

Segmentation Informed by Manifold Learning

Qilong Zhang, Richard Souvenir, and Robert Pless

Department of Computer Science and Engineering,
Washington University, St. Louis, MO, 63130, USA
{zql, rms2, pless}@cse.wustl.edu

Abstract. In many biomedical imaging applications, video sequences are captured with low resolution and low contrast challenging conditions in which to detect, segment, or track features. When image deformations have just a few underlying causes, such as continuously captured cardiac MRI without breath-holds or gating, the captured images lie on a low-dimensional, non-linear manifold. The manifold structure of such image sets can be extracted by automated methods for manifold learning. Furthermore, the manifold structure of these images offers new constraints for tracking and segmentation of relevant image regions. We illustrate how to incorporate these new constraints within a snake-based energy minimization approach, and demonstrate improvements in using snakes to segment a set of cardiac MRI images in challenging conditions.

1 Introduction

Many diagnostic and medical applications require segmenting particular tissue structures in every frame of a long 2D or 3D data set. This is challenging because medical video images often have low resolution and low contrast. Combining cues between frames is difficult because often tissues move significantly between frames, and this motion may include complicated deformations that do not lend themselves to simple parameterized models. Creating automated tools to understand and parameterize image data that is affected by a small set of deformations has the potential to impact a large set of relevant medical imaging problems.

For the purpose of segmentation or boundary detection, a collection of tools support imposing various priors on the expected solution. Snakes [1], for instance, are a tool for integrating cues from image data with priors on the expected smoothness of a contour. Within video sequences, these smoothness constraints can be extended to enforce temporal consistency, minimizing variation between consecutive frames, providing cues for segmentation in image regions that have particularly low contrast or high noise.

However, there is often additional structure in a set of images beyond consistency between consecutive frames. In particular, medical image sequences often vary due to a small number of factors. For example, in cardio-pulmonary imaging, the patient breathing cycle causes a deformation of the chest cavity, and the heartbeat leads to large deformations of the shape of the heart. It is complicated

to define parametric models of these deformations that fit multiple patients. For any one patient, however, a collection of cardio-pulmonary images forms a 2 dimensional manifold, where each image is indexed by the current phase of the breathing and heartbeat cycles.

The observation within this paper is that the manifold structure of these data sets provides *stronger* constraints between images than temporal consistency. In [2], a method is proposed to incorporate a statistical prior on the shape of the segmenting contour, while our work can be viewed as a method to use manifold learning to help enforce priors on the changes of shape. We offer an example mechanism to exploit automated manifold learning tools as a pre-processing step to provide a new multi-image constraint to be used in energy-minimization based segmentation procedures. We implement these tools for the specific case of extracting left ventricle wall contours in cardio-pulmonary MRI. In this example domain, a variation in the breathing cycle leads to a uniform translation of the heart. Variation in the heartbeat cycle leads to variations in the shape of the heart, but, largely, not its position. These deformations suggest that strong constraints can be placed on the expected variation of the heart contour between images — constraints more specific than general smoothness constraints.

The following section gives a brief overview of related work in manifold learning and image segmentation. This is followed by an explanation of how to use Isomap in order to extract the cardiopulmonary manifold structure. Then, the primary contribution of this paper is presented – the classical snake-based energy function is extended in order to exploit using this manifold structure to fit snakes simultaneously in many images. We conclude by demonstrating the efficacy of these constraints on both simulated and real data.

1.1 Background and Previous Work

This work integrates ideas from snake-based energy minimization and manifold learning. To our knowledge, these ideas have not been explicitly considered together before. In order to ground our later presentation, we first introduce, very briefly, some recent research in the use of snakes in biomedical image analysis and an overview of manifold learning.

Snakes for Medical Image Segmentation. The enormous amount of prior work on snakes is a testimony to the effectiveness of active contours on a wide variety of problems in medical imagery (for general reviews, see example [3, 4]). With respect to tracking in cardiovascular imagery, an important use of deformable models is to measure the dynamic behavior of the human heart, especially the left ventricle. Accurately characterizing heart wall motion is necessary to diagnose and estimate the severity and extent of diseases such as ischemia [5].

For this task, one approach uses a 2D deformable contour model to segment the LV boundary in each slice of an initial image volume. These contours are then used as the initial approximation of the LV boundaries in corresponding slices of the image volume at the next time instant and are then deformed to extract the new set of LV boundaries [6, 7].

A traditional snake is a parametric contour embedded in the image plane, represented as a closed curve $C(s) = (x(s), y(s))^T$, where x and y are the coordinate functions and $s \in [0, 1]$ is the parametric domain. The shape of the contour minimizes the following functional:

$$E = \int_0^1 E_{int}(C(s)) + E_{img}(C(s)) + E_{con}(C(s)) ds, \quad (1)$$

where E_{int} represents the internal energy of the snake due to bending, E_{img} is the image energy derived from the image, and its local minima coincide with intensity extrema, edges, and other image features of interest, and E_{con} is a place holder for additional constraints appropriate for a given context, including prior shape models and limitations on changes between consecutive images. Our contribution within this work is to offer a method for having these additional constraints depend upon the automatically extracted manifold structure of an image set.

Once the energy function is specified, one can obtain the snake minimizing the functional in Equation 1 by solving the following Euler-Lagrange equation:

$$C_t = -\alpha \frac{\partial^2 C}{\partial s^2} + \beta \frac{\partial^4 C}{\partial s^4} + \nabla E_{img} \quad (2)$$

where C_t is the partial derivative of $C(s, t)$ with respect to the introduced time variable t , which tracks the evolution of the snake. Equilibrium is achieved when the internal force and image force balance and the left-hand side term C_t vanishes.

The *internal energy* can be written as:

$$E_{int} = \frac{1}{2} \left(\alpha \left| \frac{\partial C}{\partial s} \right|^2 + \beta \left| \frac{\partial^2 C}{\partial s^2} \right|^2 \right) \quad (3)$$

where α and β are blending parameters that control the snake tension and rigidity, respectively.

A commonly used *external* image force is gradient vector flow (GVF) [8, 9]. This is a bidirectional image force ∇E_{img} that can capture the object boundaries from either side and can deal with concave regions. A GVF field $\mathbf{v}(x, y)$ is defined as the equilibrium solution of the following system of partial equations:

$$\begin{aligned} \mathbf{v}_t &= g(|\nabla f|) \nabla^2 \mathbf{v} - h(|\nabla f|) (\mathbf{v} - \nabla f) \\ \mathbf{v}(x, y, 0) &= \nabla f \end{aligned} \quad (4)$$

where \mathbf{v}_t is the partial derivative of $\mathbf{v}(x, y, t)$ with respect to t , and ∇f is the gradient of the image edge map. The steady state of this update equation depends on the scalar parameter κ which affects the relative weight of the smoothness term $g(|\nabla f|) = \exp\{-\frac{|\nabla f|}{\kappa}\}$ and the term that fits the GVF to the image gradient $h(|\nabla f|) = 1 - g(|\nabla f|)$.

In section 3 we offer a definition of a new E_{con} term that is used to enforce constraints that are available from an understanding the intrinsic manifold structure of an image set. The next section gives a brief background on these automated manifold learning methods, and their recent specialization for medical imagery.

Manifold Learning. Image data can be naturally represented as points in a high dimensional data space (one dimension for each pixel). Often, however, a set of images has a lower *intrinsic* dimensionality, and the image data set can be mapped onto a lower dimensional space. Classical dimensionality reduction techniques for image sets rely on Principle Component Analysis (PCA) [10] and Independent Component Analysis (ICA) [11]. These seek to represent data as linear combinations of a small number of basis vectors. However, many natural image data sets have an intrinsic dimensionality that is much less than the number of basis images required to linearly reconstruct them.

This has led to a number of methods seeking to parameterize low-dimensional, non-linear manifolds. These methods measure local distances or approximate geodesic distances between points in the original data set, and seek low-dimensional embeddings that preserve these properties. Isomap [12] extends classic multidimensional scaling (MDS) by substituting an estimate of the geodesic distance along the image manifold for the inter-image Euclidean distance as input. LLE [13] attempts to represent the image manifold locally by reconstructing each image as weighted combination of its neighbors. SDE [14] applies semi-definite programming to learn kernel matrices which can be used to create isometric embeddings.

Isomap performs well for image sets sampled from convex manifolds. LLE and SDE do not fail in the case of non-convexity, but do not provide minimal parameterizations for cyclic manifolds (i.e., they give points on a sphere three coordinates instead of two). One algorithm which explicitly addresses cyclic manifolds is [15]. These algorithms, and others [16, 17, 18] have been used in various applications, including classification, recognition, tracking, and to a limited extent, biomedical image analysis [19]. Using image distance measures that explicitly reflect the variations within the image set (for instance, using estimates of the local deformation instead of pixel intensity differences) has been shown to be advantageous for medical imagery, and leads to low-dimensional embeddings that more accurately reflect the underlying intrinsic degrees of freedom [20, 21].

2 Cardio-Pulmonary Image Manifolds

Cardio-pulmonary MRI imagery is an outstanding candidate for manifold analysis. The appearance and deformations of MRI imagery of the chest varies greatly from patient to patient. However, images of a particular patient vary with the global deformation of the chest cavity due to breathing, deformation of the heart and nearby tissues due to heartbeats, and image noise.

The analysis of these images, even the *capture* of MR imagery, is affected by these parameters. Diagnostic imaging often requires “held-breath” and cardiac-gated imaging. These protocols offer two distinct methods to minimize variation in the captured imagery. The “held-breath” protocols require the patient to minimize variation due to breathing (by capturing all the imagery while the patient holds their breath and minimizes lung motion). Cardiac-gated imaging triggers the image capture to always occur at the same part of the cardiac cycle.

When the diagnostic value of imaging comes from analysis of motion, the imaging process takes unconstrained samples of the cardiopulmonary image manifold — a set of images that vary due to the heart and lung deformation. In this section we consider the post hoc analysis of samples of the cardiopulmonary image manifold. Although heartbeat monitors are relatively unobtrusive and common in diagnostic environments, we explore methods that operate only on the images without any additional knowledge, and focus on tools that will be effective in the presence of significant noise. To make this presentation concrete, we focus on a particular, noisy MRI image set, which is illustrated in Figures 1,4,5.

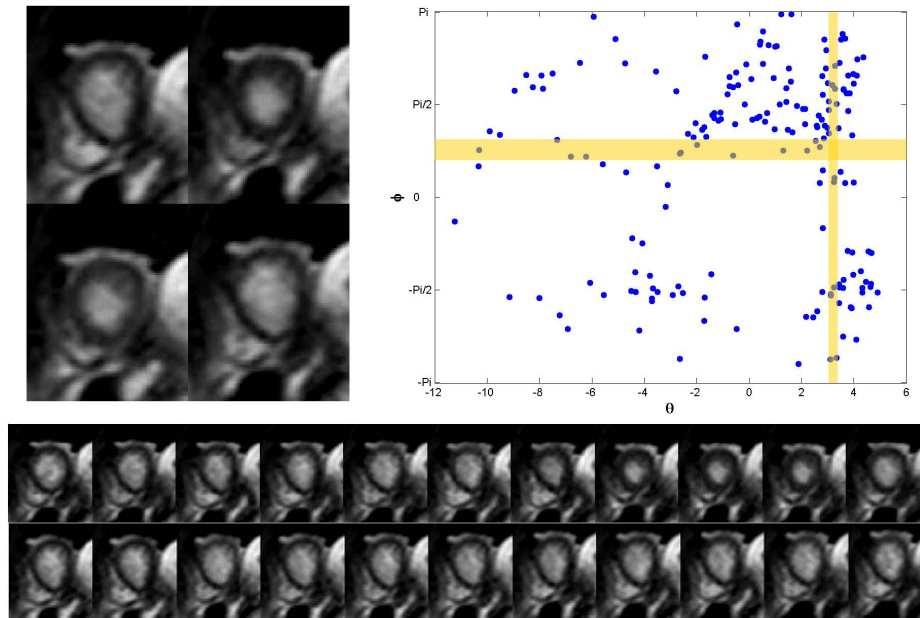


Fig. 1. The top left shows four images from a 200 image cardio-pulmonary MRI cine sequence of the heart. Note the variation both in the shape of the left ventricle (the white blob roughly centered in the image) and the position of the heart (shifting vertically). The top right shows the cylindrical Isomap embedding of this data set (using the algorithm of [15]). The scale of the axes is the same, and is proportional to the distance between images (measured, as discussed in the text, as the sum of the phase difference of a grid of complex Gabor filters). Cylindrical Isomap imposes a scale factor so that the y-axis (the cyclic axis) is embedded in the range $[-\pi, \pi]$. The bottom of the figure shows two sequences, example images from two paths through the Isomap embedding. The top sequence shows ordered image from a vertical path (drawn in yellow) at $\theta \approx 3.1$. Notice that in these images the heart deforms due to its beating, but there is no translation of the heart. The bottom sequences shows image for which $\phi \approx \pi/4$. Notice here that the heart is not deforming, but rather there is a translation (due to breathing).

2.1 Isomap Embedding of Cardiopulmonary Manifolds

The Isomap procedure for dimensionality reduction starts by computing the distance between all pairs of images (using some distance function such as SSD pixel intensities). Then, a graph is defined with each image as a node and undirected edges connecting each image to its k -closest neighbors (usually choosing k between 5 and 10). A complete pair-wise distance matrix is calculated by solving for the all-pairs shortest paths in this sparse graph. Finally, this complete distance matrix is embedded into low dimensions, by solving an Eigenvalue problem using a technique called Multidimensional Scaling (MDS) [22]. The low-dimensional embedding can be chosen as desired, but ideally is the number of degrees of freedom in the image set, in our case 2 (the two intrinsic dimensions of variability are the heartbeat and breathing).

One previous work that applied manifold learning to biomedical image analysis suggests modifying Isomap to use image distance functions other than pixel intensity differences [21]. For data sets with deformable motion, the suggested distance function is computed as the phase difference of local Gabor filters, where the filters have a reasonable magnitude:

$$\|I_1 - I_2\|_{motion} = \sum_{x,y} \Psi(G_{(\omega,V,\sigma)} \otimes I_1, G_{(\omega,V,\sigma)} \otimes I_2) + \Psi(G_{(\omega,H,\sigma)} \otimes I_1, G_{(\omega,H,\sigma)} \otimes I_2)$$

where $G_{(\omega,\{V|H\},\sigma)}$ is defined to be the 2D complex Gabor filter with frequency ω , oriented either vertically or horizontally, with σ as the variance of the modulating Gaussian, and Ψ returns the phase difference of the pair of complex Gabor responses above some threshold τ ; we choose τ to be the 50-th percentile filter magnitude. A technical modification to the Isomap procedure also allows the images to be embedded on a cylindrical manifold instead of a flat plane [15]. Figure 1 illustrates the cylindrical embedding of 185 frame cardiac MRI image set, captured each 72 ms. The figure also illustrates that the manifold embedding separates the non-rigid deformation of the heart from the translation of the heart due to breathing. In the next section, we consider how to exploit this manifold structure to provide new constraints for defining contours over every image of a data set undergoing such deformations.

3 Segmentation Constraints

The manifold embedding provides an automated tool to parameterize the cardiac image data, in terms of the motion caused by the heartbeat and the breathing. In this section we propose a method to solve for contours of the left ventricle in all images in the data set simultaneously. This solution uses two new constraints, first, a generic smoothness constraint that penalizes variation in contours that fit images nearby in the manifold, and second, a term that uses the specific nature of the image changes along different manifold directions to provide stronger constraints on manifold shape.

For a cardiopulmonary image sequence, heart deformation through time is defined by cardiac phase ϕ and pulmonary phase θ . Therefore, we seek to define

the shape of the heart contour as a function of ϕ and θ , described by $C(s, \phi, \theta) = (x(s, \phi, \theta), y(s, \phi, \theta))^T$. For a fixed ϕ, θ varying the arc-length parameter s traces out the contour boundary. Varying ϕ, θ defines all possible contours of the heart observed. A given cardiopulmonary image sequence specifies this contour by minimizing the following functional:

$$E = \oint E_{int}(C(s, \phi, \theta)) + E_{img}(C(s, \phi, \theta)) + E_{con}(C(s, \phi, \theta)) ds d\phi d\theta \quad (5)$$

The generic smoothness constraints (between nearby images on the manifold), can be written naturally as a parallel to the internal energy of the snake model, as follows:

$$E_{int} = \frac{1}{2} \left[\alpha \left| \frac{\partial C}{\partial s} \right|^2 + \beta \left| \frac{\partial^2 C}{\partial s^2} \right|^2 + \mu \left(\left| \frac{\partial C}{\partial \phi} \right|^2 + \left| \frac{\partial C}{\partial \theta} \right|^2 \right) + \gamma \left(\left| \frac{\partial^2 C}{\partial \phi^2} \right|^2 + \left| \frac{\partial^2 C}{\partial \theta^2} \right|^2 \right) \right] \quad (6)$$

where parameter μ and γ control the snake's tension and rigidity along ϕ and θ , respectively.

More specific constraints are available when the manifold dimensions correspond to specific kinds of motion. The breathing of the patient, while causing a complicated deformation of the chest cavity as a whole, results, largely, in a translation of the heart. The cardiac cycle, absent motion caused by breathing, causes deformation with minimal overall translation. Both of these types of motion allow stronger constraints on the relationship of a contour between frames than simple temporal continuity. Therefore, there are additional constraints between images that are embedded at either the same ϕ coordinate or the same θ coordinate. These constraints can be written as:

$$E_{con} = \frac{\eta}{2} \left| \frac{\partial C}{\partial \theta} - \int_0^1 \frac{\partial C}{\partial \theta} ds \right|^2 + \frac{\rho}{2} \left| \int_0^1 \frac{\partial C}{\partial \phi} ds \right|^2, \quad (7)$$

where the first term penalizes non-rigid changes in the snake (by integrating the squared difference between the motion of points on the contour and the mean motion of the contour), and the second term penalizes the overall mean translational motion of the snake, which is minimal when motion is caused only by the heartbeat. The rest of this section details our implementation of these constraints.

3.1 Implementation

Splines are a widely used function approximation tool [23]. A snake can be modeled by a closed cubic b-spline with N control points $\{p_i = (x_i, y_i)^T, i = 1 \dots N\}$, and a closed curve $C(s)$ as a collection of n curve segments $g_i(s), s \in [0, 1]$. Each curve segment is controlled by four nearby control points, as follows:

$$g_i(s) = \sum_{j=0}^3 b_j(s) p_{i+j-1}, i = 1, \dots, N \quad (8)$$

where $p_0 = p_N, p_{N+1} = p_1$ and $p_{N+2} = p_2$. Uniform cubic B-spline basis functions, $b_0 \sim b_3$, are defined by:

$$\begin{aligned}
 b_0(s) &= \frac{1}{6}(1-s)^3 \\
 b_1(s) &= \frac{1}{6}(3s^3 - 6s^2 + 4) \\
 b_2(s) &= \frac{1}{6}(-3s^3 + 3s^2 + 3s + 1) \\
 b_3(s) &= \frac{1}{6}s^3
 \end{aligned} \tag{9}$$

Then, the snake contour $C(s)$ is represented by multiplying a vector \mathbf{P} of snake control points with its associated b-spline basis functions:

$$\underbrace{\begin{bmatrix} g_1 \\ g_2 \\ \vdots \\ g_N \end{bmatrix}}_{\mathbf{C}(s)} = \underbrace{\begin{bmatrix} b_1 & b_2 & b_3 & \cdots & b_0 \\ b_0 & b_1 & b_2 & b_3 & \cdots \\ & b_0 & b_1 & b_2 & b_3 & \cdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \\ & \cdots & b_0 & b_1 & b_2 & b_3 \\ b_3 & \cdots & b_0 & b_1 & b_2 \\ b_2 & b_3 & \cdots & b_0 & b_1 \end{bmatrix}}_{\mathbf{H}(s)} \underbrace{\begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_N \end{bmatrix}}_{\mathbf{P}} \tag{10}$$

Using b-spline representation, one can analytically compute the snake’s k -th order of derivatives with respect to s as:

$$\frac{\partial^k \mathbf{C}(s)}{\partial s^k} = \frac{\partial^k \mathbf{H}(s)}{\partial s^k} \mathbf{P}. \tag{11}$$

To create a set of samples along each curve segment p_i , we can choose a set of parametric variables $\{s_j, j = 1, \dots, k\}$, such that $0 \leq s_1 < s_2 < \dots < s_k < 1$. Then the snake can be written in a discrete form:

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}(s_1) \\ \mathbf{C}(s_2) \\ \vdots \\ \mathbf{C}(s_k) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{H}(s_1) \\ \mathbf{H}(s_2) \\ \vdots \\ \mathbf{H}(s_k) \end{bmatrix}}_{\mathbf{H}} \mathbf{P} \tag{12}$$

Because it is our intention to solve simultaneously for snakes in every image, we parameterize the snake control points in each image as a function of position of each image on the manifold. Given an image sequence of M frames, we use the manifold learning procedure to estimate the breathing phase θ and cardiac phase ϕ for each image. For the i -th frame, let these values be (ϕ_i, θ_i) . Let

$\mathbf{C}(\phi_i, \theta_i)$ denote the snake defined by a discrete set of control points $\mathbf{P}(\phi_i, \theta_i) = [p_1(\phi_i, \theta_i), p_2(\phi_i, \theta_i), \dots, p_n(\phi_i, \theta_i)]^\top$. Then we can express the set of all snakes in all images as:

$$\underbrace{\begin{bmatrix} \mathbf{C}(\phi_1, \theta_1) \\ \mathbf{C}(\phi_2, \theta_2) \\ \vdots \\ \mathbf{C}(\phi_M, \theta_M) \end{bmatrix}}_{\mathbf{C}} = \underbrace{\begin{bmatrix} \mathbf{H} & & \\ & \mathbf{H} & \\ & & \ddots \\ & & & \mathbf{H} \end{bmatrix}}_{\mathbf{H}} \underbrace{\begin{bmatrix} \mathbf{P}(\phi_1, \theta_1) \\ \mathbf{P}(\phi_2, \theta_2) \\ \vdots \\ \mathbf{P}(\phi_M, \theta_M) \end{bmatrix}}_{\mathbf{P}} \quad (13)$$

For any control point $p_i(\phi, \theta)$, its position changes in different frames as a function of ϕ and θ . To model the relationship of the positions of control point p_i between different frames, we use a cubic B-spline surface to represent the position of each point p_i . Any point on that surface presents the position of control point p_i in the frame specified by (ϕ, θ) . Therefore, the change of the control point position should be locally small and continuous.

The cubic b-spline surface for the i -th snake control point is defined by a two-dimensional set of control points $\{q_{u,v}^{(i)}, u = 1, \dots, n; v = 1, \dots, m\}$. The following is the equation of a cubic b-spline surface defined by n rows and m columns of surface control points:

$$\begin{aligned} p_i(\phi, \theta) &= \sum_{u=1}^n \sum_{v=1}^m B_u(\phi) B_v(\theta) q_{u,v}^{(i)} \\ &= \underbrace{\begin{bmatrix} B_1(\phi) B_1(\theta) \\ B_2(\phi) B_1(\theta) \\ \vdots \\ B_n(\phi) B_1(\theta) \\ B_1(\phi) B_2(\theta) \\ \vdots \\ B_n(\phi) B_m(\theta) \end{bmatrix}^\top}_{\mathbf{B}(\phi, \theta)} \underbrace{\begin{bmatrix} q_{1,1}^{(i)} \\ q_{2,1}^{(i)} \\ \vdots \\ q_{m,1}^{(i)} \\ q_{1,2}^{(i)} \\ \vdots \\ q_{m,n}^{(i)} \end{bmatrix}}_{\mathbf{Q}_i} \end{aligned} \quad (14)$$

where $B_i(\phi)$ and $B_j(\theta)$ are cubic b-spline basis functions. Then we can analytically compute the k -th order of derivatives of control point p_i with respect to ϕ (or θ) as

$$\frac{\partial^k p_i(\phi, \theta)}{\partial \phi^k} = \frac{\partial^k \mathbf{B}(\phi, \theta)}{\partial \phi^k} \mathbf{Q}_i, \quad \frac{\partial^k p_i(\phi, \theta)}{\partial \theta^k} = \frac{\partial^k \mathbf{B}(\phi, \theta)}{\partial \theta^k} \mathbf{Q}_i \quad (15)$$

Considering all control points over all frames, we have

$$\mathbb{P} = \begin{bmatrix} \mathbf{P}(\phi_1, \theta_1) \\ \mathbf{P}(\phi_2, \theta_2) \\ \vdots \\ \mathbf{P}(\phi_n, \theta_n) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{B}(\phi_1, \theta_1) & & & \\ & \mathbf{B}(\phi_1, \theta_1) & & \\ & & \ddots & \mathbf{B}(\phi_1, \theta_1) \\ \vdots & & \vdots & \\ \mathbf{B}(\phi_m, \theta_m) & & & \\ & \mathbf{B}(\phi_m, \theta_m) & & \\ & & \ddots & \mathbf{B}(\phi_m, \theta_m) \end{bmatrix}}_{\mathbb{B}} \underbrace{\begin{bmatrix} \mathbf{Q}_1 \\ \mathbf{Q}_2 \\ \vdots \\ \mathbf{Q}_m \end{bmatrix}}_{\mathbb{Q}} \quad (16)$$

Finally, we have a single form which expresses the contours in all images:

$$\mathbb{C} = \mathbb{H}\mathbb{P} = \mathbb{H}\mathbb{B}\mathbb{Q}. \quad (17)$$

Using the calculus of variations, the snake minimizing the functional in (5) can be found by solving for the following Euler-Lagrange equation:

$$\begin{aligned} \frac{\partial E}{\partial C} &= \frac{\partial E_{int}}{\partial C} + \nabla E_{img} + \frac{\partial E_{con}}{\partial C} = 0 \\ \Rightarrow -\alpha \frac{\partial^2 C}{\partial s^2} + \beta \frac{\partial^4 C}{\partial s^4} - \mu \left(\frac{\partial^2 C}{\partial \phi^2} + \frac{\partial^2 C}{\partial \theta^2} \right) + \gamma \left(\frac{\partial^4 C}{\partial \phi^4} + \frac{\partial^4 C}{\partial \theta^4} \right) \\ &+ \nabla E_{img} - \rho \left(\frac{\partial^2 C}{\partial \theta^2} - \int_0^1 \frac{\partial^2 C}{\partial \theta^2} ds \right) - \eta \int_0^1 \frac{\partial^2 C}{\partial \phi^2} ds = 0 \\ \Rightarrow \mathbf{A}_{int}\mathbb{C} + \mathbf{A}_{con}\mathbb{C} - \mathbf{V} &= 0 \end{aligned} \quad (18)$$

where \mathbf{A}_{int} and \mathbf{A}_{con} are matrices corresponding to internal energy and external constraint energy term, respectively, and they can be directly computed using equations in (17), (15), and (11). \mathbf{V} is the matrix presenting the collection of GVF \mathbf{v} sampled along $C(s, \phi, \theta)$ over all images. Since the snake is implemented using cubic b-spline, all 4-th order of partial derivatives of C with respect to s , ϕ and θ are zero. Hence, in the later sections, we will ignore the blending parameters β and γ in the equations in (18).

In order to obtain desired solution of the Euler equation (18), the snake $C(s, \phi, \theta)$ is treated as a function evolves with respect to time variable t , and the resulting equation is

$$\mathbf{A}_{int}\mathbb{C}^{t+1} - \mathbf{V} + \mathbf{A}_{con}\mathbb{C}^{t+1} = -\delta(\mathbb{C}^{t+1} - \mathbb{C}^t) \quad (19)$$

where δ denotes a step size. At equilibrium, the time derivative vanishes and we end up from equation (19) and (17) the following update rules as:

$$\begin{aligned} \mathbb{C}^{t+1} &= (\mathbf{A}_{int} + \mathbf{A}_{con} + \delta\mathbf{I})^{-1}(\delta\mathbb{C}^t + \mathbf{V}) \\ \Rightarrow \mathbb{Q}^{t+1} &= [(\mathbf{A}_{int} + \mathbf{A}_{con} + \delta\mathbf{I})\mathbb{H}\mathbb{B}]^+ (\delta\mathbb{H}\mathbb{B}\mathbb{Q}^t + \mathbf{V}) \end{aligned} \quad (20)$$

where \mathbf{I} denotes identity matrix.

4 Experimental Results

This section describes preliminary results of a system that implements the constraints defined in the last section. We first consider a collection of artificially generated data, for which we can control the deformation and noise parameters, and test our approach versus standard snake approaches. We follow this with an application to finding the left ventricle wall shape in a noise cardiac MRI sequence. In both cases, we used $\alpha = 0.1$ for all snakes, and $\mu = 0.1$, $\rho = 0.1$ and $\eta = 0.01$ for manifold constraints. Deriving optimal choices for these parameters, or other methods to automate the process of finding good parameters is in further investigation.

4.1 Simulation Data

We construct an artificial data set by defining a shape and deforming it with a composition of a non-rigid deformation and a rigid translation. Thus, this data set has a 2D manifold structure, indexed by the magnitude of each deformation. One hundred images were created, each was then corrupted by additive white Gaussian noise, and the contrast was decreased in a randomly selected patch of the image. The two deformation types are depicted in the top of Figure 2, and 8 selected frames among the 100 generated images are shown at the top right.

The noise in the image and the low contrast patches make this a challenging data set for snakes to converge to the correct boundary. The second column of Figure 2 gives results for classical snakes (using only a single image) with the starting condition shown in the first column. While some optimization is possible to improve these results for this data set, for this work it is our goal to illustrate the advantages of using the manifold structure of these images.

The third column of Figure 2 gives the contours which are the results of applying the algorithm of the previous section, which exploits the manifold structure of these images. These can be compared to the ground truth contours shown in the rightmost column.

4.2 Example Application

In cardiovascular imagery, an important application is to measure the dynamic behavior of the human heart, especially the left ventricle. Figure 3 illustrates one example of coupled snakes [24, 25] that outline the interior and exterior wall of the left ventricle. Often, For continuously captured cine-MRI images, the available resolution and contrast is more limited than this example shows, and extracting ventricle wall contours on each image individually is difficult. By imposing manifold constraints, we solved the problem with a modified version of [25].

Figure 4 shows examples of cine-MRI images (from the same data set as shown in Figure 1). The rectangular region of the heart is blown up and the results of fitting pairs of snakes to the inner and outer ventricle wall on individual images is shown in the middle column. On the right of this figure are




Fig. 2. An artificial data set constructed by composing 2 deformations (shown in the top row), a non-rigid variation to a shape and a rigid translation. The bottom four rows show results fitting snakes to these images which are corrupted by zero mean Gaussian noise and by a patch of reduced contrast. Column (a) shows the image and the initial condition, column (b) shows the classical snake result on this image, column (c) shows the result (for that image) of using the manifold based constraints to solve for all snakes simultaneously, and (d) is the ground truth contour used to generate each image.

the results when enforcing the additional constraints from understanding the manifold image structure, following the algorithm outlined in Section 3. While these results are not perfect, they are an improvement and encourage further




Fig. 3. Important applications in cardiac imagery include measuring the thickness of ventricle walls. For relatively high contrast, and high resolution images, a pair of snakes can find both the inner and outer wall. Applying this algorithm to lower resolution images (in Figure 4) fails because of insufficient image resolution and contrast.




Fig. 4. For continuously captured cine-MRI images, there are currently limits on the available resolution and contrast. Extracting ventricle wall contours on each image individually is difficult. The original image is shown on the left, and a subwindow shows an expansion of the heart region, and the initial snake contour (for the interior wall). The contours extracted from a single image are shown in the middle. On the right are the results when enforcing the additional shape constraints from the manifold structure extracted in Figure 1.




Fig. 5. Additional segmentation results obtained by imposing our manifold constraints over the whole image sequence. Here we show 8 consecutive frames of the 200 frame sequence. Notice that in these cardio-plumonary images, the heart deformation is sometimes large between consecutive images due to the relatively low sampling rate. In these cases, naive temporal smoothness constraints on heart deformation may fail.

work in integrating manifold constraints with energy minimization tools for contour fitting. Additional results are shown in Figure 5, for 8 consecutive images of the cine sequence. The frame to frame deformation is in some cases quite large, so naive smoothness constraints between consecutive images may not be successful.

5 Conclusion

This work presents preliminary efforts towards incorporating manifold learning as a tool to provide additional constraints for finding contours using energy minimization tools. The advantages of combining these two techniques were illustrated for a collection of simulation data and demonstrated on a low-resolution, high-noise cardiac MRI video sequence.

It is also possible to extend the manifold constraints to a level set based segmentation framework for handling topological changes and numerical stability of evolving curves [26]. We believe that many algorithms may be improved through better understanding and exploitation of non-linear image manifold learning algorithms, and tight integration of these with classical analysis tools.

Acknowledgements

This research was supported by NSF grant IIS-0413291.

References

1. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *International Journal of Computer Vision* **1** (1988) 321–331
2. Cremers, D., Tischhäuser, F., Weickert, J., Schnörr, C.: Diffusion snakes: Introducing statistical shape knowledge into the mumford-shah functional. *International Journal of Computer Vision* **50** (2002) 295–313
3. Duncan, J.C., Ayache, N.: Medical image analysis: Progress over two decades and the challenges ahead. *IEEE Trans. Pattern Anal. Mach. Intell.* **22** (2000) 85–106
4. Frangi, A.F., Rueckert, D., Duncan, J.S.: Three-dimensional cardiovascular image analysis. *IEEE Trans. Med. Imaging* **21** (2002) 1005–1010
5. McInerney, T., Terzopoulos, D.: Deformable models in medical images analysis: a survey. *Medical Image Analysis* **1** (1996) 91–108
6. Ueda, N., Mase, K.: Tracking moving contours using energy-minimizing elastic contour models. *International Journal of Pattern Recognition and Artificial Intelligence* **9** (1995) 465–484
7. Geiger, D., Gupta, A., Costa, L.A., Vlontzos, J.: Dynamic programming for detecting, tracking, and matching deformable contours. *IEEE Trans. Pattern Anal. Mach. Intell.* **17** (1995) 294–302
8. Xu, C., Prince, J.L.: Gradient vector flow: A new external force for snakes. In: *Proceedings of the Conference on Computer Vision and Pattern Recognition*, Washington, DC, USA (1997) 66
9. Xu, C., Prince, J.L.: Generalized gradient vector flow external forces for active contours. *Signal Process.* **71** (1998) 131–139
10. Jolliffe, I.T.: *Principal Component Analysis*. Springer-Verlag (1986)
11. Hyvärinen, A., Karhunen, J., Oja, E.: *Independent Component Analysis*. John Wiley and Sons (2001)
12. Tenenbaum, J.B., de Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *Science* **290** (2000) 2319–2323
13. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290** (2000) 2323–2326
14. Weinberger, K.Q., Saul, L.K.: Unsupervised learning of image manifolds by semidefinite programming. In: *Computer Vision and Pattern Recognition*. (2004)
15. Pless, R., Simon, I.: Embedding images in non-flat spaces. In: *Proc. of the International Conference on Imaging Science, Systems, and Technology*. (2002)
16. Donoho, D.L., Grimes, C.: Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *PNAS* **100** (2003) 5591–5596
17. Brand, M.: Charting a manifold. In S. Becker, S.T., Obermayer, K., eds.: *Advances in Neural Information Processing Systems 15*. MIT Press, Cambridge, MA (2003) 961–968
18. Belkin, M., Niyogi, P.: Laplacian eigenmaps and spectral techniques for embedding and clustering. In: *Advances in Neural Information Processing Systems*. (2002)
19. Lim, I.S., Ciechomski, P.d.H., Sarni, S., Thalmann, D.: Planar arrangement of high-dimensional biomedical data sets by isomap coordinates. In: *Proceedings of the 16th IEEE Symposium on Computer-Based Medical Systems*. (2003) 50–55
20. Pless, R.: Differential structure in non-linear image embedding functions. In: *Articulated and Nonrigid Motion*. (2004)
21. Souvenir, R., Pless, R.: Isomap and non-parametric models of image deformation. In: *In Proc. IEEE Workshop on Motion and Video Computing*, Breckenridge, CO (2005)

22. Borg, I., Groenen, P.: *Modern Multidimensional Scaling: Theory and Applications*. Springer-Verlag (1997)
23. Dierckx, P.: *Curve and surface fitting with splines*. Oxford University Press, Inc., New York, NY, USA (1993)
24. Paragios, N.: A variational approach for the segmentation of the left ventricle in cardiac image analysis. *International Journal of Computer Vision* **50** (2002) 345–362
25. Malpicaa, N., Ledesma-Carbayoa, M., Santosa, A., Prezb, E., Garca-Fernandezb, M., Descob, M.: A coupled active contour model for myocardial tracking in contrast echocardiography. In: *Image Understanding and Analysis*, Imperial College London, UK (2004)
26. Zhang, Q., Pless, R.: Segmenting cardiopulmonary images using manifold learning with level sets. In: *ICCV workshop on Computer Vision for Biomedical Image Applications: Current Techniques and Future Trends*, Beijing, China (2005)