Shape-Based Nonrigid Correspondence with Application to Heart Motion Analysis

Hemant D. Tagare, *Member, IEEE*

***Abstract*— A common problem in many biomedical imaging studies is that of finding a correspondence between two plane curves which aligns their shapes. A mathematical formulation and solutions to this problem is proposed in this paper. The formulation exhibits desirable properties. It allows for one-to-one as well as non-one-to-one correspondences, it consistently compares shape, even in nonrigid situations, and it is completely symmetric with respect to the two curves.**

A numerical implementation of the algorithm for finding the optimal correspondence is also reported. The algorithm is used to estimate nonrigid motion of the endocardium in MRI image sequences of normal and post-infarct dog hearts. The return error (the difference between the starting and ending positions of a point) is used as a performance measure to evaluate the technique. Since heart motion is periodic, the return error is a measure of consistency of the algorithm.

Preliminary applications to other data sets are reported as well.

I. Introduction

FINDING a correspondence (i.e., a pairing of points) between curves is a crucial step in analyzing many medical images. Given two closed plane curves, say C_1 and C_2, we have to pair points of C_1 with points of C_2 in such a way that the following holds.

1) The pairing is continuous. Continuity means nearby points of one curve are paired with nearby points of the other.
2) The pairing is nonrigid. That is, points of one curve which are separated by a certain distance may be paired with points separated by a different distance in the other curve. The phrase nonrigid correspondence really means not-necessarily-rigid correspondence. Rigid correspondences are a subset of the set of nonrigid correspondences.
3) The pairing aligns local shape. Loosely speaking, this means that strongly curved regions of one curve are paired with strongly curved regions of the other and flat regions are paired with flat regions.

In addition, the correspondence (pairing) is often required to be one-to-one. However, this requirement is not strict. There are cases where a non-one-to-one correspondence may be appropriate.

Manuscript received November 25, 1998; revised March 30, 1999. This work was supported in part by the National Library of Medicine under Grant 1-R01-LM0500. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was B. Vemuri.

The author is with the Department of Diagnostic Radiology and the Department of Electrical Engineering, Yale University, New Haven, CT 06520 USA.

Publisher Item Identifier S 0278-0062(99)07425-X.

There are many applications of shape-based correspondences [18]. For example, shape-based nonrigid correspondence is used for tracking two-dimensional (2-D) heart motion [3], [9], [13], [14], [24], [25], in neural growth studies [16], in morphometrics [5], for the analysis of brain structure [10], [23], and in chromosome analysis [19].

The aim of this paper is to propose a theory and an algorithm for shape-based nonrigid correspondences. A brief outline of the paper is as follows. First, a useful set of correspondences between two curves is identified. Next, an objective function is created to measure how well any correspondence in the set aligns the shapes of the two curves. Finally, a numerical procedure is proposed to search through the set to find the best correspondence according to the objective function.

The correspondences, called bimorphisms, are defined in Section II. The objective function for comparing local shape (i.e., local bending) of the two curves is explained in Section III. The numerical algorithm is explained in Section IV. Section V contains experimental results where the algorithm is used to analyze heart motion and Section VI concludes the paper.

The treatment in this paper is expository. All proofs are omitted. The reader is referred to [31] for proofs and generalizations of the main mathematical results.

A. Previous Work

Shape analysis and comparison is a vast field and reviewing it exhaustively is beyond the scope of this paper. Excellent reviews are available in [21], [22], and [26]. Relevant subsets of shape analysis algorithms are those that work by establishing a correspondence between curves. Many of these algorithms work in this way. Let s_1 be the arc-length parameter of curve C_1, s_2 be the arc-length parameter of curve C_2, and ϕ be a nondecreasing function mapping s_1 to s_2. Of all feasible ϕ's, the one which minimizes the following objective function is chosen as the best correspondence between C_1 and C_2:

$$\int_{C_1} \{\kappa_1(s_1) - \kappa_2(\phi(s_1))\}^2 \, ds_1 \tag{1}$$

where, $\kappa_1(\,)$ and $\kappa_2(\,)$ are the curvature arc-length functions of the two curves.

Although this formulation captures the intuition behind shape-based nonrigid correspondence, it has two serious drawbacks.

1) The objective function of (1) is not symmetric with respect to C_1 and C_2. Since the objective function is

an integral over C_1, changing the labels of the two curves causes the domain of integration (and hence the value of the integral) to change. A correspondence which is optimal with respect to the first objective function will not necessarily remain optimal after the change of domain. To see this, consider the following calculation which changes the variable of integration from s_1 to s_2:

$$\int_{C_1} \{\kappa_1(s_1) - \kappa_2(\phi(s_1))\}^2 \, ds_1 = \int_{C_2} \{\kappa_1(\phi^{-1}(s_2)) - \kappa_2(s_2)\}^2 \, \frac{d\phi^{-1}(s_2)}{ds_2} \, ds_2.$$

This shows that the minimizing correspondence $\chi = \phi^{-1}$ can be found from C_2 if the energy function $\int_{C_2} w(s)\{\kappa_1(\chi(s_2) - \kappa_2(s_2)\}^2 \, ds_2$ is minimized with respect to χ where the weight $w(s)$ is set equal to $d\chi(s_2)/ds_2$. Because the weight $w(s)$ is not necessarily uniform, this minimum need not occur at the same χ for which $\int_{C_2} \{\kappa_1(\chi(s_2) - \kappa_2(s_2)\}^2 \, ds_2$ is minimized.

2) The use of the numerical value of curvature in a nonrigid situation is problematic since the numerical value of curvature is only a rigid invariant.

The technique proposed in this paper eliminates these drawbacks.

Also related to this paper, but different, is the development in morphometrics of the statistics of shape [5]. An excellent exposition of this is available in a book by Small [29]. Developed independently by Bookstein and Kendall, shape statistics are gathered from point landmarks and are not curve based. However, there is a recent attempt to generalize landmark-based shape statistics to curves [6].

The specific application to heart motion analysis is important enough that many approaches have been proposed. Some are specific to MR and use MR physics to create tags which can be tracked [2], [33]. Heart motion analysis has also been attempted with ultrasound (e.g., [7], [9], and [12]) and nuclear images (e.g., [27]).

The importance of shape in tracking heart motion is shown in [28], where implanted markers are shown to follow trajectories which are consistent with tracking shape. This is the motivation for applying shape-based correspondence algorithms to heart motion analysis.

We now turn to the discussion of correspondences between curves.

II. Correspondences and Bimorphisms

Our aim in this section is to define continuous nonrigid correspondences which may be one-to-one or non-one-to-one. Fig. 1(a)–(d) illustrates what we have in mind. Fig. 1(a) is a one-to-one correspondence and Fig. 1(b)–(d) is non-one-to-one correspondences which shrink segments to points or expand points to segments. Their utility is that they can be used to model the emergence and disappearance of local shape features [as is shown in Fig. 1(b)–(d)]. To define such correspondences, we will first define a (general) correspondence and then add constraints to the definition in such a way that

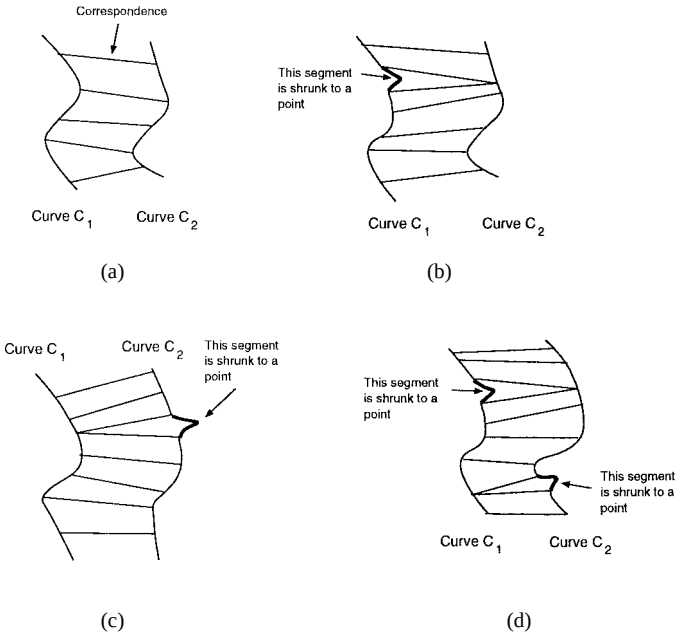

Fig. 1. Examples of bimorphisms. (a) One-to-one correspondence. (b) A non-one-to-one correspondence which shrinks a segment of C_1 down to a point. (c) A non-one-to-one correspondence which shrinks a segment of C_2 down to a point. (d) A non-one-to-one correspondence which shrinks segments of both curves to a point.

the constrained correspondence captures exactly the intuition of Fig. 1(a)–(d). The definitions are taken from [31].

In set theoretic terms, a correspondence between two curves C_1 and C_2 is simply a set $\{(u, v)\}$ of pairs of points, with the requirement that point u belong to C_1 and point v belong to C_2 and that every point of C_1 and C_2 appear in some pair. A more precise definition requires the notion of the product space of two curves. The product space of two closed curves C_1 and C_2 is the set $C_1 \times C_2$ of all ordered pairs (x, y) such that x belongs to C_1 and y to C_2. Elementary topology (e.g., [1]) informs us that the product space $C_1 \times C_2$ of two closed curves is a torus.

The product space $C_1 \times C_2$ has two natural projections. These are functions $p_1 \colon C_1 \times C_2 \to C_1$ and $p_2 \colon C_1 \times C_2 \to C_2$, which project an element of $C_1 \times C_2$ onto its first and second component, $p_1((x, y)) = x$ and $p_2((x, y)) = y$.

Pairing an element x of C_1 with an element y of C_2 is the same as choosing the element (x, y) of $C_1 \times C_2$. Making a set of such pairings is the same as choosing a subset of the product space. Thus, we obtain the following definition.

Definition: A correspondence Φ between closed curves C_1 and C_2 is a subset of $C_1 \times C_2$ whose projections on C_1 and C_2 are onto. The image of a point x under a correspondence is the set of points that x is paired with.

We can now proceed to define correspondences of the type shown in Fig. 1. Notice that the correspondence in Fig. 1(d) cannot be expressed either by a function from C_1 to C_2 or by a function from C_2 to C_1 (since a function can only produce a point as an output). However, locally, the correspondence looks like a function from one of the curves to the other. That is, near the segment of C_1, which is shrunk to a point in C_2, the correspondence looks like a function from C_1 to C_2. On

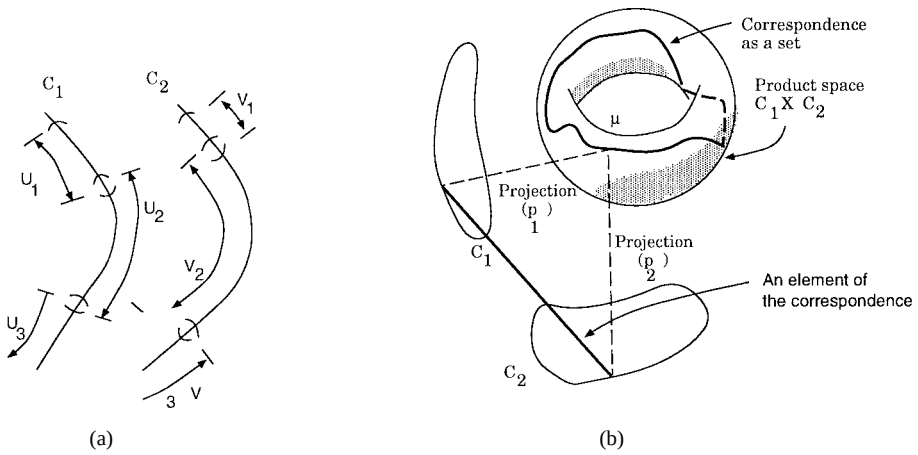

Fig. 2. Structure of a bimorphism.

the other hand, near the segment of C_2 which is shrunk to a point in C_1, the correspondence looks like a function from C_2 to C_1. With this observation, it should be intuitively clear that we can define all of the correspondences in Fig. 1 by letting the correspondence locally look like a function, but allowing the local functions to go in either direction. Furthermore, we can restrict the functions so that they can only shrink segments to points.

To convert this intuition into a definition, we cover C_1 and C_2 with pairs of open segments U_α and V_α [Fig. 2(a)] and require that the correspondence restricted to the subset $U_\alpha \times V_\alpha$ look like a function from one of the curves to the other.

Definition: Let α belong to some index set and let U_α and V_α be open segments covering C_1 and C_2 [Fig. 2(a)]. Further, let Φ be a correspondence between C_1 and C_2. The correspondence is a bimorphism if it has the following properties.

1) The image of any point under the correspondence is connected.
2) The image of U_α is V_α and the image of V_α is U_α.
3) For each α, the correspondence Φ restricted to $U_\alpha \times V_\alpha$ can be written either as $\{(x, \phi_\alpha(x))\}$ for some differentiable function ϕ_α: $U_\alpha \to V_\alpha$, or as $\{(\phi_\alpha(y), y)\}$ for some differentiable function ϕ_α: $V_\alpha \to U_\alpha$.

Here is what the definition means. The first property states that the image of any point under Φ must be connected. That is, the image of a point must be either a point or a connected segment. Nothing else is legal. The second and third properties state how correspondences of individual points are related to each other. The second property requires that the image of all points in U_α be the set V_α and vice versa. This means that it is possible to analyze the correspondence in $U_\alpha \times V_\alpha$ without worrying about the rest. The last property states that within $U_\alpha \times V_\alpha$, the correspondence is given by a differentiable function from $U_\alpha \to V_\alpha$ or from $V_\alpha \to U_\alpha$.

The definition certainly includes one-to-one correspondences. To see this, suppose that f: $C_1 \to C_2$ is a one-to-one onto differentiable function. Then $\{(x, f(x)) | x \in C_1\}$ is the correspondence due to f. First, notice that the image of any point under the correspondence is a point (because f is one-to-one). Now set the index set of α to $\{1\}$ and let $U_1 = C_1$, $V_1 = C_2$, and $\phi_1 = f$. This shows that the correspondence fits the definition of a bimorphism. In the same way, any one-to-one onto differentiable function from C_2 to C_1 is also a bimorphism.

So far we have not considered whether a bimorphism preserves the orientation of the two curves or not. That is, we have not considered how the image of a point behaves under the correspondence as we move the point along the orientation of one of the curves. It is easy to show that bimorphisms are of two types: one type of bimorphisms preserve the orientation and the other reverses the orientation. For applications, the first type is useful.

A. How to Create a Bimorphism

For algorithmic purposes we need a technique for numerically creating bimorphisms. It might appear from the definition of a bimorphism that this is hard, since it requires us to find a correspondence Φ and open covers U_α and V_α of a specific type. Fortunately, there is way around this. There are two simpler recipes for constructing bimorphisms and neither of them requires us to find open covers. The first recipe depends on the following result [31].

Proposition 1a: Φ is a bimorphism between C_1 and C_2 if and only if Φ is a regular curve in $C_1 \times C_2$ such that its projections on C_1 and C_2 are the onto and the inverse projection of any point of C_1 and C_2 to Φ is connected.

Fig. 2(b) illustrates the proposition. The product space of two curves is shown as a torus and a bimorphism is shown as a regular curve[1] on the torus. The proposition states that the curve is constrained in a way that the inverse projection of any point in C_1 or C_2 is a connected set. This simply means that the inverse projection of a point is either a point or a segment on the curve. Fig. 3 illustrates this. Proposition 1 implies that at any point, the curve can move horizontally or vertically or with some finite slope as shown in Fig. 3(a). But the curve cannot double back, as shown in Fig. 3(b). When the curve is horizontal, it pairs a connected segment of U_α with V_α. When the curve is vertical, it pairs a connected segment of V_α with

[1] A curve is regular if it has a nonzero tangent vector at every point.

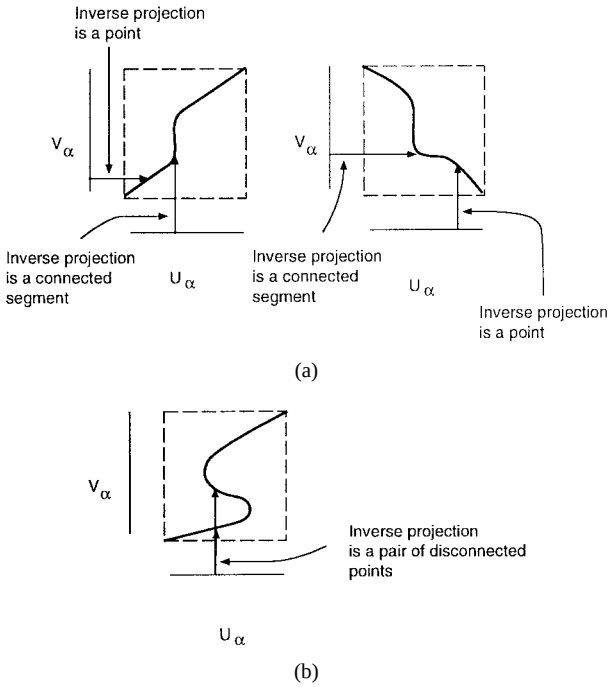

Fig. 3. Local structure of a bimorphism.

U_α. The entire bimorphism is a curve that winds around the torus without doubling back on itself.

Proposition 1a is valid for orientation-preserving, as well as orientation-reversing bimorphisms. It can be simplified if we restrict our attention to orientation-preserving bimorphisms:

Proposition 1b: Φ is an orientation-preserving bimorphism between C_1 and C_2 if and only if Φ is a regular curve in $C_1 \times C_2$ such that its projections on C_1 and C_2 are onto and are nondecreasing with respect to arc length.

By "its projections [are] nondecreasing with respect to arc length" we mean that as a point moves along the bimorphism, the arc lengths of its projections on C_1 and C_2 are nondecreasing.

Proposition 1b gives the first recipe for constructing an orientation-preserving bimorphism. We find a regular curve in $C_1 \times C_2$, keeping in mind that its projection on C_1 and C_2 must include every point of both curves and must be nondecreasing with respect to the arc lengths of both curves. Proposition 1b guarantees that the set of all such curves is identical to the set of orientation-preserving bimorphisms between C_1 and C_2.

To simplify this recipe further we give the bimorphism an arc length and express all constraints on the bimorphism in terms of the arc-length parameter. Since the bimorphism is a regular curve, it has a regular parameterization $\mu(t)$, $\mu : [0,1] \rightarrow C_1 \times C_2$ such that projections $p_1(\mu(t))$ and $p_2(\mu(t))$ are nondecreasing with respect to arc lengths of C_1 and C_2. Therefore

$$\frac{ds_1}{dt} = \|p_1(\mu'(t))\|, \quad \text{and} \quad \frac{ds_2}{dt} = \|p_2(\mu'(t))\|$$

where s_1 and s_2 are the arc-length parameters of C_1 and C_2, respectively, and $\|\ \|$ is the Euclidean norm. Note that $p_1(\mu'(t))$ and $p_2(\mu'(t))$ cannot be simultaneously zero because bimorphisms are regular curves. The arc length ds of the bimorphism between the points $\mu(t)$ and $\mu(t+dt)$ is defined as

$$ds = \sqrt{\left(\frac{1}{L_{C_1}}\frac{ds_1}{dt}\right)^2 + \left(\frac{1}{L_{C_2}}\frac{ds_2}{dt}\right)^2}\, dt \tag{2}$$

where, L_{C_1}, L_{C_2} are the total arc lengths of C_1 and C_2, respectively. The reason for dividing by L_{C_1} and L_{C_2} will become clear shortly. The condition that the bimorphism is a regular curve, simply translates into the existence of the arc-length parameterization of the bimorphism by the variable s defined above.

Furthermore, the condition that the projections of the bimorphism on C_1 and C_2 are nondecreasing with respect to arc length becomes

$$\frac{ds_1}{ds} \geq 0, \quad \text{and} \quad \frac{ds_2}{ds} \geq 0. \tag{3}$$

We now have a second recipe for creating a bimorphism. Create any regular curve in $C_1 \times C_2$ and arc-length parameterize it with the variable s, defined in (2). If the projections of the curve on C_1 and C_2 are onto and inequalities (3) hold, then the curve is a bimorphism. The set of curves generated this way is identical to the set of orientation-preserving bimorphisms between C_1 and C_2.

B. Uniform Mapping Bimorphisms

A special kind of bimorphism is particularly useful in the numerical algorithm and we will pause here briefly to define it. A function $f\colon [0, L_{C_1}] \rightarrow [0, L_{C_2}]$ *uniformly maps* the arc length of C_1 to the arc length of C_2 if

$$f(s_1) = \begin{cases} \dfrac{L_{C_2}}{L_{C_1}}\, s_1 + \delta, & \text{if } \dfrac{L_{C_1}}{L_{C_2}}\, s_1 + \delta \leq L_{C_2} \\ \dfrac{L_{C_2}}{L_{C_1}}\, s_1 + \delta - L_{C_2}, & \text{otherwise} \end{cases}$$

for any $L_{C_2} > \delta \geq 0$. That is, f simply maps the arc length of C_1 linearly to the arc length of C_2 in such a way that $s_1 = 0$ is mapped to $s_2 = \delta$.

The correspondence $\{C_1(s_1), C_2(f(s_2))\}$, given by the uniform map f, is called the uniform mapping bimorphism. These correspondences are used as initial conditions in the gradient descent part of the numerical algorithm in Section IV.

It is easy to show that, for a uniform mapping correspondence, ds_1/ds and ds_2/ds are constant and equal to $L_{C_1}/\sqrt{2}$ and $L_{C_2}/\sqrt{2}$, respectively.

III. Comparing Local Structure with a Correspondence

Having characterized the kind of correspondences we want, we now turn to the issue of creating an objective function which uses a bimorphism for comparing the local structure of the curves. Recall that the aim of the objective function is to indicate how desirable a correspondence is by comparing the local shape and local stretching of the curves.

A. Shape Comparison

Let $\Theta(s)$ be the angular orientation of the normal to a curve at arc length s. Then, the curvature of the curve at s is $\kappa(s) = (d\Theta'(s)/ds)$. If $\mu(s)$ is an element of a bimorphism Φ at arc length s along Φ and $p_1(\mu(s))$ and $p_2(\mu(s))$ are its projections on C_1 and C_2, then the angular orientations of the normals at these points are $\Theta_1(p_1(\mu(s)))$ and $\Theta_2(p_2(\mu(s)))$. Since a bimorphism is a curve, as μ moves along it the derivatives of $\Theta_1(p_1(\mu(s)))$ and $\Theta_2(p_2(\mu(s)))$, with respect to the arc length of Φ, give the local shapes of C_1 and C_2 as viewed from the bimorphism. The difference in the derivatives expresses the difference in local shapes.

The difference in the derivatives is

$$\frac{d\Theta_1(p_1(\mu(s)))}{ds} - \frac{d\Theta_2(p_2(\mu(s)))}{ds} = p_1'(\mu(s))\kappa_1(p_1(\mu(s))) - p_2'(\mu(s))\kappa_2(p_2(\mu(s)))$$

where s is the arc length along the bimorphism and κ_1 and κ_2 are the curvature arc-length functions of C_1 and C_2, respectively.

Integrating the square of the above quantity measures the overall dissimilarity of the shapes of C_1 and C_2, as viewed from Φ. Denoting this by $J(C_1, C_2; \Phi)$, we have

$$J(C_1, C_2; \Phi) \tag{4}$$

$$= \int_\Phi \left(\frac{d\Theta_1(p_1(\mu(s)))}{ds} - \frac{d\Theta_2(p_2(\mu(s)))}{ds} \right)^2 ds. \tag{5}$$

A correspondence which gives a lower value of $J(C_1, C_2; \Phi)$ has aligned the shapes of C_1 and C_2 more closely and is more desirable.

B. Arc-Length Comparison

In many applications, detailed prior knowledge about the non-rigidity is not available and the best that can be said is that, in addition to being shape-based, the correspondence should be as close to a uniform mapping as possible. Recall that ds_1/ds and ds_2/ds are constant for a uniform mapping bimorphism and are equal to $L_{C_1}/\sqrt{2}$ and $L_{C_2}/\sqrt{2}$, respectively. We can measure the closeness of the bimorphism to a uniform mapping by using the second derivatives of s_1 and s_2 with respect to s

$$H(C_1, C_2; \Phi) = 2 \int_\Phi \left(\frac{1}{L_{C_1}} \frac{d^2 s_1}{ds^2} \right)^2 + \left(\frac{1}{L_{C_2}} \frac{d^2 s_2}{ds^2} \right)^2 ds.$$

Since the factor of 2 is just a scaling factor, we can drop it and define $H(C_1, C_2; \Phi)$ as

$$H(C_1, C_2; \Phi) = \int_\Phi \left(\frac{1}{L_{C_1}} \frac{d^2 s_1}{ds^2} \right)^2 + \left(\frac{1}{L_{C_2}} \frac{d^2 s_2}{ds^2} \right)^2 ds. \tag{6}$$

A correspondence which gives a lower value of $H(C_1, C_2; \Phi)$ is closer to the uniform mapping and is more desirable.

Occasionally, we have additional prior knowledge about the limits of local arc-length stretching and shrinking. Since $ds_1/ds_2 = (ds_1/ds)/(ds_2/ds)$, this knowledge can be expressed as the constraint

$$T_1 \le \frac{ds_1}{ds_2} \le T_2, \qquad \text{which is}$$

$$T_1 \frac{ds_2}{ds} \le \frac{ds_1}{ds} \le T_2 \frac{ds_2}{ds}. \tag{7}$$

Here, T_1 and T_2 are the upper and lower bounds on ds_1/ds_2.

C. The Objective Function

Since lower values of $J(C_1, C_2; \Phi)$ and $H(C_1, C_2; \Phi)$ are desirable, we can construct a composite objective function $Q(C_1, C_2; \Phi)$

$$Q(C_1, C_2; \Phi) = J(C_1, C_2; \Phi) + \lambda H(C_1, C_2; \Phi) \tag{8}$$

where, λ is a nonnegative constant. A bimorphism which minimizes $Q(C_1, C_2; \Phi)$ subject to the constraints of (3) and possibly those of (7) is the most desirable one.

For strictly shape-based correspondence, λ can be set to a very low value. In that case, H acts as a regularizing term.

1) Properties of the Objective Function: The objective function $Q(C_1, C_2; \Phi)$ has the following properties [31].

1) It is symmetric with respect to C_1 and C_2. That is, the value of the objective function remains the same if we switch the labels of C_1 and C_2 while keeping the same pairing of points.
2) Suppose that two curves happen to have equal arc length and have identical curvature arc-length functions. Then a simple calculation shows that $Q(C_1, C_2; \Phi)$ is minimized by a uniform mapping bimorphism which aligns the curvature functions of the two curves. So we can think of $Q(C_1, C_2; \Phi)$ as the appropriate generalization of curvature comparison (which is valid only in rigid transformations) to nonrigid situations.
3) Finally, suppose that we scale C_1 while dragging the correspondence along. Scaling changes the infinitesimal arc length ds_1 but, since the net arc length L_{C_1} also scales accordingly, the arc length of the bimorphism ds does not change. Further, there is no change in the directions of the corresponding normals (since we have dragged the correspondence along while scaling), so the value of the $J(C_1, C_2; \Phi)$ is invariant to the scale of C_1 and, by a similar argument, also invariant to the scale of C_2. That is, $J(C_1, C_2; \Phi)$ measures shape and is invariant to size.

   This property is the reason for defining the bimorphism arc length as a weighted Euclidean arc length in (2).

These properties assure us that $Q(C_1, C_2; \Phi)$ works the way we want it to work: that it is not dependent on curve labeling, that it properly generalizes the notion of local shape in nonrigid situations, and that its shape comparison part is not size dependent.

IV. Numerical Algorithm

We now choose a numerical procedure for calculating $Q(C_1, C_2; \Phi)$ and for finding an optimal bimorphism with

respect to it. We have to settle two numerical issues. The first is that of robustly calculating the curvature arc-length functions $\kappa_1(s_1)$ and $\kappa_2(s_2)$. The second issue is that of minimizing $Q(C_1, C_2; \Phi)$. We will address the curvature calculation first.

A. Curvature Calculation

Curvature depends on the second derivatives of the curve parameterization and is not robust when it is calculated from raw data. If a noise model is available for the data, it should be exploited to filter the curve and provide some noise immunity. A different approach is followed here. The curve is adaptively smoothed before curvature calculation to reduce the effect of noise [10], [23].

Briefly, this technique works as follows. The curve is sampled uniformly along its arc length at a high resolution. At each sample, a window is centered and its size is adaptively varied, depending on the residue of a quadratic fit to the x and y coordinates of the samples in the window. Let the coordinates and the arc length of the kth sample point be (x_k, y_k) and s_k, respectively. Let a window of size $2N+1$ be centered on the jth sample. Then coefficients $\alpha_x, \beta_x, \gamma_x$ and $\alpha_y, \beta_y, \gamma_y$ are found, which minimize

$$E = \sum_{i=-N}^{N} \{x_{j+i} - \alpha_x(s_{j+i} - s_j)^2 - \beta_x(s_{j+i} - s_j) - \gamma_x\}^2 + \{y_{j+i} - \alpha_y(s_{j+i} - s_j)^2 - \beta_y(s_{j+i} - s_j) - \gamma_y\}^2.$$

The optimizing coefficients are given by the equation at the bottom of the page where $\Delta s_i = s_{j+i} - s_j$.

The root-mean-square (rms) error $(\sqrt{E/2N+1})$ is evaluated at the optimal coefficients. The window size is systematically increased and the rms error recalculated until it increases to two pixels. At that point, the window size and the optimal coefficients are taken to be the best adaptive fit to the curve in the window. The curvature of the quadratic approximation at the center of the window is calculated according to

$$\kappa = \frac{2(\beta_x \alpha_y - \beta_y \alpha_x)}{(\beta_x^2 + \beta_y^2)^{1/3}}.$$

Repeating this procedure for all sample points gives the adaptive curvature at those samples. The value of the curvature between sample points is found by linear interpolation from the values at the sample points.

The procedure is adaptive because it adjusts the width of the approximation window, depending on the quality of the quadratic fit. Larger window sizes are used over relatively flatter regions, while smaller window sizes are over relatively curved regions.

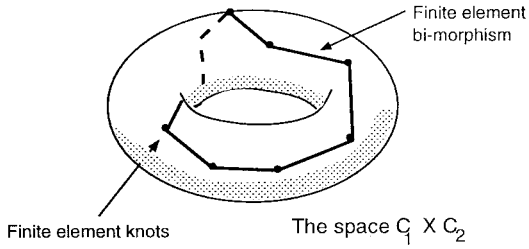

Fig. 4. Approximation to the bimorphism.

With the values of the curvature arc-length function in hand, the value of $Q(C_1, C_2; \Phi)$ can be calculated for any bimorphism by numerical quadrature, using (4), (6), and (8).

B. Optimization

For numerical optimization, the bimorphism is approximated by piecewise linear finite elements (Fig. 4) with K knot points (i.e., the bimorphism is approximated as a polyhedral curve). This decreases the size of the search space without greatly sacrificing accuracy.

The optimization has two phases. The first phase corresponds to finding an optimal uniform mapping bimorphism. This is done by starting with the uniform bimorphism for $\delta = 0$ and increasing δ in small steps over its range of values. At each step, $Q(C_1, C_2; \Phi)$ is evaluated numerically and the minimizing bimorphism is retained.

The minimizing bimorphism is used as the initial condition for a coordinate-wise gradient descent of the finite element bimorphism using the objective function $Q(C_1, C_2; \Phi)$ and subject to constraints of (3). For most of the experimental data processed by this algorithm, the coordinate descent step terminates in about 20 s on a SGI Indigo^2.

V. Experiments

This section contains two sets of experiments with biomedical images. The first set demonstrates the application of the algorithm to tracking heart motion and the second set illustrates its application to other data sets.

A. Heart Motion

The algorithm was used to track the endocardium of dog hearts in MRI sequences. A total of 21 studies were available.

$$\begin{pmatrix} \alpha_x \\ \beta_x \\ \gamma_x \\ \alpha_y \\ \beta_y \\ \gamma_y \end{pmatrix} = \begin{pmatrix} \sum \Delta s_i^4 & & & & & & \\ \sum \Delta s_i^3 & & & & & & \\ \sum \Delta s_i^2 & 0 & 0 & 0 & & & \\ \sum \Delta s_i^3 & \sum \Delta s_i^2 & \sum \Delta s_i & 0 & 0 & 0 & \\ \sum \Delta s_i^2 & \sum \Delta s_i & 2N+1 & 0 & 0 & 0 & \\ 0 & 0 & 0 & \sum \Delta s_i^4 & \sum \Delta s_i^3 & \sum \Delta s_i^2 & \\ 0 & 0 & 0 & \sum \Delta s_i^3 & \sum \Delta s_i^2 & \sum \Delta s_i & \\ 0 & 0 & 0 & \sum \Delta s_i^2 & \sum \Delta s_i & 2N+1 & \end{pmatrix}^{-1} \times \begin{pmatrix} \sum x_i \Delta s_i^2 \\ \sum x_i \Delta s_i \\ \sum x_i \\ \sum y_i \Delta s_i^2 \\ \sum y_i \Delta s_i \\ \sum y_i \end{pmatrix}$$

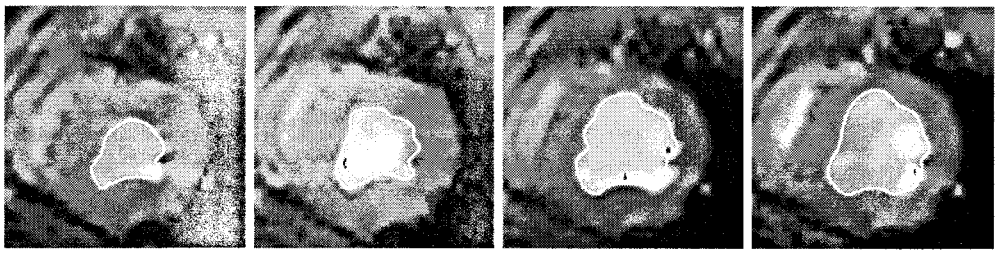

Fig. 5. MRI Images of a beating heart.

TABLE I
CONSTANTS USED IN EXPLORATORY NUMERICAL EXPERIMENTS

Symbol	Name	Value
K	Knot Points for Bimorphism	10
λ	Regularization Constant	0.1

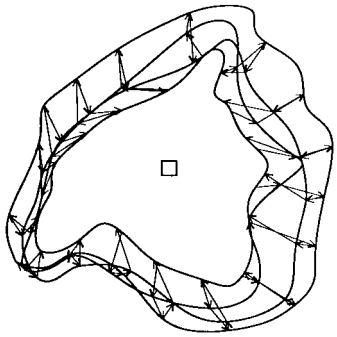

Fig. 6. Evaluation of the algorithm.

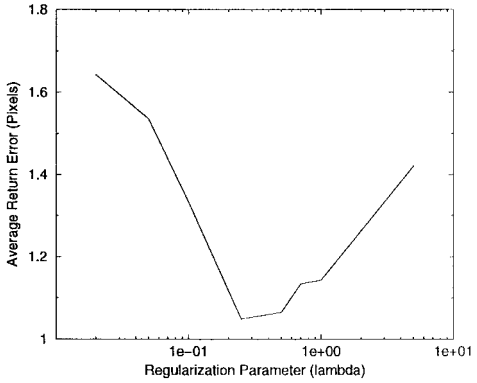

Fig. 7. Distribution of the return error as a function of regularization.

In each study, the heart was imaged in three dimensions, with 16 frames per heart beat. The study contains images of motion under baseline as well as post-infarct conditions. The infarct was produced surgically by occluding a coronary artery [the left anterior descending (LAD)].

The image plane just below the papillary muscle was chosen as the site to evaluate motion. This convention has been used by others [24]. In each study, starting from the end systole, four images (uniformly separated in time) were chosen from the sequence of 16 images. From now on, we will call these images image one, two, three, and four. Each image was enlarged six times and a B-spline active contour algorithm (snake) was used to interactively outline the endocardium. The contours are called curve one, curve two, curve three, and curve four, respectively. Fig. 5 shows the images in one sequence with overlayed contours.

To check that the proposed algorithm was suitable for analyzing this data, it was applied with the parameters shown in Table I. Since no prior knowledge was available about arc-length constraints, the $H(C_1, C_2; \Phi)$ component in the objective function acted as a regularization penalty. The algorithm found the optimal bimorphisms between curves one and two, two and three, three and four, and four and one. Fig. 6 shows the results. The square in the center of the figure shows the size of single pixel (before magnification to six times).

Having established that the algorithm looked promising, a more systematic investigation was carried out. There were two goals in this investigation. The first was to systematically select the regularization parameter λ. The second goal was to evaluate the accuracy of the algorithm.

With current technology, it is impossible to obtain ground truth about the motion of all points on the endocardium. In the absence of ground truth, a performance measure called return error was created and used. It measures the consistency of the algorithm.

B. The Return Error

Since heart motion is periodic, at the end of the heart beat period each tracked point should return to its initial location. Any error due to the algorithm or noise will return the point to a different location and the difference, which we will call return error, measures the consistency of the algorithm and its robustness to noise.

The return error can also be used to determine λ. To see this, consider how the return error changes as a function of λ. When λ is very small, there is little regularization and noise propagates through the tracking algorithm, causing high return errors. As λ increases, regularization increases and the effect of noise is reduced. Thus, the return error is likely to decrease with increasing λ. On the other hand, at very large values of λ, the regularization term in $Q(C_1, C_2; \Phi)$ will dominate its value. Since the regularization term depends on second derivatives, to keep it low, the optimal bimorphism will be very close to a linear mapping of the first curve onto the next. If the true underlying motion is nonrigid, estimating it by a close-to-linear mapping may not give good results and the return error will be high again. Thus, the return error should exhibit a minimum as a function of λ and the value of λ at which the minimum occurs is an appropriate value for the regularization parameter.

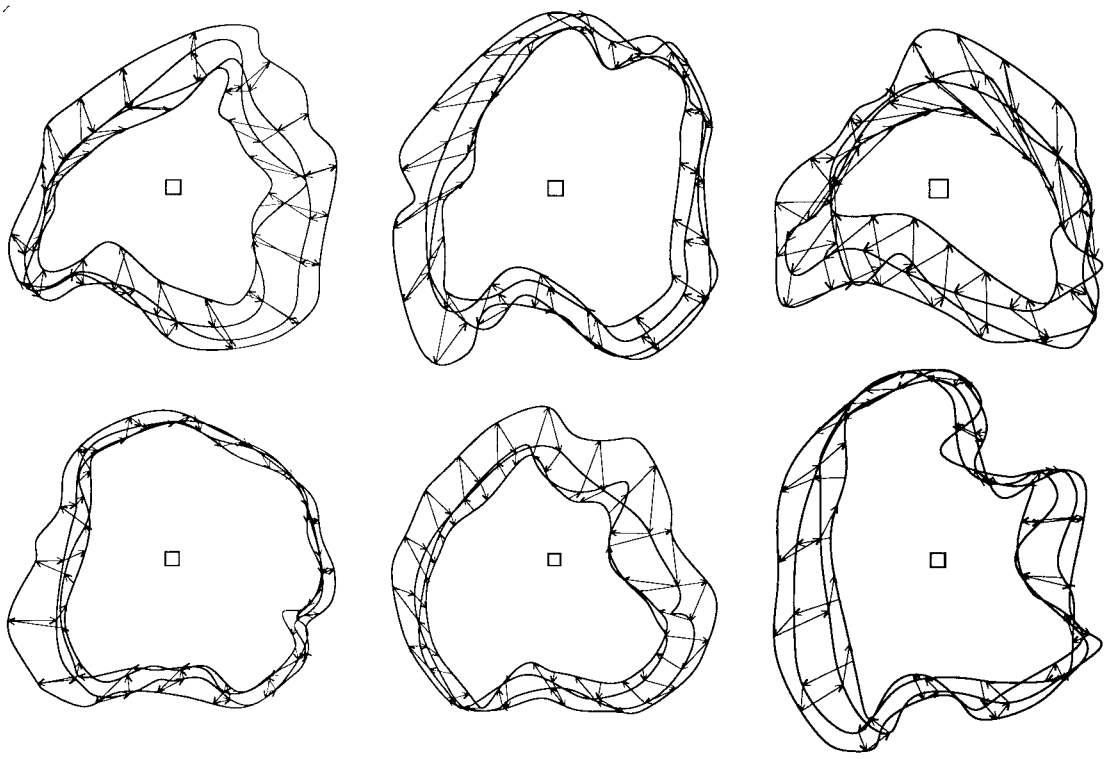

Fig. 8. Tracking the endocardium.

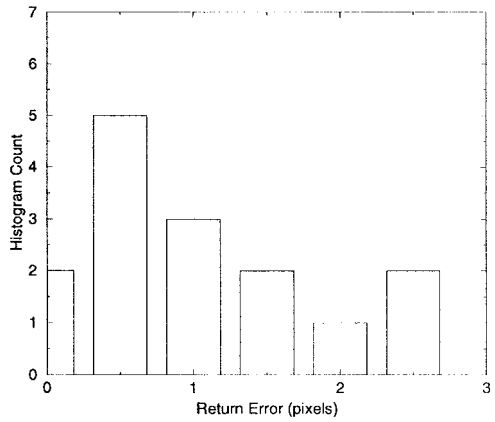

Fig. 9. Distribution of the return error for the remaining cases.

Fig. 7 shows the plot of the average return error on a training set of six randomly chosen sequences from the 21 MRI sequences. Numerically, the return error was calculated by uniformly placing ten points along curve one and tracking the points via the optimal bimorphisms. The return error in Fig. 7 is expressed in terms of the resolution of the imaging system (the unit is a pixel). The regularization parameter was varied over two orders of magnitude from 0.02 to 5.0. As predicted, the return error does exhibit a minimum. It occurs at $\lambda = 0.25$.

The value of $\lambda = 0.25$ was chosen as the appropriate regularizing value for the analysis of the remaining 15 images. Some of the results are shown in Fig. 8.

Return errors were calculated for all the 15 cases. The average return error was 1.32 pixels, which is comparable to the resolution of the imaging system. Fig. 9 shows the histogram of the return error for all 15 cases.

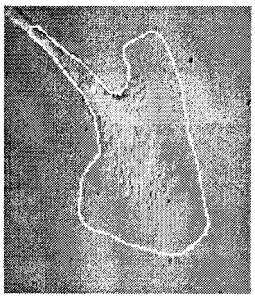

Fig. 10. Growth cone: initial image.

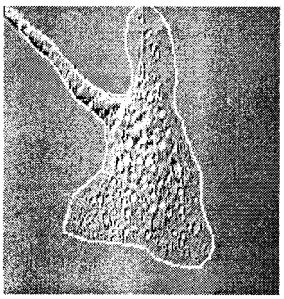

Fig. 11. Growth cone: final image.

C. Other Data Sets

Currently, the performance of the algorithm is being evaluated on other data sets. One example is shown in Figs. 10–12. Figs. 10 and 11 show the initial and final images of the growth

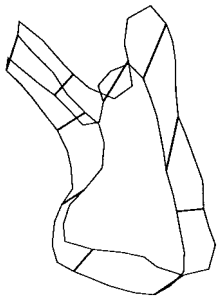

Fig. 12. Correspondence between outlines.

cone lamellipodium of an aplysia neuron with the outlines drawn on the image. Fig. 12 shows the nonrigid correspondence obtained by the algorithm. The correspondence strongly suggests the development of the growth cone.

D. Discussion

Our experience with this and other datasets suggests the following criteria for determining whether this algorithm is suitable to an application.

1) Is the desired correspondence shape-based? That is, is the pairing of curved regions with curved regions and flat regions with flat regions the appropriate solution to the problem? In some applications, e.g., correspondence based on landmarks, shape may not be relevant. If the desired correspondence is shape based, then this algorithm is applicable.
2) Can the curvature arc-length function be calculated reliably? This is a key factor since curvature arc-length functions are central to the calculation of the objective function $Q(\,)$. The adaptive procedure of Section IV-A is one way of calculating curvature. Any other reliable way of calculating curvature can also be used.
3) Is there a way of determining, or learning, the regularization coefficient λ? Often this requires the user to evaluate the performance of the algorithm on a training set for different values of λ and then decide the optimal λ. Often this may be done informally, with the user tweaking the parameter until acceptable performance is obtained on a small data set. We recommend a more objective procedure, which takes the context of the application into account. Generic criteria for determining regularization parameters are also available [32].

VI. Conclusion

A theory and algorithm for shape-based nonrigid correspondence between plane curves is developed in this paper. The algorithm allows for one-to-one as well as non-one-to-one correspondences. The latter are of a restricted type and allow for curve segments to shrink to points or points to expand to curve segments. An objective function for comparing the local shape and stretching of curves is also proposed. This objective function has a number of desirable properties: it is symmetric with respective to the two curves; it uses local shape information in a manner that is consistent with nonrigid correspondences; and it allows easy addition of certain common constraints. Finally, numerical techniques for calculation and optimization are also proposed.

The application of this algorithm to heart motion was also reported. The performance of the algorithm was evaluated, using return error of the tracked points.

Acknowledgment

The author acknowledges the contributions of D. O'Shea of Mt. Holyoke College, MA in developing the theory presented here. Discussions with A. Rangarajan and L. Berlinger, both of Yale University, and with E. Bardinet of Universidad de Granada helped enormously. A. Sinusis and E. Heller, also of Yale University, generously provided the MR data. C. Davatzikos of Johns Hopkins University suggested the adaptive technique for calculating curvature arc-length functions.

References

[1] M. A. Armstrong, *Basic Topology*. Berlin, Germany: Springer-Verlag, 1983.
[2] L. Axel and L. Dougherty, "MR imaging of motion with spatial modulation of magnetization," *Radiology*, vol. 17, no. 1, pp. 841–845, 1989.
[3] E. Bardinet, L. Cohen, and N. Ayache, "Tracking and motion analysis of the left ventricle with deformable superQuadrics," *Med. Image Anal.*, vol. 1, pp. 129–149, 1996.
[4] M. Black and A. Rangarajan, "The outlier process: Unifying line processes and robust statistics," in *Proc. IEEE Conf. CVPR*, 1994.
[5] F. L. Bookstein, *Morphometric Tools for Landmark Data*. Cambridge, U.K.: Cambridge Univ. Press, 1991.
[6] F. L. Bookstein, "Landmark methods for forms without landmarks: Morphometrics of groups differences in outline shape," *Med. Image Anal.*, vol. 1, no. 3, pp. 225–243, 1996.
[7] A. Buda and E. Delp, "Digital two-dimensional echocardiography," in *Digital Cardiac Imaging*, A. Buda and E. Delp, Eds. Boston, MA: Martinus Nijhoff, 1985, pp. 41–60.
[8] S. W. Chen, S. T. Tung, C. Y. Fang, S. Cherng, and A. K. Jain, "Extended attributed string matching for shape recognition," *Comput. Vision Image Understanding*, vol. 70, no. 1, pp. 36–50, Apr. 1998.
[9] I. Cohen, N. Ayache, and P. Sulger, "Tracking points on deformable objects using curvature information," in *Lecture Notes in Computer Science: ECCV '92*. Berlin, Germany: Springer-Verlag, 1992, vol. 588.
[10] C. Davatzikos and R. N. Bryan, "Using a deformable surface model to obtain a shape representation of the cortex," *IEEE Trans. Med. Imag.*, vol. 15, pp. 785–795, Dec. 1996.
[11] M. Demi, R. Calami, G. Coppini, and G. Valli, "A visual framework for the study of cardiac motion," *Comput. Cardiol.*, 1990, pp. 30–34.
[12] J. M. B. Dias and J. M. N. Leitao, "Wall position and thickness estimation from sequences of echocardiographic images," *IEEE Trans. Med. Imag.*, vol. 15, pp. 25–38, Feb. 1996.
[13] J. S. Duncan and R. L. Owen *et al.*, "Measurement of nonrigid motion using contour shape descriptors," in *Proc. CVPR*, Maui, HI, 1991, pp. 318–324.
[14] J. Feldman and N. Ayache, "Locally affine registration of free-form surfaces," in *Proc. CVPR*, Seattle, WA, 1994, pp. 496–501.
[15] N. Friedland and D. Adam, "Automatic ventricular boundary detection from sequential ultrasound images using mathematical morphology," *IEEE Trans. Med. Imag.*, vol. 4, pp. 344–353, 1989.
[16] S. H. Gwydir, H. M. Buettner, and S. M. Dunn, "Non-rigid motion analysis and feature labeling of the growth cone," in *Proc. IEEE Workshop Biomedical Image Analysis*, Seattle, WA, June 24–25, 1994.
[17] C. Kambhamettu and D. Goldgof, "Point correspondence recovery in nonrigid motion," in *Proc. CVPR*, 1992, pp. 222–227.
[18] R. S. Ledley, "High-speed automatic analysis of biomedical pictures," *Science*, vol. 146, no. 3461, pp. 216–223, 1964.
[19] E. T. Lee, "Shape-oriented chromosome classification," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-5, pp. 629–632, 1975.
[20] M. Leyton, "A process-grammar for shape," *Artif. Intell.*, vol. 34, pp. 213–247, 1988.
[21] S. Loncaric, "A survey of shape analysis techniques," *Pattern Recognit.*, vol. 31, no. 8, pp. 983–1001, 1998.

[22] D. Mumford, "Mathematical theories of shape: Do they model perception?," in *Proc. Conf. 1570 Soc. Photo-Optical Industrial Engineers*, 1991, pp. 2–10.
[23] A. Manceaux-Demiau, R. N. Bryan, and C. Davatzikos, "A probabilistic ribbon model for shape analysis of the cerebral sulci: Application to the central sulcus," *J. Comput. Assist. Tomogr.*, vol. 22, 1998.
[24] J. C. McEachen and J. S. Duncan, "A constrained analytic solution for tracking nonrigid motion of the left ventricle," in *Proc. IEEE 18th Ann. Northeast Bioengineering Conf.*, University of Rhode Island, Mar. 1992, pp. 137–138.
[25] ———, "Shape-based tracking of left ventricular wall motion," *IEEE Trans. Med. Imag.*, vol. 16, pp. 270–283, June 1997.
[26] T. Pavlidis, "A review of algorithms for shape analysis," *Compu. Vision, Graph., Image Processing*, vol. 7, pp. 243–258, 1978.
[27] F. H. Shehan, E. L. Bolson, H. T. Dodge, D. G. Mathey, J. Schofer, and H. W. Woo, "Advantages and applications of the centerline method for characterizing regional ventricular function," *Circulation*, vol. 74, pp. 293–305, 1986.
[28] C. Slager, T. Hooghoudt, P. Serruys, J. Schuurbiers, J. Reiber, G. Meester, P. Verdouw, and R. Hugenholtz, "Quantitative assessment of regional left ventricular motion using endocardial landmarks," *JACC*, vol. 7, no. 2, pp. 317–326, 1986.
[29] C. G. Small, *The Statistical Theory of Shape*. Berlin, Germany: Springer-Verlag, 1996.
[30] H. D. Tagare, "Non-rigid curve correspondence for estimating heart motion," in *Information Processing in Medical Imaging* (Lecture Notes in Computer Science). Berlin, Germany: Springer-Verlag, 1997, vol. 1320.
[31] H. D. Tagare and D. O'Shea, "Non-rigid shape comparison of plane curves in images," *SIAM J. Appl. Math.*, submitted for publication.
[32] G. Wahba, "Spline models for observational data, *Soc. Indus. Appl. Math.*, 1990.
[33] E. A. Zerhouni, D. M. Parish, W. J. Rogers, A. Yang, and P. Shapiro, "Human heart: Tagging with MR imaging—A method for noninvasive assessment of myocardial motion," *Radiology*, vol. 169, no. 1, pp. 59–63, 1988.