude take place in a gradual manner over space, a regularizing smoothing functional which includes both an initial estimate adherence term and a local smoothness term is constructed. Since the reliability of the initial correspondences vary based on their confidence measures, it should be clear that good, unique matches from the initial bending energy matching process are preserved, while ambiguous matches are smoothed over by their neighboring good and unique matches. The confidence measures from the initial match are used here as smoothing coefficients.

In the continuous domain, the regularizing smoothing function is constructed to be

$$\mathbf{d}^*(\mathbf{u}) = \arg \min_{\mathbf{d}} \int_S \{c(\mathbf{u})[\mathbf{d}(\mathbf{u}) - \mathbf{d}_{init}(\mathbf{u})]^2 + (\frac{\partial \mathbf{d}(\mathbf{u})}{\partial \mathbf{u}})^2\} d\mathbf{u} \quad (11)$$

Here, S is the surface space, \mathbf{u} is the parameter vector representing a point in S , \mathbf{d}_{init} and \mathbf{d}^* are the initial and final displacement vectors, respectively, and $c(\mathbf{u})$ is the confidence measure. Since we are dealing with discrete triangulated polygonal surfaces, the derivative of the displacement vector $\mathbf{d}(\mathbf{u})$ at \mathbf{u} is approximated by the weighted average finite differences performed on the neighboring points around \mathbf{u} :

$$\frac{\partial \mathbf{d}(\mathbf{u})}{\partial \mathbf{u}} = \sum_{i=0}^{N-1} \omega_i \{\mathbf{d}(\mathbf{u}) - \mathbf{d}(\mathbf{u}_i)\}, \quad i = 0, \dots, N-1 \quad (12)$$

$$\sum_{i=0}^{N-1} \omega_i = 1 \quad (13)$$

where $\{\mathbf{d}(\mathbf{u}_i)\}$ is the displacement vector set associated with the 1-order natural neighbor points, and $\{\omega_i\}$ is the natural neighbor weighting set associated with these neighbor points [54]. Furthermore, the discrete version of the smoothing process can be conveniently posed as a series of linear equations, and can be solved iteratively:

$$\mathbf{d}(\mathbf{u})_{new} = (1 - c'(\mathbf{u}))\mathbf{d}(\mathbf{u})_{old} + c'(\mathbf{u}) \sum_{i=0}^{N-1} \omega_i \mathbf{d}(\mathbf{u}_i) \quad (14)$$

$$c'(\mathbf{u}) = \frac{1}{c(\mathbf{u}) + 1} \quad (15)$$

In the same spirit of the two-stage, non-shrinkage Gaussian smoothing of the triangulated surface described above, the regularization process is also solved by a two-stage iterative procedure:

$$\mathbf{d}(\mathbf{u})_{new1} = (1 - c'_1(\mathbf{u}))\mathbf{d}(\mathbf{u})_{old} + c'_1(\mathbf{u}) \sum_{i=0}^{N-1} \omega_i \mathbf{d}(\mathbf{u}_i) \quad (16)$$

$$\mathbf{d}(\mathbf{u})_{new2} = (1 - c'_2(\mathbf{u}))\mathbf{d}(\mathbf{u})_{new1} + c'_2(\mathbf{u}) \sum_{i=0}^{N-1} \omega_i \mathbf{d}(\mathbf{u}_i) \quad (17)$$

where c'_1 and c'_2 are derived from the confidence coefficient with $0 < c'_1 < -c'_2$, and their difference

are determined by a bounded band [54, 57].

The solution of this system of linear equations will yield a dense set of smoothed motion vectors that adhere to the reliable initial shape-based matches, yet vary smoothly on the surface space. Repeating the process for all the LV wall surfaces and assembling the vectors end-to-end over the sixteen-frame temporal sequence, we have a dense set of motion trajectories of the myocardial surface points throughout the cardiac cycle. Figure 8 shows the dense endocardial displacement vector field from ED to ES, sub-sampled for visualization purposes. The trajectories are shown against the rendered endocardial surface at ES. We usually have about 1,800 endocardial and 2,600 epicardial points for a typical MRI dataset, and around 8,000 endocardial and 15,000 epicardial points for a typical DSR dataset.

5 Experimental Validation and Evaluation

To evaluate the efficacy of using image-derived *in vivo* measures of regional LV function for predicting the extent of regional myocardial injury, we conducted experiments on fasting anesthetized, open chest, adult mongrel dogs at the Yale University School of Medicine and the Mayo Clinic. The results reported here are obtained from eight sets of canine magnetic resonance imaging (MRI) data under both baseline and post-infarction conditions, and four sets of canine Dynamic Spatial Reconstructor (DSR) data.

The purpose of the animal study is to validate our shape-based boundary motion tracking approach, and to obtain the most accurate and reproducible measures for characterizing *in vivo* measures that best predict *post mortem* myocardial injury. The first step has to do with validating the algorithm-derived motion versus the implanted-marker-derived motion (using both MRI and DSR data). The trajectories of surface points which are near implanted imaging-opaque markers are compared to those of the markers. The second step is to evaluate how well the algorithm-derived measures predict the spatial extent of *post mortem* myocardial injury (MRI data only). The changes of motion measurements between baseline and post-infarct data are quantified to estimate the extent of myocardial injury, and are compared to post-mortem histo-chemical staining maps of the actual injury (TTC-stained post-mortem optical photographs).

The clinical usefulness of the proposed motion analysis framework lies with its ability to identify disease states from derived cardiac regional function measures, such as path length and wall thickening. Since baseline values are most likely to be unavailable for comparison, it is worthwhile

to point out that effort is underway to establish a database of baseline distributions of these physiological measures, against which conditions depicted by a new dataset can be evaluated. It will be the subject of a forthcoming paper.

5.1 Image Data

5.1.1 Magnetic Resonance Imaging Data

Pairs of endocardial and epicardial markers were implanted. They were loosely tethered, paired combinations of small copper beads (which show up as small dark signal voids in the images) at the endocardial wall and small Plexiglas capsules filled with a 200:1 mixture of water to Gd-DTPA at the epicardial wall (which show up as small bright signal spots in the images). The markers within each pair were attached by an elastic string, which maintained the endocardial and epicardial markers on their respective surfaces during the entire cardiac cycle without restricting myocardial thickening. Marker pairs were placed in four locations on the canine heart: 1) central infarct zone, 2) normal zone, 3) peri-infarct border zone, and 4) the site of coronary occlusion. These zones were identified by the experimental cardiologists who were performing the animal surgeries. An example of an MRI image which contains both the endocardial and epicardial markers is shown in Figure 9.

ECG-gated magnetic resonance imaging was performed on a GE Signa 1.5 Tesla scanner at Yale. Axial images through the LV were obtained with the gradient echo cine technique. The imaging parameters were as the following: section thickness 5 mm, no inter-section gap, 40 cm field of view, TE 13 msec, TR 28 msec, flip angle 30 degrees, flow compensation in the slice and read gradient directions, 256×128 matrix and 2 excitation. The resulting 3D image set consists of sixteen 2D image slices per temporal frame, and sixteen temporal 3D frames per cardiac cycle.

Dogs were positioned in the magnetic resonance scanner for initial imaging under baseline conditions. The left anterior descending coronary artery was then occluded, without noticeable movement of the dog relative to the imaging planes. Images were again acquired after 10 minutes of coronary occlusion. This way, it was assured that the images were roughly registered with each other pre- and post-occlusion. Figure 10 shows sixteen spatial slices which cover the entire left ventricle, from apex to base. Sixteen temporal frames of a mid-ventricle slice are shown in Figure 11, from ED to ES and then back to ED.

The location of each implanted marker was determined in each temporal frame by first manually delineating a point within the marker area in one of the 2D image slices, grouping all

evidence of the brightest (or darkest) values in the vicinity of that point over all slices within the 3D frame using a connected components algorithm and then computing the 3D centroid of that cluster of points, weighted by the gray level of each point. The markers were spatially placed far apart enough so that trajectories can be formed by simple nearest Euclidean distance correspondence between centroids in adjacent temporal frames. It should be noted that because of the complexity of the implantation of the markers, not all the markers are found all the time. However, for the eight studies presented here, 101 out of 128 possible markers have been found at all the sixteen time frames. Finally, we note that in a 2D moving phantom study, the trajectories traced out by tracking the identified marker from a cine gradient-echo MRI acquisition using the technique just described and the known physical trajectory of the phantom were virtually identical.

5.1.2 Dynamic Spatial Reconstructor Data

The Dynamic Spatial Reconstructor is a three-dimensional computerized tomography scanner at Mayo Clinic [48, 47]. It can provide accurate, stop-action images of moving organs of the body. The canine data we are using was acquired at 33 msec frame intervals in real time, with the spatial resolution of 0.91mm in all three dimensions. During the DSR imaging, four pairs of untethered small lead beads (X-ray opaque) were placed on the endocardium and epicardium, and they appeared as small bright spots in the images. The markers were hand-traced as the average positions of the most probable points. Since the DSR has very high spatial resolution in all three dimensions, the image data usually contain fifty to sixty slices to cover the LV, and have sixteen to eighteen temporal volumes. Sixteen sample spatial slices which cover the left ventricle are shown in Figure 12, and sixteen temporal frames of a mid-ventricle slice are shown in Figure 13. Here, the spatial slices cover the LV from apex to base, and the temporal frames cover the cardiac cycle from ED to ES and then back to ED.

5.2 Quantitative Measures of LV Motion and Function

The complete trajectories of all the myocardial surface points are built by concatenating motion vectors between pairs of surfaces found from consecutive temporal frames, using the proposed method. Flow vectors running through the surface points over time are connected to form trajectories that track each individual point's motion over the entire cardiac cycle. Indices for LV regional function can be derived from these trajectories, including endocardial-epicardial surface motion measures, LV thickening measures, and LV cross-wall strain measures, etc. Each of these

measures are recorded for each point trajectory on each heart image data set. A subscript α will be used to differentiate algorithm-derived measures ($\alpha = alg$) from gold-standard-derived measures ($\alpha = gold$). The gold standard measurements for the validation are derived from the trajectories of the implanted markers described in the *in vivo* acute animal model. From these two types of measures, correlations are observed between the shape-tracked functional indices and the marker-derived indices.

We have proposed several functional measures for the analysis of the left ventricle. The single surface measures (on either endocardial or epicardial surface) include the overall path length (sum of the magnitude of the flow vectors) and the displacement vector with respect to some known cardiac state such as end-diastole. The LV thickening measures (the cross-wall measures) are computed using cross-wall vectors which relate points on the endocardial surface to corresponding points on the epicardial surface. The magnitude of each of these connecting vectors is taken as a measure of wall thickness, and its temporal change is an indication of the wall thickening.

5.3 Motion Validation

In order to validate the motion trajectories resulting from the shape-based tracking algorithm, the gold-standard marker trajectories and the algorithm-derived point trajectories are compared across the acute dog studies (eight MRI datasets and four DSR datasets). To avoid the possible distortion of the boundary surfaces by the implanted markers which may cause false shape features, the image intensity around each marker is interpolated to be the average gray level of the surrounding same type tissue after the markers' positions have been detected. This procedure is performed before the contour segmentation process. Since the position of each marker is the 3D centroid of a cluster of candidate points in the image, the markers are usually not exactly on the myocardial surfaces. For both MRI and DSR studies, the point which is the nearest surface point to each marker is chosen as the *surface marker point* for the given physical marker. After offsetting the initial positional difference, the algorithm-derived trajectory of this surface marker point is considered as the algorithm-derived trajectory of the physical marker. Then, point-by-point positional errors $\|\mathbf{x}_{alg}(i) - \mathbf{x}_{gold}(i)\|$ between the algorithm and gold standard trajectories are compiled across all 16 time frames for each marker, where $\mathbf{x}_\alpha(i)$ is the 3D position of physical marker ($\alpha = alg$) and the surface marker point ($\alpha = gold$) at frame i . The mean and the standard deviations of the positional errors for all the markers of both MRI and DSR data are represented in Table 1.

For the baseline MRI studies, endocardial markers have a mean error of about one pixel and

the epicardial markers have a mean error around two-thirds of a pixel. Since the endocardial points usually move two-to-three times the distance of the corresponding epicardial points, this mean error disparity is somewhat expected. However, the endocardial surface often has more shape features than the epicardial surface, which may explain that both of them have similar error standard deviations.

For the post-infarct MRI studies, both endocardial and epicardial markers have mean errors in the half pixel range, with their error standard deviations also much smaller than those under baseline condition. This is because the infarction dramatically reduces the motion of the myocardium, especially at the endocardium (a thirty percent drop of average path length from baseline to post-infarct for the endocardial surface, a ten percent drop of average path length from baseline to post-infarct for the epicardial surface). This reduction of motion makes it less likely to make large matching errors. Even though it still has larger motion, the feature-rich endocardial surface now has similar error statistics as the epicardial surface.

Since the motion tracking algorithm is based on matching local shape features, it is not surprising that the higher resolution DSR baseline studies produce better endocardial tracking results than the MRI studies. The high resolution DSR images pick up more local detail of the endocardium, thus have smaller mean errors as validated by the markers. Meanwhile, the accuracy of the relatively feature-less epicardium does not improve, in terms of pixel units, from the MRI studies, and it is slightly worse than the endocardial results. Evidently, the rich shape features of the endocardial surface are sufficient to track the larger motion presented in this case.

Overall, this table illustrates that the motion trajectories predicted by the algorithm at points near the implanted markers are very close to the gold-standard motion of the markers, and are more accurate with better image resolution. No mean errors are statistically significantly different from the in-plane resolution of the image data. A visual illustration of the marker position comparison for one MRI baseline study and one DSR baseline study is shown in Figure 14. Here, the real physical marker (green) trajectories and surface marker trajectories (blue), both from ED to ES, are plotted together, against the endocardial surfaces at ES. It is evident from these figures that the algorithm trajectories are usually very close to the corresponding marker trajectories.

5.4 Injury Validation

For five MRI studies under both baseline and post-infarct conditions, two *in vivo* measures (path length and wall thickening) computed from the algorithm-derived trajectories are used to evaluate

how well the shape-based motion tracking approach can distinguish and predict myocardial regional injury. Three other MRI studies are not included in the study because the implanted markers in their normal and/or infarct zone are missing from some of the image frames which make it difficult to identify the zones unambiguously.

5.4.1 Path Length

The path length of any surface point is the sum of the magnitudes of all sixteen displacement vectors of this point over the cardiac cycle, and it measures the overall motion of the point. One physiological hypothesis of myocardial injury is that the post-infarct heart will have dramatically decreased motion in the injured zone from its baseline condition. Hence, it is reasonable to expect that we can measure myocardial regional injury from the comparison of the baseline and post-infarct path length maps of the myocardial surfaces.

Two validation tests are performed on the path length data. The first test verifies the validity of using path length changes to differentiate the normal zone and the infarct zone motion behavior. For both the baseline and the post-infarct data of five MRI studies, the mean and standard deviation of the path lengths are computed for the antero-apical infarct zone and the remote normal zone from sampled points near the two implanted markers. These data are shown in Table 2. Two-way analysis of variance is performed for the path length measure, where the F -statistics and p values are computed for comparisons between a) baseline and post-infarct conditions, b) infarct zone and normal zone, and the interaction of a) and b). As can be seen in the table, the path length measure shows significant differences in the interaction term, indicating that the algorithm can discriminate between baseline and infarction conditions, and it can differentiate the motion at the location of the normal and infarct markers. More detailed testing showed that the significance of the interaction is due to the difference of the measures in the infarct zone at baseline and post-infarct conditions ($p < 0.028$), and the difference between the infarct zone and the normal zone at post-infarct condition ($p < 0.022$).

The second test verifies the algorithm-predicted extent (area) of the infarct zone on the endocardial surface against the post-mortem definition of myocardial injury. Three-dimensional *post mortem* endocardial injury zones are found by digitizing color photographs of TTC-stained *post mortem* myocardium slices (5mm thick) of the excised hearts. The endocardial, epicardial and infarction zone boundaries of each *post mortem* left ventricular slice are hand-traced, aligned, and stacked to reconstruct the LV surface representations. The areas of the endocardial infarction zones

are computed and compared to the areas defined by thresholding the magnitudes of the *in vivo* algorithm-derived path length measures. Within a search window defined by a sphere centered at the central infarct marker I with radius being the distance between I and the occlusion marker O , points are considered to be inside the infarct zone if their post-infarct path length fell to levels below 80% of the mean infarct zone baseline value, computed from sample points near the infarct zone marker. Figure 15 shows the injury zones depicted by *post mortem* TTC-staining (left) and *in vivo* algorithm-derived trajectories (right). The mean and standard deviation of both *post mortem* and *in vivo* infarct areas are reported in Table 3, as is their Pearson product-moment correlation. The high correlation value points to the utility of these measures for measuring the *extent* of endocardial injury.

5.4.2 Wall Thickening

We have shown that the algorithm-derived path length changes of the myocardial surface points predict the existence and extent of the infarct zone. However, the path length measure is a gross overall measure of local deformation *and* global motion, which may result in over- or under- estimation of the injury. This is because the motion of the injury zone may interact with adjacent normal tissue. Meanwhile, myocardial thickening, measured by the change of thickness of LV wall between ED and ES, may provide additional insight into myocardial viability since it is a more local measure similar to radial strain ¹. For every epicardial point, its thickening measure is defined as:

$$\tau' = \frac{\tau_{ED} - \tau_{ES}}{\tau_{ED}} \quad (18)$$

where τ_{ED} and τ_{ES} are the thickness of the LV wall at ED and ES. The thickness of the wall is computed by emitting a ray from the epicardial point in its inward normal direction until it intersects the endocardial surface, and the distance between the intersection and the epicardial point is considered to be the wall thickness.

For both the baseline and the post infarct data from five MRI studies, the mean and standard deviation of the thickening measures are computed for the antero-apical infarct zone and the remote

¹In fact, this lays the groundwork for our further efforts in estimating transmural properties of the left ventricle, see [51, 53]

normal zone from sampled points near the two implanted markers. The data are shown in Table 4. Two-way analysis of variance is again performed for the change of thickening measure, and the F -statistics and p values are computed for comparisons between a) baseline and post-infarct conditions, b) infarct zone and normal zone, and the interaction of a) and b). As can be seen in the table, the thickening measure shows significant differences in the interaction term, indicating that the algorithm can distinguish infarction. More detailed testing showed that the significance of the interaction is due to the difference of the measure in the infarct zone at baseline and post-infarct conditions ($p < 0.010$), and the difference between the infarct zone and the normal zone at post-infarct conditions ($p < 0.045$). This shows the validity of using our algorithm-derived, point-tracked LV wall thickening measure to differentiate the normal zone and the infarct zone motion behavior. And since thickening is a gross approximation to the radial strain of the wall, it shows that shape-based tracking may be quite useful for estimating strains near or across the endocardial and epicardial surfaces.

6 Conclusions

In this paper, we have described a shape-based methodology aimed at accurately and objectively determining and quantifying the local, regional and global dynamic function of the left ventricle of the heart from 3D image data. It is our hypothesis that we can use the shape properties of the myocardial surfaces to track the motion trajectories of a dense field of points over the entire cardiac cycle, which can then be used to predict the existence and extent of myocardial injury.

Our shape-based tracking algorithm has been validated in a canine model with eight MRI studies under both baseline and post-infarct conditions, and on four DSR baseline studies. The accuracy of the algorithm point tracking results has been validated using the gold-standard motion of implanted imaging-opaque markers with very favorable results. We have shown that the changes in motion parameters (path length and wall thickening measures) between baseline and post-infarct data are indicative of the myocardial function, and that we can use these changes to diagnose the location and extent of myocardial injury, which is validated using *post mortem* TTC staining technique.

Further and future research involves incorporation of temporal periodic characteristics of the heart motion within the current framework, integration of the transmural information from phase contrast MR images and/or tagged MR images, and developing richer surface models. We

are also investigating the proper ways to build a standardized canine, human in the future, heart atlas of representative baseline function distributions using this method.

Study Type	Image Resolution. (<i>mm/pixel</i>)	Average Errors	
		Endocardial Markers	Epicardial Markers
MRI Baseline	1.64	0.993 ± 0.522 (<i>pixel</i>)	0.685 ± 0.509 (<i>pixel</i>)
		1.623 ± 0.856 (<i>mm</i>)	1.123 ± 0.835 (<i>mm</i>)
MRI Post-Infarct	1.64	0.488 ± 0.322 (<i>pixel</i>)	0.498 ± 0.339 (<i>pixel</i>)
		0.800 ± 0.528 (<i>mm</i>)	0.817 ± 0.556 (<i>mm</i>)
DSR	0.91	0.556 ± 0.305 (<i>pixel</i>)	0.645 ± 0.338 (<i>pixel</i>)
		0.506 ± 0.278 (<i>mm</i>)	0.587 ± 0.308 (<i>mm</i>)

Table 1: Average positional errors (in *pixel* units and in *mm*) between implanted markers and algorithm-derived point positions for eight MRI baseline, eight MRI post-infarct, and four DSR studies. Each data cell represents the mean \pm standard deviation errors for all the same type markers over the entire temporal frames.

	Path Length	a Factor		a		b		axb	
		Infarct Zone	Normal Zone	F	p	F	p	F	p
b Factor	Baseline	22.1 ± 3.8	21.7 ± 3.5	3.46	0.14	4.50	0.10	10.44	0.03
	Post-Infarct	11.0 ± 7.7	19.9 ± 4.5						

Table 2: Table of means and standard deviations of algorithm-derived *in vivo* path length measures (in *mm*), observed in the infarct zone and in a remote normal zone for five acute canine studies under both baseline and post-infarct conditions.

Study	Infarct Area (mm^2)						Correlation
	1	2	3	4	5	average	
<i>Post Mortem</i>	963	811	833	749	623	796±124	0.968
<i>In Vivo</i>	939	855	807	710	602	783±131	

Table 3: Correlation between *post mortem* and *in vivo* infarct areas.

	Thickening Measure(τ')	a Factor		a		b		axb	
		Infarct Zone	Normal Zone	F	p	F	p	F	p
b	Baseline	44.6 ± 26.2	34.7 ± 19.5	0.82	0.42	4.90	0.09	44.7	0.003
Factor	Post-Infarct	6.1 ± 13.2	37.4 ± 13.4						

Table 4: Table of means and standard deviations of algorithm-derived *in vivo* measure of radial thickness change (in percentage), observed in the infarct zone and in a remote normal zone for five acute canine studies under both baseline and post-infarct conditions.




Figure 1: LV boundary segmentation. Left: segmented endocardial and epicardial boundaries superimposed on one baseline MR image; Middle: stacked endocardial contours; Right: stacked epicardial contours.

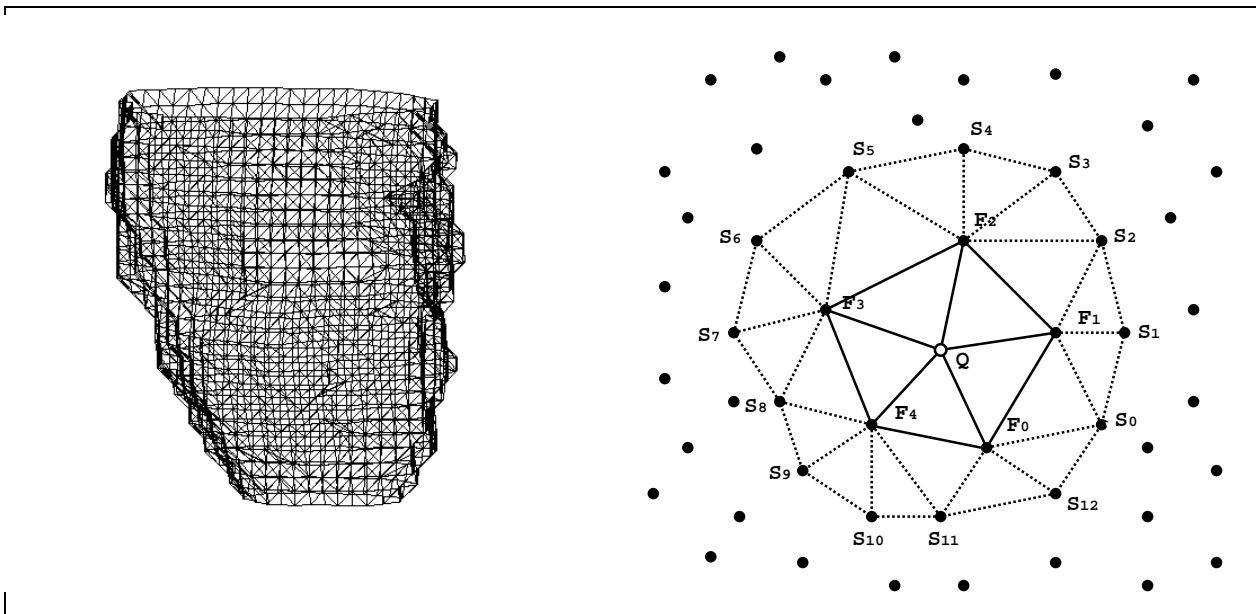


Figure 2: Left: A Delaunay tessellated MR endocardial surface representation. Right: Local surface patch defined by multiple *orders* of natural neighboring points. For point Q , points $F_i, i = 0, \dots, 4$ are its 1-order natural neighbor points, points $S_i, i = 0, \dots, 12$ are its 2-order natural neighbor points.




Figure 3: Regular and non-shrinkage Gaussian smoothing of triangulated surface. Top: the original triangulated endocardial surface (left: rendered; right: wireframe); Lower: the smoothed triangulated endocardial surface (left: the non-shrinkage Gaussian smoothing result with $\lambda'_1 = 0.33$, $\lambda'_2 = -0.34$ and iteration number $n = 10$; right: regular Gaussian smoothing result with $\lambda' = 0.33$, and iteration number $n = 10$).




Figure 4: Left: Possible multiple mappings in the graph representation of surface patch. The surface is the solid line, with the surface points the filled dots (inside the neighborhood sphere). Obviously, the graph function $z = h(x, y)$ is not a unique mapping. Right: Invalid neighborhood points. The valid surface points are the filled dots, while the unfilled dots are invalid points even though they are inside the neighborhood sphere (the unfilled dots are disconnected from the valid surface points, and belong to another surface patch).




Figure 5: Curvature maps of an endocardial surface derived from multi-scale graph functions. The first column is for 1-scale, the second column is for 2-scale, and the third column is for 3-scale graph functions. From top to bottom row: Gaussian curvature, mean curvature, maximum principal curvature, and minimum principal curvature maps. For the curvature maps, the values from negative to positive are represented by a continuous color spectrum from red to green. The value-to-color scale is the same for all three maps within each row, but it is different between different curvatures.




Figure 6: Potential energy maps of an endocardial surface over the entire cardiac cycle. The temporal sequence of the maps runs from top to bottom, left to right. The potential energy maps have the same color scale: darker for larger values, brighter for smaller values.

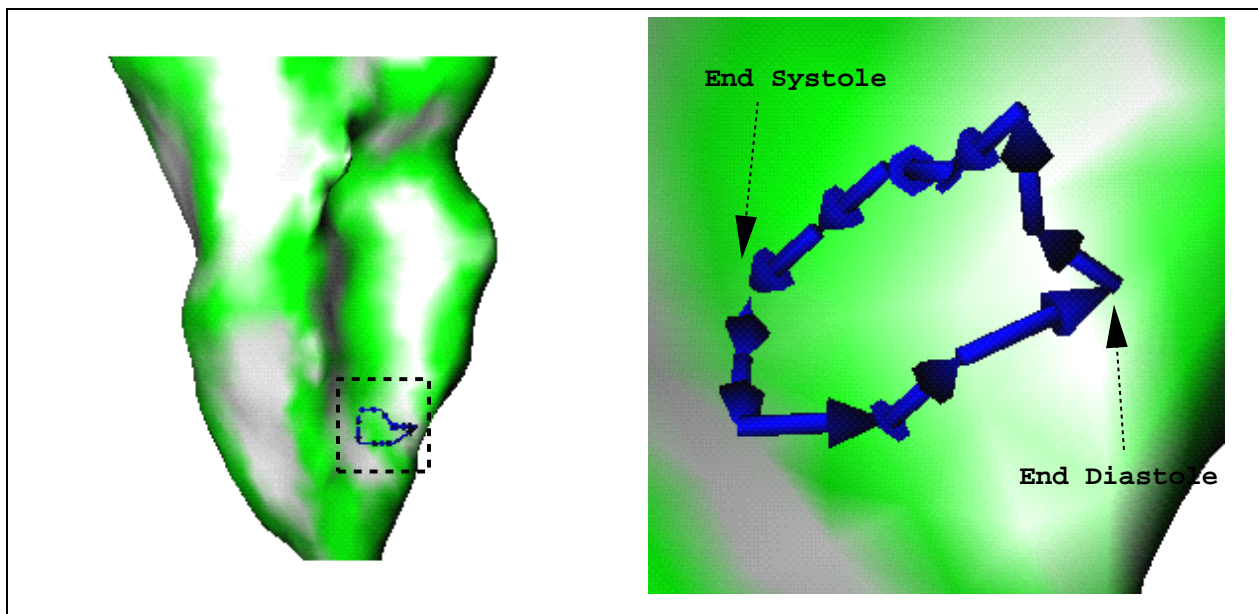


Figure 7: Complete trajectory of one endocardial point over the cardiac cycle (sixteen temporal frames), superimposed onto the potential energy maps the endocardial surface at end systole. Top: an overview ; Bottom: a blowup of the trajectory.




Figure 8: Dense (sub-sampled for visualization purposes) endocardial displacement vector field from ED to ES. The trajectory arrows are shown against the rendered endocardial surface at ES.

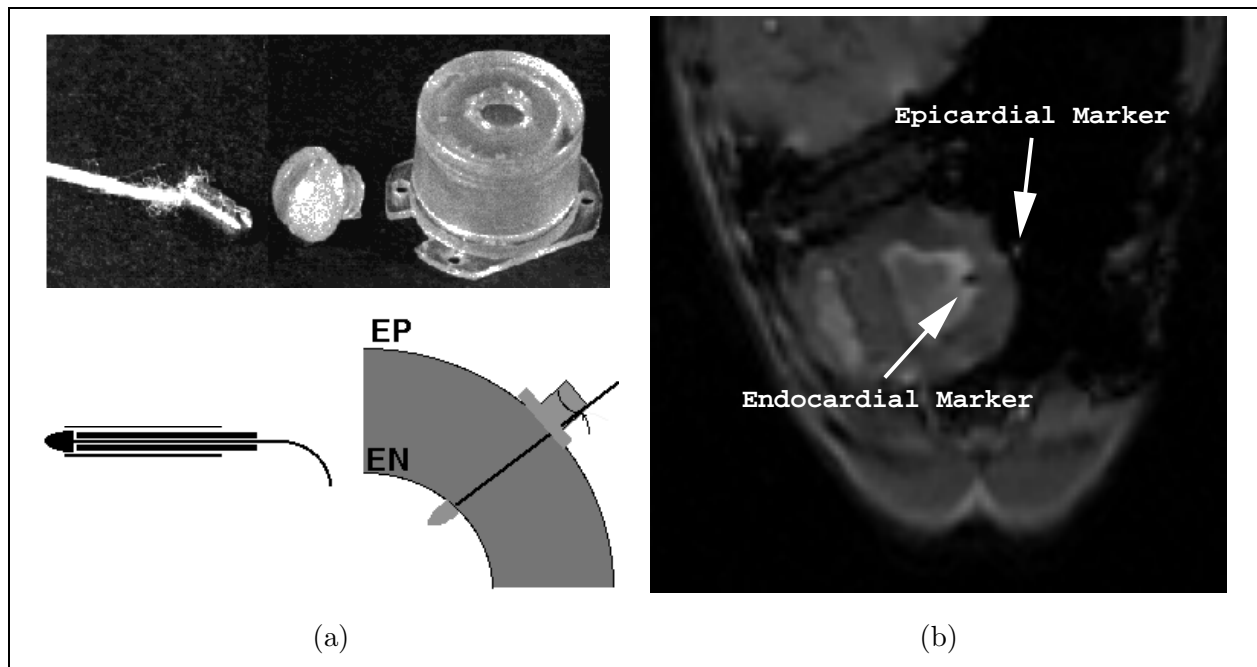


Figure 9: MRI markers and their relative positions. Left: the upper picture shows the endocardial marker (copper bead on the left) and the epicardial marker (capsule on the right), the lower diagram illustrates the procedure to insert the endocardial marker into the myocardium through a needle as well as the relative positions of the markers on the myocardium; Right: MRI image with endocardial and epicardial markers.




Figure 10: Sixteen spatial MRI images which cover the entire left ventricle. The images run from top to bottom, left to right, to cover the LV from apex to base.




Figure 11: Sixteen temporal MRI images of a mid-ventricle slice. The images run from top to bottom, left to right, for the temporal sequence of ED to ES to ED.




Figure 12: Sixteen spatial DSR images which cover the left ventricle. The images run from top to bottom, left to right, to cover the LV from apex to base. Note the implanted markers which appear as the bright dots in the images.




Figure 13: Sixteen temporal DSR images of a mid-ventricle slice. The images run from top to bottom, left to right, for the temporal sequence of ED to ES to ED.

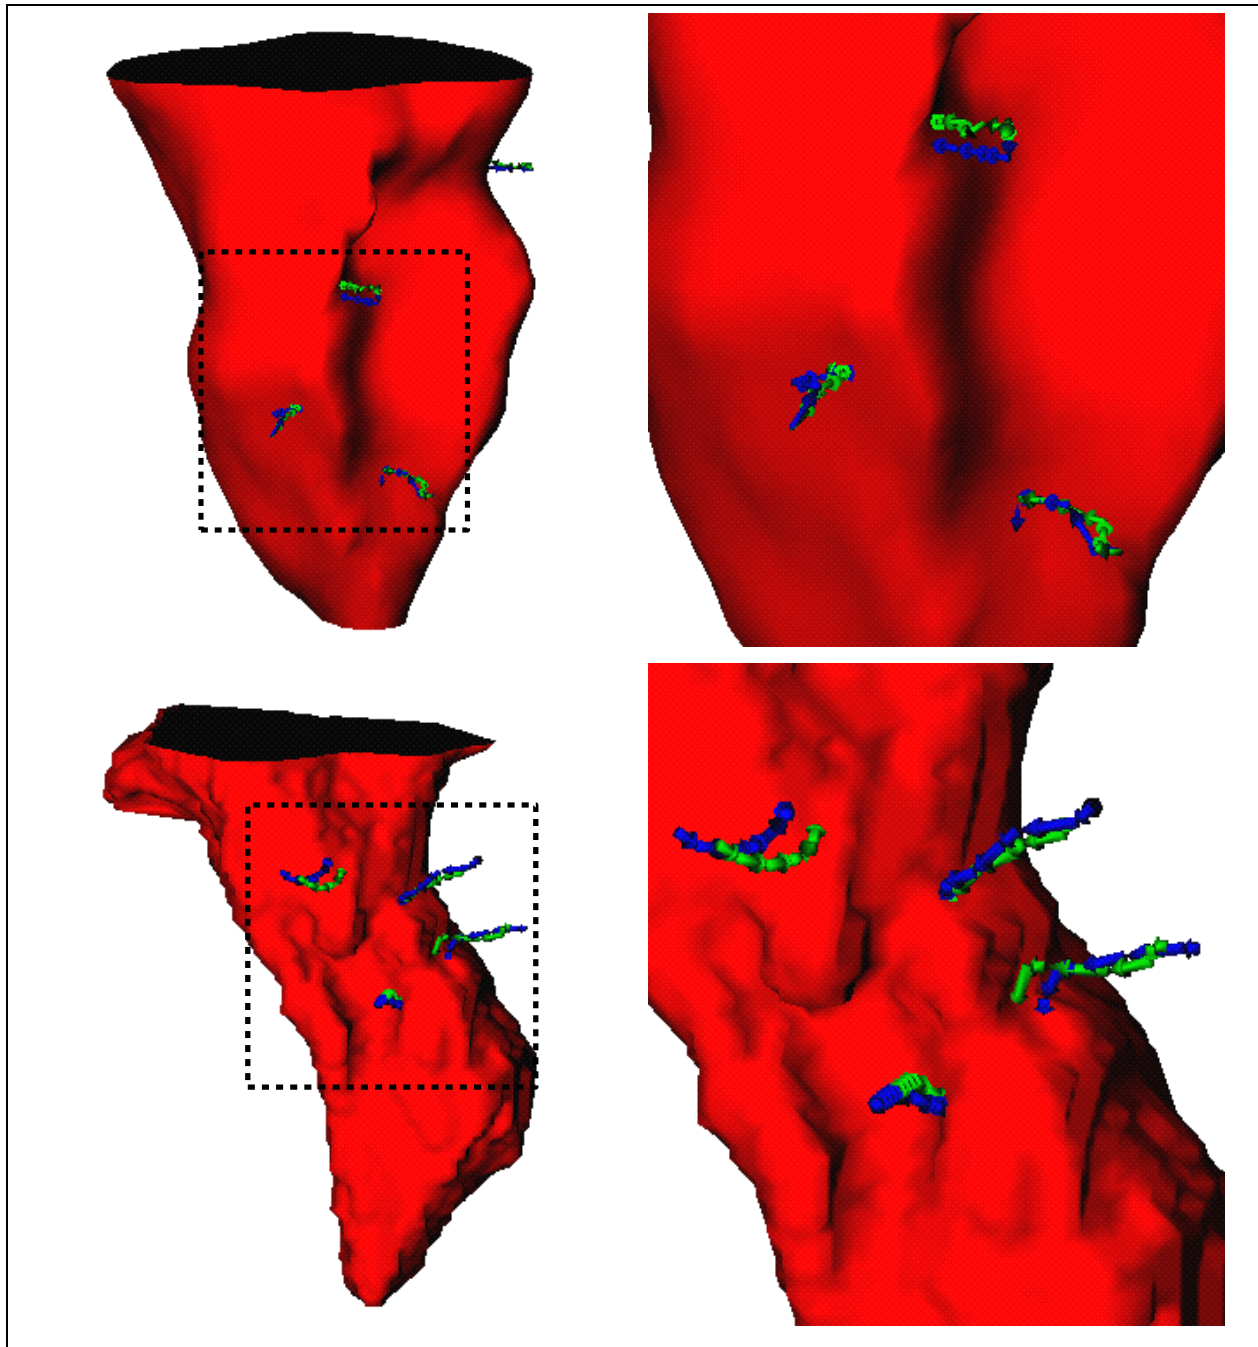


Figure 14: Comparison of algorithm and gold standard marker trajectories. Top (MRI study): algorithm trajectory (blue) and the implanted marker trajectory (green); Bottom (DSR study): algorithm trajectory (blue) and the implanted marker trajectory (green). In both studies, the paths are shown relative to the endocardial surfaces rendered at end-systole.




Figure 15: Comparison of TTC-staining and motion algorithm derived injury zones. Left: rendered endocardial surface along with the injury zone (dark) from *post mortem* TTC staining; Right: rendered endocardial surface along with the injury zone (dark) from algorithm-derived *in vivo* motion trajectories.

References

- [1] A. Amini, R. Curwen, and J. Gore. Snakes and splines for tracking non-rigid heart motion. In *European Conference on Computer Vision*, Cambridge, England, 1996.
- [2] A. A. Amini and J. S. Duncan. Bending and stretching models for LV wall motion analysis from curves and surfaces. *Image and Vision Computing*, 10(6):418–430, 1992.
- [3] A. A. Amini and J. S. Duncan. Differential geometry for characterizing 3D shape change. *SPIE- Mathematical Methods in Medical Imaging*, 1768:170–181, 1992.
- [4] P. Anandan. A computational framework and an algorithm for the measurement of visual motion. *International Journal of Computer Vision*, 2:283–310, 1989.
- [5] L. Axel and L. Dougherty. MR imaging of motion with spatial modulation of magnetization. *Radiology*, 171:841–845, 1989.
- [6] Azhari et al. Noninvasive quantification of principal strains in normal canine hearts using tagged MRI images in 3D. *American Journal of Physiology*, 264:H205–H216, 1993.
- [7] S. T. Barnard and W. B. Thompson. Disparity analysis of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2(4):333–340, 1980.
- [8] E. L. Bolson, S. Kliman, F. Sheehan, and H. T. Dodge. Left ventricular segmental wall motion - A new method using local direction information. In *Proceedings of Computers in Cardiology*, 1980.
- [9] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 567–585, 1989.
- [10] G. Borgefors. Distance transformations in arbitrary dimensions. *Computer Vision, Graphics and Image Processing*, 27:321–345, 1984.
- [11] G. Borgefors. Distance transformations in digital images. *Computer Vision, Graphics and Image Processing*, 34:344–371, 1986.
- [12] A. Chakraborty, L. H. Staib, and J. S. Duncan. An integrated approach to boundary finding in medical images. In *Proceedings of the IEEE Workshop on Biomedical Image Analysis*, pages 13–22, 1994.

- [13] Amit Chakraborty. *Feature and Module Integration for Image Segmentation*. PhD thesis, Yale University, New Haven, CT, 1996.
- [14] J. Chesebro, G. Knatterud, and others. Thrombolysis in myocardial infarction (TIMI) trial phase I: A comparison between intravenous tissue plasminogen activator and intravenous streptokinase. *Circulation*, 76:142–154, 1987.
- [15] R.T. Constable, K. Rath, A. Sinusas, and J. 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.
- [16] R. Courant and D. Hilbert. *Methods of Mathematical Physics*. Interscience, London, 1957.
- [17] C. Davatzikos and R. N. Bryan. Using a deformable surface model to obtain a shape representation of the cortex. *IEEE Transactions on Medical Imaging*, 15:785–795, December 1996.
- [18] T. Denney and J. L. Prince. 3D displacement field reconstruction from planar tagged cardiac MR images. In *Proceedings of the IEEE Workshop on Biomedical Image Analysis*, pages 51–60, 1994.
- [19] M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, New Jersey, 1976.
- [20] 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.
- [21] F. Glazer, G. Reynolds, and P. Anandan. Scene matching by hierarchical correlation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 432–441, 1983.
- [22] W. Grossman. Assessment of regional myocardial function. *Journal of American College of Cardiology*, 7(2):327–328, 1986.
- [23] R. J. Herfkens, N. J. Pelc, L. R. Pelc, and J. Sayre. Right ventricular strain measured by phase contrast MRI. In *Proceedings of the 10th Annual Society of Magnetic Resonance Medicine*, page 163, San Francisco, 1991.
- [24] G. T. Herman, J. Zheng, and C. A. Bucholtz. Shape-based interpolation. *IEEE Computer Graphics and Applications*, pages 69–79, 1992.

- [25] E. C. Hildreth. *The Measurement of Visual Motion*. MIT Press, Cambridge, 1984.
- [26] B. K. P. Horn. *Robot Vision*. MIT Press, Cambridge, 1986.
- [27] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [28] B. Horowitz and S. Pentland. Recovery of non-rigid motion and structure. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 325–330, Maui, June 1991.
- [29] W. C. Huang and D. B. Goldgof. Adaptive-size physically-based models for nonrigid motion analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 833–835, 1992.
- [30] C. Kambhamettu and D. Goldgof. Point correspondence recovery in non-rigid motion. In *Computer Vision and Pattern Recognition*, pages 222–227, June 1992.
- [31] G. E. Mailloux, A. Bleau, M. Bertrand, and R. Petitclerc. Computer analysis of heart motion from two dimensional echocardiograms. *IEEE Transactions on Biomedical Engineering*, BME-34(5):356–364, May 1987.
- [32] J. C. McEachen, A. Nehorai, and J. S. Duncan. A recursive filter for temporal analysis of cardiac motion. In *Proceedings of the IEEE Workshop on Biomedical Image Analysis*, pages 124–133, 1994.
- [33] D. Metaxas and D. Terzopoulos. Constrained deformable superquadrics and nonrigid motion tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 337–343, 1991.
- [34] F. G. Meyer, R. T. Constable, A. G. Sinusas, and J. S. Duncan. Tracking myocardial deformation using spatially constrained velocities. In *Information Processing in Medical Imaging*. Kluwer, 1995.
- [35] S. Mishra, D. Goldgof, and T. Huang. Motion analysis and epicardial deformation estimation from angiography data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 331–336, Maui, June 1991.

- [36] H. H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.
- [37] C. Nastar and N. Ayache. Non-rigid motion analysis in medical images: a physically based approach. In *Information Processing in Medical Imaging*. Springer-Verlag, 1993.
- [38] G. Naylor, N. Firmin, and D. Longmore. Blood flow imaging by cine magnetic resonance. *Journal of Computer Assisted Tomography*, 10:715–722, 1986.
- [39] W.G. O’Dell, C.C. Moore, W. Hunter, E.A. Zerhouni, and E.R. McVeigh. Displacement field fitting for calculating 3D myocardial deformations from tagged MR images. *Radiology*, 195:829–835, 1995.
- [40] R. Owen, L. Staib, P. Anandan, and J. Duncan. Measurement of non-rigid motion in images using contour shape descriptors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1991.
- [41] J. Park, D. Metaxas, and L. Axel. Volumetric deformable models with parameter functions: a new approach to the 3D motion analysis of the LV from MRI-SPAMM. In *Fifth International Conference on Computer Vision*, pages 700–705, 1995.
- [42] J. Park, D. Metaxas, and A. Young. Deformable models with parameter functions: application to heart-wall modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 437–442, 1994.
- [43] N. J. Pelc and M. A. Bernstein. Optimized encoding for phase contrast flow measurement. In *Proceedings of the 9th Annual Society of Magnetic Resonance Medicine*, page 475, New York, 1990.
- [44] N. J. Pelc, R. J. Herfkens, and L. R. Pelc. 3D analysis of myocardial motion and deformation with phase contrast cine MRI. In *Proceedings of the 11th Annual Society of Magnetic Resonance Medicine*, page 18, Berlin, 1992.
- [45] N. J. Pelc, R. J. Herfkens, A. Shimakawa, and D. Enzmann. Phase contrast cine magnetic resonance imaging. *Magnetic Resonance Quarterly*, 7(4):229–254, 1991.

- [46] N. J. Pelc, A. Shimakawa, and G. H. Glover. Phase contrast cine MRI. In *Proceedings of the 8th Annual Society of Magnetic Resonance Medicine*, page 101, Amsterdam, 1989.
- [47] E. L. Ritman. State of the art and future perspectives-fast CT for quantitative cardiac analysis. *Mayo Clinic Proceedings*, 65:1336–1349, 1990.
- [48] E. L. Ritman, J. H. Kinsey, L. D. Harris, and R. A. Robb. Some imaging requirements for quantification of structure and function of the heart. In M. D. Short, D. A. Pay, S. Leeman, and R. M. Harrison, editors, *Physical Techniques in Cardiological Imaging*, pages 189–198. Adam Hilger, Ltd., 1982.
- [49] P. T. Sander and S. W. Zucker. Inferring surface trace and differential structure from 3D images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(9):833–854, 1990.
- [50] S. Sclaroff and A. P. Pentland. Modal matching for correspondence and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(6):545–561, 1995.
- [51] P. Shi, G. Robinson, R. T. Constable, A. Sinusas, and J. Duncan. A model-based integrated approach to track myocardial deformation using displacement and velocity constraints. In *Fifth International Conference on Computer Vision*, pages 687–692, 1995.
- [52] P. Shi, G. Robinson, and J. Duncan. Myocardial motion and function assessment using 4D images. In *Visualization in Biomedical Computing*, pages 148–159, 1994.
- [53] P. Shi, A. J. Sinusas, R. T. Constable, and J. S. Duncan. Volumetric deformation analysis using mechanics-based data fusion: Applications in cardiac motion recovery. *International Journal of Computer Vision*, in press.
- [54] Pengcheng Shi. *Image Analysis of 3D Cardiac Motion Using Physical and Geometrical Models*. PhD thesis, Yale University, New Haven, CT, 1996.
- [55] C. Slager, T. Hooghoudt, P. Serruys, J. Schuurbiens, and J. Reiber et al. Quantitative assessment of regional left ventricular motion using endocardial landmarks. *Journal of American College of Cardiology*, 7(2):317–326, 1986.
- [56] S. Song and R. Leahy. Computation of 3D velocity fields from 3D cine CT images. *IEEE Transactions on Medical Imaging*, 10:295–306, Sept 1991.

- [57] G. Taubin. Curve and surface smoothing without shrinkage. In *Proceedings of the Fifth International Conference on Computer Vision*, pages 852–857, 1995.
- [58] D. Terzopoulos and D. Metaxas. Dynamic 3D models with local and global deformation: Deformable superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(17), 1991.
- [59] D. Terzopoulos and M. Vasilescu. Sampling and reconstruction with adaptive meshes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 70–75, 1991.
- [60] M. Tistarelli and G. Marcenaro. Using optical flow to analyze the motion of human body organs from bioimages. In *Proceedings of the IEEE Workshop on Biomedical Image Analysis*, pages 100–109, 1994.
- [61] J. Udupa and G. Herman. *3D Imaging in Medicine*. CRC Press, Boca Raton, 1991.
- [62] D. F. Watson. Computing the n-dimensional delaunay tessellation with application to Voronoi polytopes. *The Computer Journal*, 24:167–172, 1981.
- [63] D. F. Watson and G. M. Philip. Neighborhood-based interpolation. *Geobyte*, 2:12–16, 1987.
- [64] V. Wedeen. Magnetic resonance imaging of myocardial kinematics: technique to detect, localize and quantify the strain rates of active human myocardium. *Magnetic Resonance Medicine*, 27:52–67, 1992.
- [65] A. A. Young and L. Axel. Three-dimensional motion and deformation in the heart wall: estimation from spatial modulation of magnetization – a model-based approach. *Radiology*, 185:241–247, 1992.
- [66] E. Zerhouni et al. Tagging of the human heart by multiplanar selective RF saturation for the analysis of myocardial contraction. In *Abstracts of the Annual Meeting of the Society of Magnetic Resonance in Imaging*, page 10, San Francisco, 1988.