Multiscale Deformable Model Segmentation and Statistical Shape Analysis Using Medial Descriptions

Sarang Joshi*, *Member, IEEE*, Stephen Pizer, *Senior Member, IEEE*, P. Thomas Fletcher, Paul Yushkevich, Andrew Thall, *Student Member, IEEE*, and J. S. Marron

***Abstract*—This paper presents a multiscale framework based on a medial representation for the segmentation and shape characterization of anatomical objects in medical imagery. The segmentation procedure is based on a Bayesian deformable templates methodology in which the prior information about the geometry and shape of anatomical objects is incorporated via the construction of exemplary templates. The anatomical variability is accommodated in the Bayesian framework by defining probabilistic transformations on these templates. The transformations, thus, defined are parameterized directly in terms of natural shape operations, such as *growth* and *bending*, and their locations. A preliminary validation study of the segmentation procedure is presented. We also present a novel statistical shape analysis approach based on the medial descriptions that examines shape via separate intuitive categories, such as global variability at the coarse scale and localized variability at the fine scale. We show that the method can be used to statistically describe shape variability in intuitive terms such as growing and bending.**

***Index Terms*—Deformable templates, image segmantation, medial geometry, statistical shape analysis.**

I. Introduction

MODERN anatomical imaging technologies have enabled extremely detailed study of anatomy, and the development of functional imaging modalities has provided detailed *in vivo*-associated information regarding physiological function. While modern imaging modalities provide exquisite imagery of the anatomy and its function, automatic segmentation of these images and the precise quantitative study of the biological variability exhibited within them continue to pose a challenge. In this paper, we present a multiscale medial framework based on deformable templates [7], [12], [16], [20], [22], [26], [28] for the automatic extraction and analysis of the shape of anatomical objects from the brain and abdomen imaged via magnetic resonance imaging and computed tomography (CT),respectively. The multiscale deformable template approach is based on the medial-axis representation of objects first proposed by Blum [4] for studying shape. Shape analysis is emerging as an important area of medical image processing because it has the potential to improve both the accuracy of medical diagnosis and the understanding of processes behind growth and disease [9], [10]. We present a novel shape analysis method that can quantify shape variability, explain it in intuitive shape-derived terms, and pinpoint the places where variability is most pronounced. We use our method to analyze the shape of the mid-sagittal slice of the corpus callosum. This approach is an extension of work by Pizer [23] and Fritsch [15] on deformable medial representation of objects.

We adopt a Bayesian approach of incorporating prior knowledge of the anatomical variations and the variation of the imaging modalities. Following the deformable templates paradigm, we construct exemplary templates to incorporate prior information about the geometry and shape of the anatomical objects under study. The infinite anatomical variability is accommodated in the Bayesian framework by defining probabilistic transformations on these templates [16]. The segmentation problem in this paradigm is that of finding the transformation S of the template that maximizes the posterior

$$P(S|\text{data}) \propto P(\text{data}|S)P(S)$$

where $P(S)$ is the prior probability function capturing prior knowledge of the anatomy and its variability and $P(\text{data}|S)$ is the data likelihood function capturing the image data-to-geometry relationship. For efficiency of implementation, we equivalently maximize the log-posterior given by

$$\text{Log}P(S|\text{data}) = \text{Log}P(\text{data}|S) + \text{Log}P(S), \text{ up to an additive constant.}$$

The modeling approach taken in this paper for building exemplary templates and associated transformations is based on a multiscale medial representation. The transformations defined in this framework are parameterized directly in terms of natural shape operations, such as *thickening* and *bending*, and their locations.

This multiscale approach has many stages of scale; at each the geometric primitives are intuitive for that scale. At each scale, the spacing is proportional to the size of space (modeling aperture) that they summarize, this leads to a spatial tolerance that successively decreases as the scale is refined. A Markov random field approach is used for defining the energetics of the log probabilities needed for the posterior. The log probabilities at a given scale are not only conditioned on a neighborhood at that scale, but conditioned on the result of the next coarser scale. The posterior at each scale can then be separately optimized, successively refining the scale. The multiscale nature of our approach allows

Manuscript received November 9, 2001; revised April 2, 2002. This work was supported by the National Institutes of Health (NIH) under Grants P01 CA47982 and Grant R01 CA67183. *Asterisk indicates corresponding author.*

*S. Joshi is with the Medical Image Display and Analysis Group, University of North Carolina at Chapel Hill, Chapel Hill NC 27514 USA (e-mail: joshi@radonc.unc.edu).

S. Pizer, P. T. Fletcher, P. Yushkevich, A. Thall, and J. S. Marron are with the Medical Image Display and Analysis Group, University of North Carolina at Chapel Hill, Chapel Hill NC 27514 USA.

Publisher Item Identifier S 0278-0062(02)05538-6.

0278-0062/02$17.00 © 2002 IEEE

for the investigation of these properties at various scales from the coarse scale of entire body sections to the fine scale on the order of the resolution of the imaging modality. The size properties derived from the medial description allow the creation of natural levels of scale, each suited for shape description at that scale level.

Sections II and III discuss the medial representation of objects; Section III, in particular, discusses the deformation of models to fit image data and the geometric measures used in the log prior term which measures geometric typicality. Section IV discusses the log-likelihood term measuring the match of a deformed model to a target image, and Section V gives segmentation results using this method. In Section VI, methodology for the statistical analysis of shape variability is presented.

II. Medial Representation of Objects

Many authors in image analysis, geometry, human vision, computer graphics, and mechanical modeling have come to the understanding that the medial relationship between points on opposite sides of a figure is an important factor in the object's shape description. Biederman [1], Marr [19], Burbeck [5], Leyton [18], and others have produced psychophysical and neurophysiological evidence for the importance of two-dimensional (2-D) medial relationships in human vision. Medial geometry has also been explored in three dimensions by Nackman [21] and Siddiqi [25] and medial-axis modeling techniques have been applied by many researchers, including Bloomenthal [3] and Sherstyuk [24], Cameron and Robb [6], Storti [27], and Blanding [2]. Of these, Bloomenthal and Sherstyuk created skeletal-based convolution surfaces; Cameron *et al.* explored skeleton-based surface deformation; and Storti and Blanding did skeletal-based solid modeling in a CAD framework.

Our representation, described in [23], expands the notion of medial relations from that of a simple medial skeleton by 1) including a width-proportional tolerance and 2) using a width-proportional discrete sampling of the medial manifold in place of a continuous representation. The advantages—relative to the traditional medial-axis methods descended from Blum [4]—are in representational and computational efficiency and in stability with respect to boundary perturbation. One of the weaknesses of medial-axis methods for shape analysis is the instability of the medial axis with respect to boundary perturbations; small perturbations of the boundary can drastically change the topology of the medial axis. In the deformable template approach presented in this paper, the medial-axis topology of objects under consideration is fixed to be that of the template. This has the drawback that a given medial-axis topology cannot define the boundary of an object to an arbitrary precision. Associating a tolerance with the boundary position provides opportunities for stages of the representation with successively smaller tolerance. Representations with large tolerance can ignore detail and focus on gross shape; in these large-tolerance stages, discrete sampling can be coarse, resulting in considerable efficiency of manipulation and presentation. Smaller-tolerance stages can focus on refinements of the larger-tolerance stages and, thus, capture more local aspects. To finally accommodate very fine structure—on the order of the voxel dimension of the imaging modality—we define a dense displacement field on the boundary of the medially implied object accommodating the fine-scale boundary features. As this vector displacement field is not medially based, it allows the fine-scale features of the boundary to be accommodated robustly without introducing topological instability into the medial-axis description.

M-reps, the medial representation used in this paper, are based on a hierarchical representation of linked figural models—protrusions, indentations, neighboring figures and included figures—which represent solid regions and their boundaries simultaneously. The linked collection of figural components imply a fuzzy (i.e., probabilistically described) boundary position with a width-proportional tolerance. At small scale these figural boundaries are made precise by displacing a dense sampling of the m-rep implied boundary. A model for a single figure is made from a net (a mesh or a chain) of medial atoms; each atom describes not only a position and width, but also a local figural frame giving figural directions and an object angle between opposing, corresponding positions (*medial involutes*) on the implied boundary. A figure can be expressed as a sequence over scale of medial nets, implying successively refined (smaller tolerance) versions of the figural boundary. At the final stage, a dense displacement field is defined on the boundary of the medially implied object that accommodates the fine-scale perturbations of the object boundary.

A. Single-Figure Description via M-Rep

The medial representation used is based on the medial framework of Blum [4]. In this framework, a geometrical object is represented as a set of connected continuous medial manifolds. These medial manifolds are formed by the centers of all spheres (circles, in two dimensions) that are interior to the object and tangent to the object's boundary at two or more points. The medial description is defined by the centers of the inscribed spheres and by the associated scalar field of their radii. Each continuous segment of the medial manifold represents a medial figure. In this paper, we focus on objects that can be represented by a single medial figure.

In two dimensions, there at two basic types of medial figural segments with medial manifolds $\mathcal{M}$ of dimension zero and one. Figural segments with a single point (zero-dimensional) represent the degenerate case of circular objects. In three dimensions, there are three basic types of medially defined figural segments with corresponding medial manifolds $\mathcal{M}$ of dimension zero, one, and two, respectively. Figural segments with 2-D medial manifolds represent slab-like segments. Tube-like segments where the medial manifold is an one-dimensional (1-D) space curve and spherical segments, where the medial manifold consists of a single point are degenerate cases. Shown in Fig. 1 are examples of slab like and tubular figures.

In nondegenerate cases, the medial manifold divides each figure into two opposing halves. The two halves of the figure's boundary come together at a manifold called the crest. In three dimensions, the crest is a curve that connects the two sides of the slab-like segment. In two dimensions, the crest is composed of two points at the tips of the object boundary. The points on the

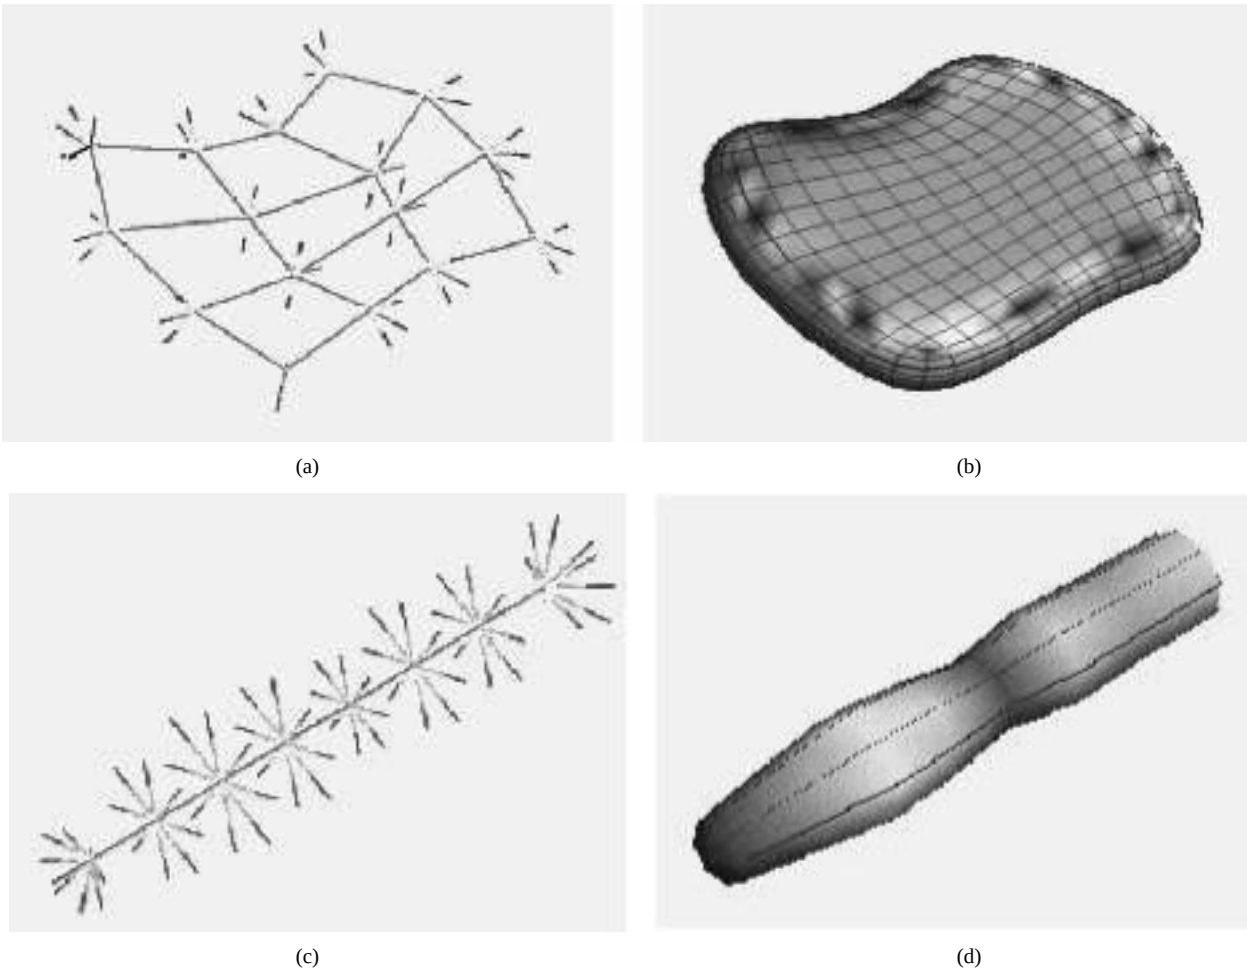

Fig. 1. The top row shows an example of a slab-like figure with a 2-D medial manifold. The bottom row shows a tubular figure with 1-D medial manifold.

crest are maxima of boundary curvature. The centers of spheres (or circles) inscribed at the crest form the edges of the figural segments of the medial manifold. The derivatives of the radial field exhibit catastrophic behavior at these edges, making them difficult to model discretely. [11]

B. Discretized Figural Segments

In this paper, we focus on single-figure objects having single continuous medial manifolds that are discretized. We sample the medial manifold $\mathcal{M}$ over a spatially regular lattice; the elements of this lattice are called *medial atoms*.

For 2-D figures, the lattice is a linear chain of medial atoms $\mathbf{m}_i, i \in [1, N]$. For three-dimensional (3-D) slab-like figures, the lattice is a quadrilateral mesh $\mathbf{m}_{ij}$, $(i, j) \in [1, N] \times [1, M]$. The number of medial atoms in a lattice is inversely proportional to the scale of the medial description.

A 2-D medial atom, shown in Fig. 2(a), carries first-order geometric information at a point on the 1-D medial manifold. A zeroth-order description consists of the position $\mathbf{x}$ and the radius r of the inscribed circle centered at $\mathbf{x}$. A first-order description includes the unit spatial tangent $\mathbf{b}$ of the medial manifold at $\mathbf{x}$ and captures first-order width information by the *object angle* θ, which describes the change in radius along the medial axis by the Blum relationship $\cos\theta = -dr/ds$, for the arc length s on the medial manifold.

The places where the inscribed circle centered at $\mathbf{x}$ touches the two halves of the boundary, indexed by 1, -1, are defined as $\mathbf{y}^1$, $\mathbf{y}^{-1}$ with respective normals $\mathbf{n}^1$, $\mathbf{n}^{-1}$, given by

$$\mathbf{n}^1 = R(\theta)\mathbf{b} \qquad \mathbf{n}^{-1} = R(-\theta)\mathbf{b} \tag{1}$$

$$\mathbf{y}^1 = \mathbf{x} + r\mathbf{n}^1 \quad \mathbf{y}^{-1} = \mathbf{x} + r\mathbf{n}^{-1} \tag{2}$$

where $R(\theta)$ is the rotation matrix.

Due to the catastrophic behavior of the radius field at the crest, we introduce a special medial atom at the ends of medial chains [Fig. 2(b)] for robust sampling. These medial atoms include an extra parameter η that defines the position of the crest point on the object boundary and is given by

$$\mathbf{y}^0 = \mathbf{x} + r\eta\mathbf{b}, \quad \eta \geq 1.0$$

where $\eta = 1.0$ represents a circular end cap and $\eta > 1$ increasing elongation.

Following the construction of 2-D medial atoms above, 3-D interior medial atoms are defined as a four-tuple $\mathbf{m} = \{\mathbf{x}, r, \mathbf{F}, \theta\}$ consisting of: $\mathbf{x} \in \mathbb{R}^3$, the center of the inscribed sphere, $r \in \mathbb{R}^+$, the local width defined as the radius of the sphere; $\mathbf{F} \in SO(3)$, the local frame parameterized by $(\mathbf{n}, \mathbf{b}, \mathbf{b}^\perp)$, where $\mathbf{n}$ is the normal to the medial manifold, $\mathbf{b}$ is the direction in the tangent plane of the fastest narrowing of the implied boundary sections; and $\theta \in [0, \pi/2]$ the object angle

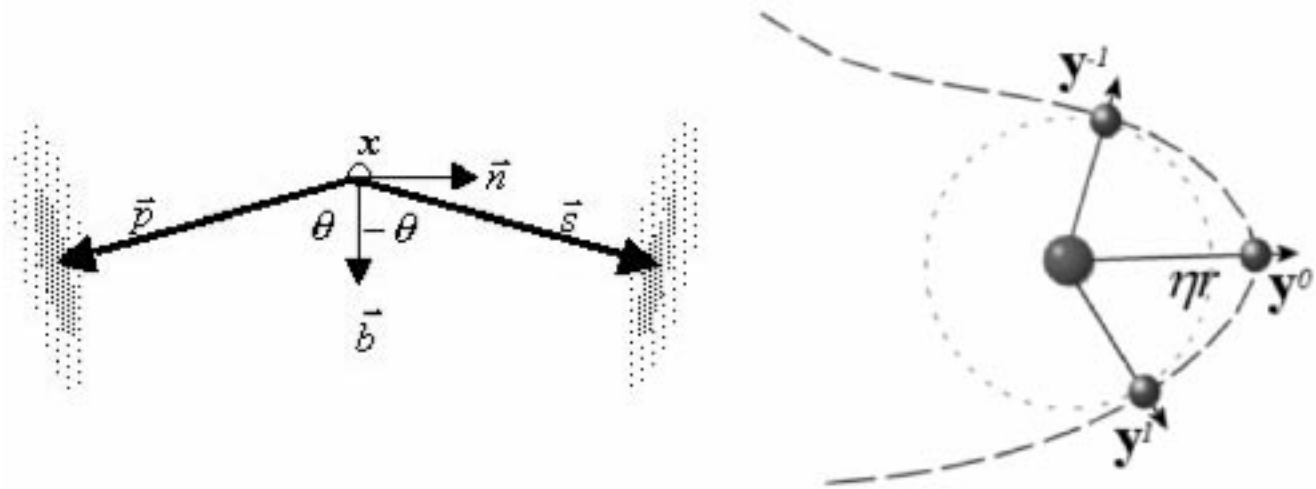

(a) (b)

Fig. 2. (a) A 2-D medial atom $\mathbf{m}$ represents a double tangency of a circle to the boundary. It is a four-tuple defined by position $\mathbf{x}$, width r, vector $\mathbf{b}$ tangent to the medial axis, and object angle θ. (b) An end atom is a medial atom with an additional component η.

determining the angulation of the implied sections of boundary relative to $\mathbf{b}$. The two opposing boundary points implied by the medial atom are given by $\mathbf{y}^1$ and $\mathbf{y}^{-1}$ and the respective normals

$$\mathbf{n}^1 = R_{(\mathbf{b},\mathbf{n})}(\theta)\mathbf{b}, \quad \mathbf{n}^{-1} = R_{(\mathbf{b},\mathbf{n})}(-\theta)\mathbf{b} \tag{3}$$

$$\mathbf{y}^1 = \mathbf{x} + r\mathbf{n}^1, \qquad \mathbf{y}^{-1} = \mathbf{x} + r\mathbf{n}^{-1} \tag{4}$$

where $R_{(\mathbf{b},\mathbf{n})}(\theta)$ is a rotation by θ in the $(\mathbf{b},\mathbf{n})$ plane. Fig. 3 shows these relationships.

For stability at the crest, as in two dimensions, medial atoms on the boundary of the medial manifold also include an extra elongation parameter η determining the crest position.

Both in two dimensions and three dimensions, the above representation $\mathbf{x}$ gives the central location of the solid section of figure that is being represented by the atom $\mathbf{m}$. The scalar r gives the local scale and size of the solid section of figure that is being represented by the atom. The object angle θ and the direction $\mathbf{b}$ also define the gradient of the scalar field r via

$$\nabla r = -\mathbf{b}\cos\theta.$$

The scalar field r also provides a local ruler for the precise statistical analysis of the object.

C. Spline Interpolation of Medial Atoms

For the remainder of this section, we focus on 3-D medial geometry; 2-D construction follows an analogous, simpler form. Given a quadrilateral mesh of medial atoms $\mathbf{m}_{i,j}$, $(i,j) \in [1,\ldots,N] \times [1,\ldots,M]$, we define a continuous medial surface via a Bézier interpolation of the discretely sampled medial atoms. The medial position $\mathbf{x}(u,v)$, $u \in [i, i+1]$, $v \in [j, j+1]$ is defined via a bicubic polynomial interpolation of the form

$$\mathbf{x}(u,v) = \sum_{m,n=0}^{3} \mathbf{d}_{m,n} u^m v^n$$

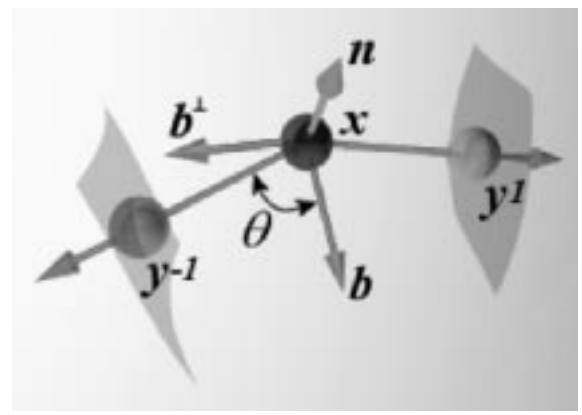

Fig. 3. A 3-D medial atom defined by the four-tuple $\{\mathbf{x}, r, \mathbf{F}, \theta\}$ with involutes perpendicular to the implied surface.

with $\mathbf{d}_{m,n}$ chosen to satisfy the known normal/tangency and continuity conditions at the sample points $\mathbf{x}_{i,j}$. The radius function $r(u,v)$ is also interpolated as a bicubic scalar field on the above interpolated medial manifold, given r and ∇r at the mesh points points $\mathbf{x}_{i,j}$. Having interpolated r and its gradient, the frame $\mathbf{F}$ and the object angle $\theta(u,v)$ are defined via the relationship $\nabla r(u,v) = -\mathbf{b}(u,v)\cos\theta(u,v)$. With the continuous medial manifold thus defined, the continuous implied boundary in the interior can be calculated as

$$\mathbf{y}^1(u,v) = \mathbf{x}(u,v) + rR_{(\mathbf{b},\mathbf{n})}\left(\theta(u,v)\right)\mathbf{b}(u,v) \tag{5}$$

$$\mathbf{y}^{-1}(u,v) = \mathbf{x}(u,v) + rR_{(\mathbf{b},\mathbf{n})}\left(-\theta(u,v)\right)\mathbf{b}(u,v). \tag{6}$$

The end section of the medially implied boundary is defined via the parametric curve $\mathbf{c}(t)$ from one involute to the other passing through the point $\mathbf{x} + \eta r\mathbf{b}$ orthogonally to $\mathbf{b}$. The curve is parameterized by $t \in [-1, 1]$ and defined by

$$\mathbf{c}(t) = \mathbf{x} + r\eta(t)R_{(\mathbf{b},\mathbf{n})}(t\theta)\mathbf{b} \text{ where}$$

$$\eta(t) = (\cos(t\pi) + 1)\,\frac{(\eta - 1)}{2} + 1.$$

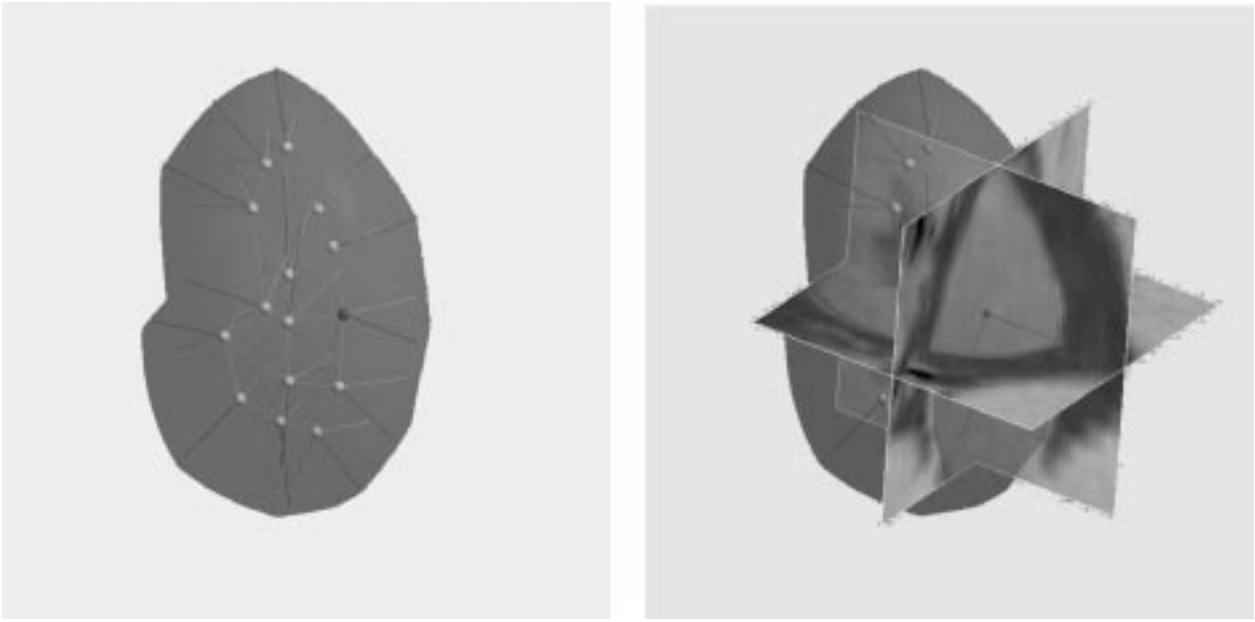

(a) (b)

Fig. 4. The m-rep model of the template kidney. (a) The medial atoms and the implied surface are shown. (b) The model overlaid on the associated CT imagery is shown.

The defined boundary is comprised of the two opposing interior sections, defined by (5) and (6) and the end section defined above which contains the crest of the object.

We have also been exploring the inverse construction by first defining a continuous spline boundary fitted to the points $\mathbf{y}^1$, $\mathbf{y}^{-1}$. The spline boundary is then used to find the exact continuous medial manifold. This method is used to produce the 2-D results.

D. Figural Coordinate System

For both image segmentation and shape analysis, we require consistent correspondences between object boundaries. Further, for image segmentation, the likelihood measure (deformed model to target image match) requires correspondence between template intensities at positions in $\mathbb{R}^3$ relative to the model and target image intensities at positions in $\mathbb{R}^3$ relative to the deformed model. Both of these correspondences are made via the medial geometry.

For 3-D objects, the 2-D continuous medial manifold of a figure is parameterized by (u, v), with u and v taking the atom index numbers at the discrete mesh positions. In two dimensions, the corresponding 1-D linear medial manifold is parameterized via a single parameter u. In both two dimensions and three dimensions, interior boundary points use a single parameter $t \in \{-1, 1\}$ designating the side of the medial manifold on which the point lies. On the end section of the boundary, t varies continually between -1 and 1 passing through 0 at the crest. For single figures, boundary correspondences are defined via the common parameterization (u, v, t) for 3-D objects and (u, t) for 2-D objects.

This also permits parameterization in the neighborhood of the m-rep. Positions in the image in the neighborhood of the implied boundary are indexed by $(u, v, t, \hat{d})$, where (u, v, t) is the parameterization of the closest point on the medially implied boundary and $\hat{d}$ is the signed distance (interior $=$ negative, exterior $=$ positive) from the boundary in multiples of the local radius r of the medial point at (u, v). We are currently studying the implications of an object-based metric based on this figural parameterization.

E. Construction of M-Rep Figures

Using the visualization and computer-aided design techniques developed for 2-D and 3-D medial modeling, we have built numerous models of anatomical objects. In this paper, we focus on the automatic segmentation of the kidney as imaged in CT for radiation treatment for cancer. Shown in Fig. 4 is the template m-rep model of the kidney built from a CT of the abdomen.

III. Transformation of M-Rep Figures

Having defined the construction of typical anatomical objects via m-rep figures, anatomical variability is accommodated by defining a cascade of transformations S^k, $k = 0, \ldots, N$, increasing in dimensionality. These transformations are applied globally to the entire object as well as locally to individual atoms at various scales. Each transformation is applied at its own level of locality to each of the primitives appearing at that level. At each level of locality—by the Markov random field framework—the primitive is related only to immediately neighboring primitives at that level. Each level's result provides both a initial value and a prior for the primitives at the next finer scale level. The transformation at the last (finest) scale level is a dense displacement field applied to the boundary of the figure on the scale of the voxel resolution of the imaging modality.

A. Object-Level Similarity Transformation

We begin with a similarity transformation: $S^0 = (\alpha, \mathbf{O}, \mathbf{t}) \in [(\mathbb{R}^+ \times SO(3)) \ltimes \mathbb{R}^3]$ is defined on the scale of the entire object

and is applied to the whole medial manifold $\mathcal{M}$. The similarity transformation S^0 scales, translates and rotates equally all the medial atoms of the object, that is

$$\mathbf{m}^1_{i,j} = S^0 \circ \mathbf{m}_{i,j} = \{\alpha \mathbf{O}\mathbf{x}_{i,j} + \mathbf{t}, \alpha r, \mathbf{O} \circ \mathbf{F}, \theta\}.$$

Notice that the similarity transformation does not affect the object angle. As the medial representation is invariant under the similarity transformation, this is equivalent to applying the similarity transformation S^0 to the implied boundary $\mathcal{B}$ of the medial mesh to yield the transformed boundary $\mathcal{B}^1$.

A prior is induced on the above transformation based on the displacement of the implied boundary of the objects. Throughout, an independent Gaussian prior on boundary displacement is used with variance proportional to the local radius $r(\mathbf{y})$ with the proportionality constant σ. For the whole object similarity transformation S^0 the log-prior becomes

$$\mathrm{Log}P\left(S^0\right) = \left[-\int_{\mathcal{B}} \frac{\left\|\mathbf{y} - S^0 \circ \mathbf{y}\right\|^2}{2\left(\sigma r(\mathbf{y})\right)^2} d\mathbf{y}\right].$$

In two dimensions, the similarity transformation and the prior are analogous with the rotation matrix being in $SO(2)$ and the translation in $\mathbb{R}^2$.

B. Atom Level Transformation

Having accomplished the gross placement of the figure, attention is now focused on the subsections of the figure defined by each of the medial atoms. At this stage local similarity transformations and rotations of the local angulation, $S^1_{i,j} = (\alpha, \mathbf{O}, \mathbf{t}, \beta)_{i,j} \in [(\mathbb{R}^+ \times SO(3)) \ltimes \mathbb{R}^3] \times [-\pi/2, \pi/2]$, are applied to the medial atom; that is

$$\begin{aligned}\mathbf{m}^2_{i,j} =& S^1_{i,j} \circ \mathbf{m}^1_{i,j} \\ =& \left(\alpha_{i,j}\mathbf{O}_{i,j}\mathbf{x}^1_{i,j} + \mathbf{t}_{i,j}, \alpha_{i,j} r^1_{i,j}, \mathbf{O}_{i,j} \circ \mathbf{F}^1_{i,j}, \theta^1_{i,j} + \beta_{i,j}\right).\end{aligned} \tag{7}$$

The resulting implied boundary is defined as $\mathcal{B}^2$. A prior on the local atom transformations $S^1_{i,j}$ is also induced based on the displacement of the implied boundary with an additional Markov random field prior on the translations, guaranteeing the smoothness of the medial manifold. In keeping with the level of locality, let $\mathcal{B}^1_{i,j}$ be the portion of the implied boundary affected by the atom $\mathbf{m}^1_{i,j}$. The prior energy on the local transformation $S^1_{i,j}$ of the atom $\mathbf{m}^1_{i,j}$ becomes

$$\mathrm{Log}P\left(S^1\right) = \left[-\int_{\mathcal{B}^1_{i,j}} \frac{\left\|\mathbf{y} - \mathbf{y}'\right\|^2}{\left(\sigma r(\mathbf{y})\right)^2} d\mathbf{y} - \sum_{i,j} \sum_{n,m=-1}^{n,m=1} \frac{\left\|\mathbf{t}_{i,j} - \mathbf{t}_{i+n,j+m}\right\|^2}{\left\|\mathbf{x}^1_{i,j} - \mathbf{x}^1_{i+n,j+m}\right\|^2}\right]$$

where $\mathbf{y}$ is the corresponding position on the figural boundary implied by the transformed atom $\mathbf{m}^2$ and $\mathbf{t}_{i,j}$ is the translation component of the local transformation $S^1_{i,j}$. Good association between points on the boundary $\mathbf{y}$ and the deformed boundary $\mathbf{y}'$ is made using the figural coordinate system describe in Section II-D. The point $\mathbf{y}'$ is the point on the deformed model having the same (u, v, t) coordinates as that of the original point $\mathbf{y}$. The integral in the above prior is implemented as a discrete sum over a set of boundary points by defining a sampling of the (u, v, t) coordinate space and calculating the associated implied boundary before and after an atom deformation. Again the 2-D construction is analogous.

C. Dense Boundary Displacement Field Transformation

At the final stage, the implied boundary of the figure is displaced in the normal direction using a dense displacement field defined on the implied boundary $\mathcal{B}^2$

$$\mathbf{y}' \in \mathcal{B}^3 = \mathbf{y} + \mathbf{n}(\mathbf{y}) d(\mathbf{y}), \mathbf{y} \in \mathcal{B}^2$$

where $\mathbf{n}(\mathbf{y})$ is the normal to the implied boundary at $\mathbf{y} \in \mathcal{B}^2$ and $d(\mathbf{y})$ is the scalar displacement field. The definition of this dense displacement field allows us to accommodate the local, small-scale boundary features. These small-scale features are the main cause of the instability in the traditional, Blum medial descriptions of objects. As the vector field is nonmedial based, the small-scale boundary features are robustly accommodated.

As with the local atom transformations, the prior is induced on the dense displacement field using a Markov random field prior derived from energetics associated with thin elastic membranes to guarantee smoothness. The log-prior on the displacement field $d(y)$ becomes

$$\mathrm{Log}P\left(d(x)\right) = \left[-\int_{\mathcal{B}^2} \frac{|d(\mathbf{y})|^2}{\left(\sigma r(\mathbf{y})\right)^2} - \int_{\mathcal{B}^2} |\nabla d(\mathbf{y})|^2 d\mathbf{y}\right]. \tag{8}$$

The above prior is implemented via a discrete approximation as follows. Let $\mathbf{y}_i \in \mathcal{B}^2$, $i = 1, \ldots, N$ be the set of discrete boundary points on the implied boundary $\mathcal{B}^2$. Let $\mathcal{N}(\mathbf{y}_i)$ be the set of neighbors of the point $\mathbf{y}_i$; then

$$\mathrm{Log}P(dx) \approx -\sum_{i=1}^{N} \frac{|d(\mathbf{y}_i)|^2}{\left(\sigma r(\mathbf{y}_i)\right)} - \sum_{i=1}^{N} \sum_{j \in \mathcal{N}(\mathbf{y}_i)} \frac{|d(y_j) - d(y_i)|^2}{\|y_j - y_i\|^2}.$$

becomes the discrete approximation to (8).

IV. Image Data Log-Likelihood

Having defined the transformation and the associated prior energetics, we now define the data likelihood function needed for defining the posterior. We construct this function by defining correlation between a predefined template image I_{temp} and the data I_{data} in the neighborhood of the boundary of the medially defined object $\mathcal{B}$. Letting δ be the size of the collar around the object in multiples of r, the local radius, the data log-likelihood function becomes

$$\int_{-\delta}^{\delta} \int_{\mathcal{B}} I_{\mathrm{temp}}\left(\mathbf{y}, \hat{d}\right) I_{\mathrm{data}}\left(\mathbf{y}', \hat{d}\right) d\mathbf{y} d\hat{d} \tag{9}$$

where $(\mathbf{y}, \hat{d}) \in \mathbb{R}^3$ is the point in the template image at distance $r\hat{d}$ from the boundary point $\mathbf{y}$ and $(\mathbf{y}', \hat{d})$ is the point in the data image at distance $r\hat{d}$ from the boundary point $\mathbf{y}'$ in the transformed object $\mathcal{B}'$. This association between points in the template image and the data image is made using the object centered coordinate system described in Section II-D. Image positions in the neighborhood of the implied boundary are indexed

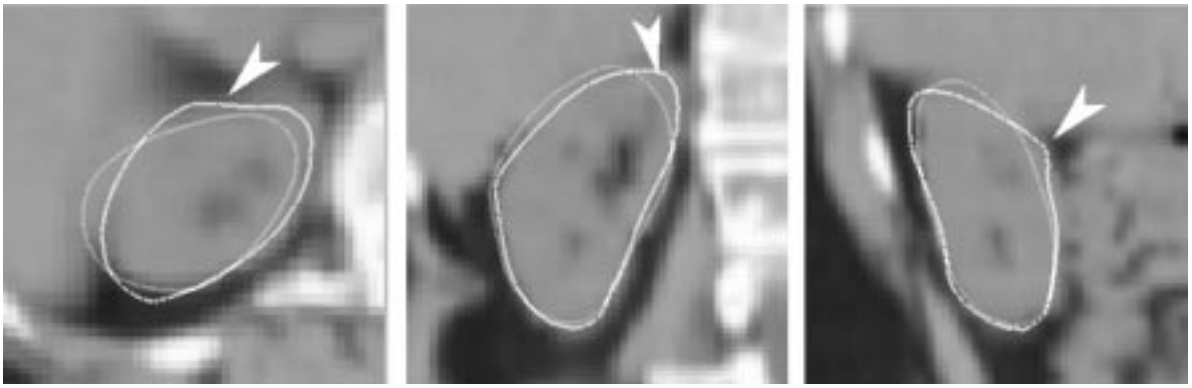

(a) (b) (c)

Fig. 5. (a) Axial, (b) coronal, and (c) sagittal slices through the subject kidney CT data set. The contours show the results of the object similarity transformation and the atom deformation. Notice the improvements in the results at the places marked.

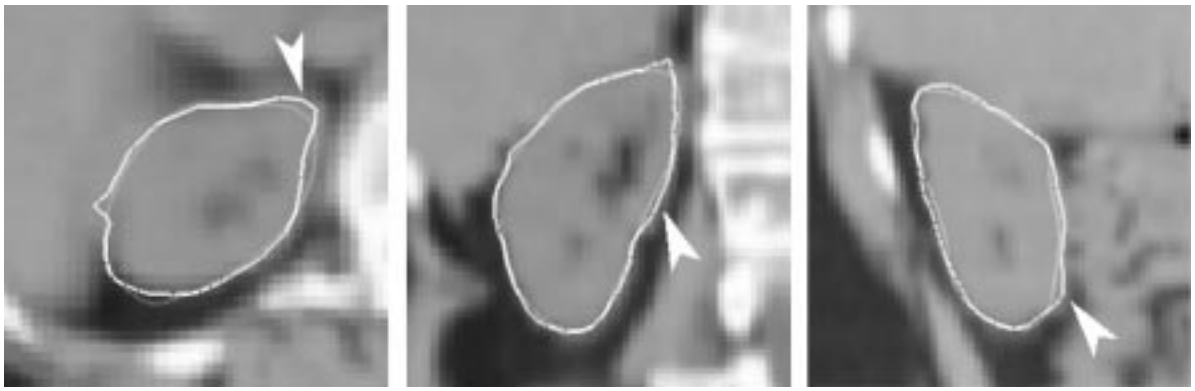

(a) (b) (c)

Fig. 6. (a) Axial, (b) coronal, and (c) sagittal slices through the subject kidney CT data set. The improvement in the segmentation of the kidney after the dense displacement field deformation. The contours show the results of the atom transformation and the dense displacement field deformation.

by $(u, v, t, \hat{d})$, where (u, v, t) is the parameterization in the object centered coordinate system of the closest point on the medially implied boundary $\mathcal{B}$ and $\hat{d}$ is the signed distance (interior = negative, exterior = positive) from the boundary in multiples of the local radius r of the medial point at (u, v). In implementing the correlation defined in (9), care must be taken in computing the surface integral by a discrete voxel summation. The template image needs to be normalized by the determinant of the Jacobian associated with the implied model surface $\mathcal{B}$. At model-building time, intensities in the template image I_{temp} are associated with their positions' $(u, v, t, \hat{d})$ values. As the model deforms, a target image position is calculated for each template $(u, v, t, \hat{d})$ value using the deformed model and the intensity interpolated at that target image position is associated with the corresponding template intensity.

We have have been using two basic types of templates: an analytical template computed from the derivative of the Gaussian and an empirical template learned from an example image from which the template medial model was built.

Using the data likelihood defined above and the prior defined in the previous section, the log posterior is defined as a weighted sum of the two terms, with weights chosen by the user. For optimizing the log-posterior with respect to the global object similarity transformation and the local atom-by-atom transformation, we use a conjugate gradient optimization algorithm; for optimizing the posterior with respect to the dense displacement field $d(\mathbf{y})$ we use a simple steepest decent algorithm.

V. Segmentation Results

We have used the automatic segmentation procedure in three dimensions for extracting the kidney parenchyma, including the renal pelvis, in subjects undergoing radiation treatment for cancer. We present, in detail, results from a series of three data sets. Taking a few seconds, the user rigidly places the template model in the subject data set. This initialization of the algorithm is followed by the automatic hierarchical segmentation which takes on the order of 5 minutes for convergence. At the first scale level, an object similarity transformation is estimated accommodating gross size and orientation differences between the template model kidney and the subject's kidney.

The next scale level in the hierarchical procedure is the atom-deformation process accommodating local shape differences at the scale of individual atoms. Fig. 5, shows the improvement in the segmentation as a result of the atom-deformation, thus accommodating more local object shape changes.

The arrow in Fig. 6 highlights the improvement due to the final stage of the deformation, as the dense displacement field accommodates the fine featured variation in the shapes of the kidney.

For quantitative comparisons of the segmentations of the method with manual segmentations, we have used two metrics from a geometric scoring package developed by Guido Gerig and Matthieu Jomier called VALMET: relative overlap and mean surface distance. The relative overlap measure is defined as the ratio of the intersection of the two segmentations divided by the union. Although the relative overlap is commonly used in the literature for scoring segmentations, it is sensitive to the size of the object and is not very effective in characterizing shape differences between two segmentations. The symmetric, mean surface distance D_s between the boundary of the two segmenations using, Euclidean distance transforms of the segmentations, is defined as follows: let $y_i^1, i = 1, \ldots, N \in \mathcal{B}^1$ and $y_j^2, j = 1, \ldots, M \in \mathcal{B}^2$ be the boundary points of two segmentations $\mathcal{B}^1$, $\mathcal{B}^2$; the mean surface distance then is

$$D_s\left(\mathcal{B}^1, \mathcal{B}^2\right) = \frac{1}{2} \cdot \left[\frac{1}{N}\sum_{i=1}^{N} \min_{j=1\cdots M} \left\|y_i^1 - y_j^2\right\| + \frac{1}{M}\sum_{j=1}^{M} \min_{i=1\cdots N} \left\|y_i^1 - y_j^2\right\|\right].$$

Shown in Table I is the summary of results from the study of the three data sets. The results shown in Figs. 5 and 6 above are typical of the three data sets and are from Data set 613. The segmentation improves at each stage of the algorithm for all three data sets.

To study the performance of the procedure, the left and the right kidneys in 12 data sets were hand segmented by an expert rater. As the aim the validation study is to compare the performance of the automatic segmentation with real clinically used segmentations, the hand segmentations where performed using the segmentation module of our clinically used radiation treatment planning system. Each of the kidneys in the data sets were than segmented using the hierarchical deformable procedure described. Fig. 7 shows the relative overlap and the mean surface distance between the hand segmentations and the result of the final stage of the hierarchical procedure.

The average relative overlap between the human segmentation and the automatic procedure for the 24 kidneys segmented is 89% with a standard deviation of 3.60%. The average mean surface distance is 0.18 cm with a standard deviation of 0.05 cm. All 12 data sets were processed at an isotropic resolution of 0.2 cm. From the results, it can be seen that the accuracy of the segmentation, as measured via the mean surface distance is on the order of the resolution of the data set and is on average, within one pixel of the hand segmentation.

In the results presented, a Gaussian derivative image template was used. The Gaussian derivative template is designed to give increased response at boundaries of objects defined by high gradient. In the CT image, the kidney sits in an environment where part of its boundary is distinctly darker, but part of its boundary region is formed by the liver. In this region, there is either very little contrast, or the liver is a bit lighter, or there is a narrow strip of dark between the kidney intensity and the liver intensity. We expect substantial improvements in the results by the use of a training image template in place of the Gaussian derivative template. This model to image match would be improved even more by a statistical model of this image template, reflecting image intensity variations relative to the geometric model. We

TABLE I
TABLE SHOWING RELATIVE OVERLAP AND MEAN SURFACE DISTANCE BETWEEN THE MANUAL AND THE AUTOMATIC SEGMENTATIONS AT THE DIFFERENT STAGES OF THE HIERARCHICAL PROCEDURE FOR THE THREE DATA SETS

Data Set (cm)	Scale Level	Relative Overlap	Surface Distance (cm)
613	Similarity Transformation	0.85	0.26
	Atom deformation	0.86	0.23
0.15 × 0.15 × 0.5	Field deformation	0.90	0.16
608	Similarity Transformation	0.88	0.22
	Atom deformation	0.89	0.19
0.2 × 0.2 × 0.4	Field deformation	0.93	0.14
1402	Similarity Transformation	0.77	0.65
	Atom deformation	0.86	0.38
0.15 × 0.15 × 0.3	Field deformation	0.90	0.38

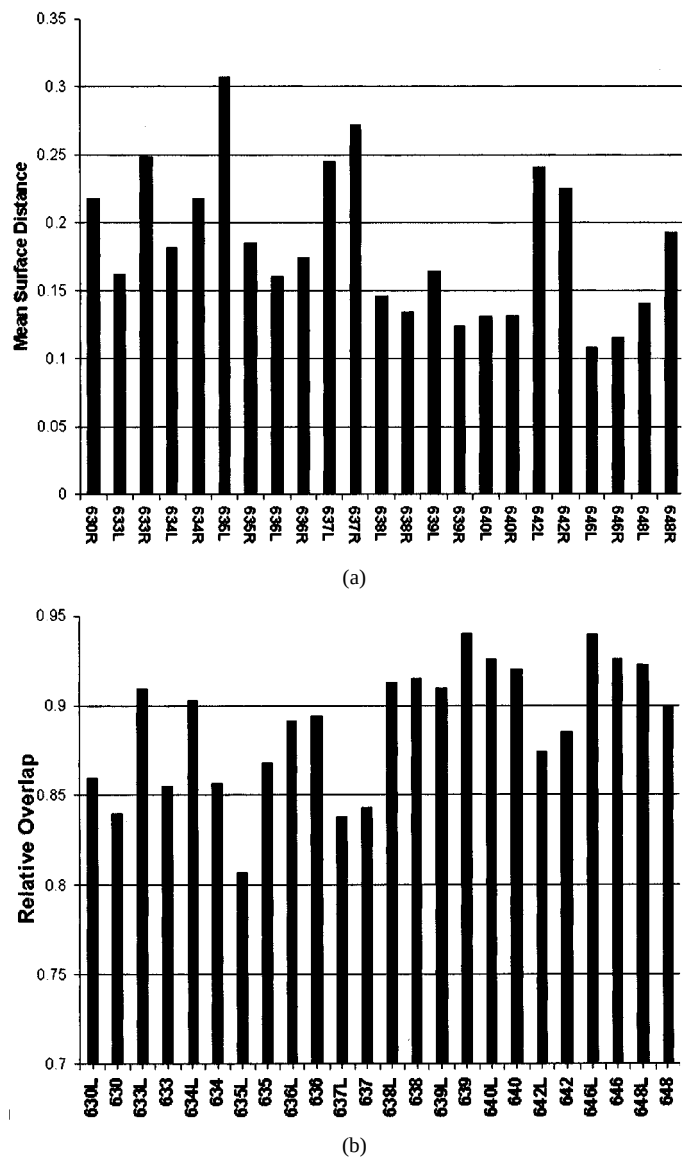

Fig. 7. The mean surface distance between the hand segmentation and the final stage of the hierarchical procedure is shown in (a). The relative overlap is shown in (b).

are in the process of implementing a statistical active appearance model following [8] to capture these intensity variations.

VI. Statistical Analysis of Shape

The goal of this work is not only to develop efficient representation and robust segmentation of individual objects, but also to develop statistical shape descriptions that capture anatomical variability across a population in an intuitive and localized fashion. The statistical shape representations that we have developed are based on the medial geometry and provide separate descriptions in terms of geometrical properties (bending or growth), scale (coarse or fine) and location (the whole object or a local region). We achieve this by building a set of medially derived features that are separable in these terms and are invariant under similarity transform. Application of principal component analysis to these features allows us to visualize and compare the different types of variability. In this section, we present our early results in two dimensions. We are currently generalizing the procedure to three dimensions.

We present results of analyzing the mid-sagittal cross sections of the corpus callosum of 71 subjects. The data consists of binary hand-segmented images divided into 31 normals and 40 schizophrenics. These images serve well for testing our method because the shape of the corpus callosum can be appropriately approximated with a single chain of medial atoms.

We now outline the basic segmentation procedure which is performed on each image in the population. Following the segmentation framework developed in the previous section, a coarse-scale single-figure 5-atom m-rep template of the corpus callosum is constructed manually. The template is deformed to fit a binary image with large tolerance producing a coarse scale segmentation M_0. The coarse m-rep M_0 is interpolated and resampled, forming a 9-atom *prediction m-rep* M_0', which is the initial estimate at the finer scale. M_0' is again deformed to fit the same image with lower tolerance, giving us the finer-scale segmentation M_1. The procedure can be repeated to produce yet finer scale descriptions $M_2 \cdots$

In order to establish correspondence, the distances between adjacent medial atoms are forced to be equal during segmentation of coarse m-reps. During the segmentation of fine m-reps, the medial atoms in the prediction m-reps are allowed to move only orthogonally to the direction of the medial manifold as given by $\mathbf{b}$.

A. Statistical Features

Having segmented the images, we now develop the statistical shape description that is fundamentally invariant under the similarity transform. In the literature, methods such as the Procrustes algorithm have been used to align objects in space to achieve this invariance [7], [13]. A drawback of these methods is that one is not able to get a truly local description of variability, as the alignment procedure itself is global in nature and based on all the data.

Instead of aligning m-reps, we compute a set of statistical features that intuitively describe shape and are both globally and locally invariant under the similarity transform. The features used are unit-free and magnification invariant, as they describe ratios between different object-related distances either explicitly or as sines or cosines of angles.

Two types of features are computed: *coarse features* describe the coarse m-rep M_0 and *refinement features* describe the differences between M_1 and M_0', thus capturing only the local fine scale information.

Coarse features capture shape information contained in medial atoms and in relationships between adjacent atoms. They, thus, describe shape in terms of either bending or local growth; growth refers to the widening or elongation of a part of the object relative to the rest. Of the five coarse features computed, the first two measure bending. The feature

$$f_i^1 = \frac{\|(\mathbf{x}_{i+1} - \mathbf{x}_i) \times (\mathbf{x}_i - \mathbf{x}_{i-1})\|}{\|(\mathbf{x}_{i+1} - \mathbf{x}_i)\| \, \|(\mathbf{x}_i - \mathbf{x}_{i-1})\|}, i = 2 \cdots n-1$$

describes how the figure bends at each of the interior medial atoms. The norm of the cross product gives the sine of the angle between the neighboring atoms.

The feature

$$f_i^2 = \begin{cases} \frac{\|\mathbf{b}_i \times (\mathbf{x}_{i+1} - \mathbf{x}_i)\|}{\|(\mathbf{x}_{i+1} - \mathbf{x}_i)\|} & i = 1 \cdots n-1 \\ \frac{\|\mathbf{b}_i \times (\mathbf{x}_n - \mathbf{x}_{n-1})\|}{\|\mathbf{x}_n - \mathbf{x}_{n-1})\|} & i = n \end{cases}$$

measures the angle between the vector $\mathbf{b}$ and the line joining the atom to its neighbor.

The remaining three features measure local growth. The feature $f_i^3 = r_i/L$, $i = 1 \cdots n$, captures the width at the medial atom, where $L = 1/(n-1) \sum_{i=1}^{n-1} \|\mathbf{x}_{i+1} - \mathbf{x}_i\|$ is the mean interatom distance. Feature $f_i^4 = \cos\theta$, $i = 1 \cdots n$, describes the change in width at the atom. The final coarse feature $f_i^5 = \eta$, $i = 1$ and n, is only defined at end atoms and measures the elongation of the object's ends.

To describe fine-scale local shape that is similarity transformation invariant, we base the fine-scale features on the differences between the coarse and fine m-reps. Since M_0 and M_0' imply the same boundary and M_0' and M_1 have an atom-to-atom correspondence, we use the differences between M_1 and M_0'.

The first two refinement features represent the displacement of the fine medial atom in M_1 with respect the corresponding atom in M_0'. Feature

$$g_i^1 = \frac{(\mathbf{x}_i - \mathbf{x}_i') \cdot (\mathbf{b}_i')^{\perp}}{L}, i = 1 \cdots n$$

where $\mathbf{x}_i'$, $\mathbf{b}_i'$ are elements of the medial atoms in M_0', measures the component of this displacement in the direction perpendicular to the medial axis of M_0'. For internal medial atoms, this component captures all of the displacement since these atoms are not allowed to move along the axis during segmentation. End atoms are allowed to move freely and feature

$$g_i^2 = \frac{(\mathbf{x}_i - \mathbf{x}'_i) \cdot (\mathbf{b}_i')}{L}, i = 1 \text{ and } n$$

is used to capture their movement along the $\mathbf{b}$ direction.

The remaining features measure differences in the direction of the medial axis ($g_i^3 = \|\mathbf{b}_i \times \mathbf{b}_i'\|, i = 1 \cdots n$), object angle ($g_i^4 = \sin(\theta_i - \theta_i'), i = 1 \cdots n$), width ($g_i^5 = (r_i - r_i')/L, i = 1, \ldots, n$) and end elongation ($g_i^6 = \eta - \eta', i = 1$ and n).

The above features completely capture the information in an m-rep which can, therefore, be reconstructed up to a similarity transform. Given the absolute position, orientation and radius of just one medial atom, we can use the information contained in the coarse features to rebuild the coarse m-rep M_0. Similarly, given a reconstructed m-rep M_0 and the refinement features, we can reconstruct M_1. The reconstruction algorithm is derived simply from the feature definitions; since the sampling distance is not one of the features, reconstruction relies on enforcing regular sampling of atoms during segmentation.

B. Statistical Analysis

Following Cootes *et al.* [7], we use Principal Component Analysis to define the shape variability. Let $\mathbf{y}_1, \ldots, \mathbf{y}_m$ be a collection of $p \times 1$ feature vectors. Each feature vector contains statistical features associated with a segmentation of one corpus callosum image. The feature vectors contain subsets of the features defined above. For example, for a statistical analysis of coarse scale bending, we use

$$\mathbf{z}_j = \left[f_2^1, \ldots, f_{n-1}^1, f_1^2, \ldots, f_n^2\right]^T.$$

Let $Z_{m\times p}$ be a feature matrix whose jth row is $\mathbf{z}_j^T$. For some integer k, the mapping from feature space to the k-dimensional coefficient space is given by the *k-dimensional reduction function* based on Z

$$\omega = \Omega_{Z,k}(\mathbf{z}) = \left[\frac{(\mathbf{z}-\bar{\mathbf{z}})\cdot \mathbf{e_1}}{\sqrt{\lambda_1}}, \ldots, \frac{(\mathbf{z}-\bar{\mathbf{z}})\cdot \mathbf{e_k}}{\sqrt{\lambda_k}}\right]^T \tag{10}$$

where $\bar{\mathbf{z}}$ is the mean of $\mathbf{z}_j$, $\lambda_1, \ldots, \lambda_n$ are the eigenvalues of Σ_Z—the covariance matrix of Z—in decreasing order and $\mathbf{e}_1, \ldots, \mathbf{e}_n$ are the corresponding eigenvectors of Σ_Z.

The corresponding mapping from coefficient space to feature space is given by

$$\mathbf{z} = \Omega_{Z,k}^{-1}(\mathbf{w}) = \bar{\mathbf{z}} + \sum_{j=1}^{k} w_j \mathbf{e}_j \sqrt{\lambda_j}. \tag{11}$$

The value k is chosen as the smallest integer that satisfies

$$\sum_{j=1}^{k} \lambda_j \geq \alpha \sum_{j=1}^{n} \lambda_j \tag{12}$$

where $\alpha \in (0,1)$ is the fraction of total variability that we want the coefficient space to capture. A typical value used for α is 0.95, discarding 5% of the shape variability, which we attribute to noise.

C. Corpus Callosum Classification

To discriminate between a pair of classes A and B, we use a technique based on the Fisher linear discriminant [14], a direction in feature space which under projection yields the the greatest separation between the means of two classes.

Fig. 8 shows the projections of the coarse m-rep features of each member of the schizophrenic class and each member of the normal class onto the Fisher linear discriminant. There is little separation between the classes.

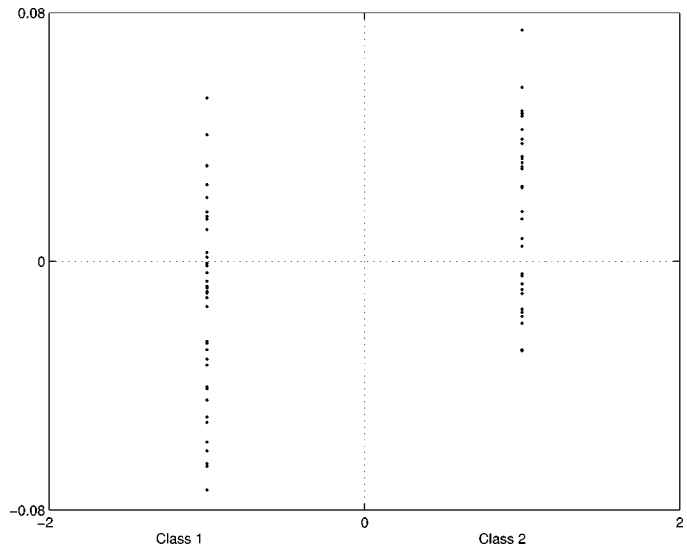

Fig. 8. Coarse features of schizophrenic and normal corpora callosa projected into the Fisher linear discriminant.

Leave-one-out applied to schizophrenic and normal corpora callosa, yields false-positive error of 0.52 and false-negative error of 0.49. The results of classification are negative, as are similar classification results reported by other researchers [17].

D. Visualization of Shape Variability

Each point in the coefficient space corresponds to a point in the feature space, which in turn corresponds to a shape that can be visualized. Important insight can be gained by reconstructing shapes that correspond to points sampled along one of the coordinate directions in coefficient space. Variability captured by the jth principal mode can be visualized by first finding the feature vectors

$$\mathbf{z}(t) = \Omega_{Y,k}^{-1}(t\beta_j) \quad (t = -2.0, -1.9, \ldots, 2.0) \tag{13}$$

where $\beta_1, \ldots, \beta_k$ are the unit basis vectors in coefficient space and then reconstructing the m-reps corresponding to $\mathbf{z}(t)$.

Alignment of the reconstructed object is necessary to produce a visually meaningful animation because the coarse statistical features lack information about the absolute placement of the m-rep. Thus, for visualization only, we apply the Procrustes algorithm [13] to the locations of the medial atoms belonging to all of the reconstructed coarse m-reps.

Fig. 9 summarizes the global variability in the coarse m-reps of the schizophrenic corpus callosum. It displays m-reps and implied boundaries corresponding to coefficient space points $-2\beta_1$, $\mathbf{0}$ and $+2\beta_1$ in Fig. 9(a) and $-2\beta_2$, $\mathbf{0}$ and $+2\beta_2$ in Fig. 9(b).

Fig. 10(a) summarizes the PCA performed on coarse features by plotting the eigenvalues of the covariance matrix and their cumulative sums. This plot shows that $k_{0.95} = 10$ dimensions are needed to capture 95% of the variability. Fig. 10(b) shows projections of the coefficient vectors ω_i onto the coordinate directions in coefficient space. Aside from a few outliers, training shapes lie within two standard deviations from the mean and their distribution is strikingly Gaussian.

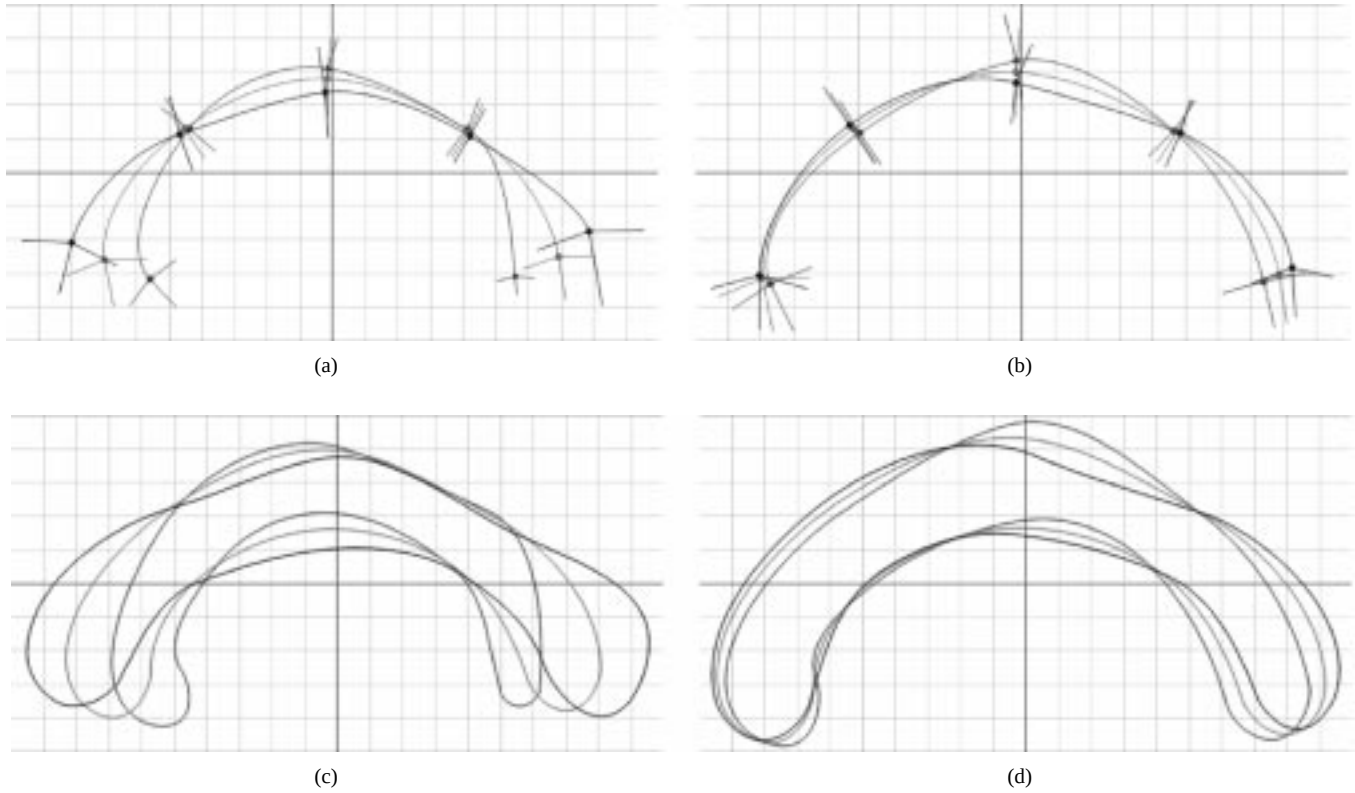

Fig. 9. Shapes reconstructed along the first two primary modes of variability in coarse features from -2σ to 2σ. All coarse features are combined in this figure, without separation into growth and bending. (a), (b) Shows the medial axis and (c), (d) shows the implied boundary.

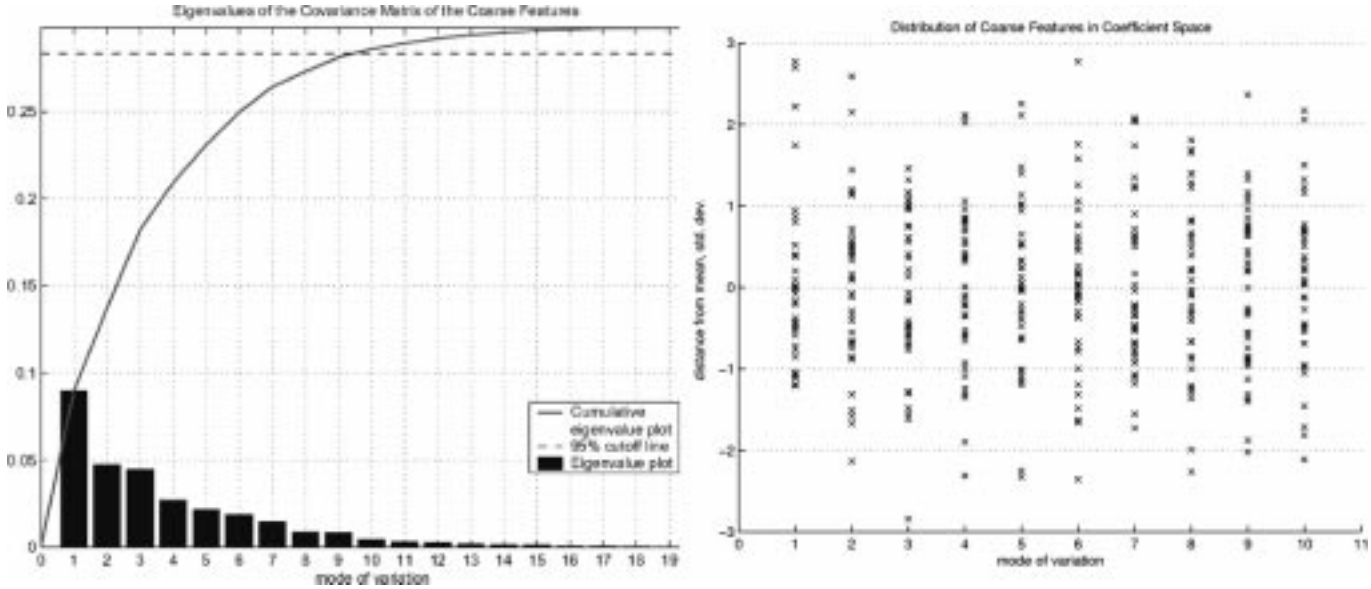

Fig. 10. Summary of PCA on coarse features. (a) Eigenvalues of the covariance matrix of features computed from coarse m-reps (schizophrenic corpus callosum). (b) Distribution of coarse shapes in PCA coefficient space.

The ability to separate variability into localized components and the ability to decompose global variability into bending and growth are strengths of our method. The bending information is captured by the first two features f^1 and f^2. Illustrated in Fig. 11 is the primary mode of variation resulting from applying the PCA to only these features. The change in bending is primarily explained by the fact that some corpora callosa are curved like a 'C' while others are straighter. In this and the subsequent two figures, the shapes at -2, 0 and $+2$ standard deviations along the mode are shown. The left panel shows medial geometry and

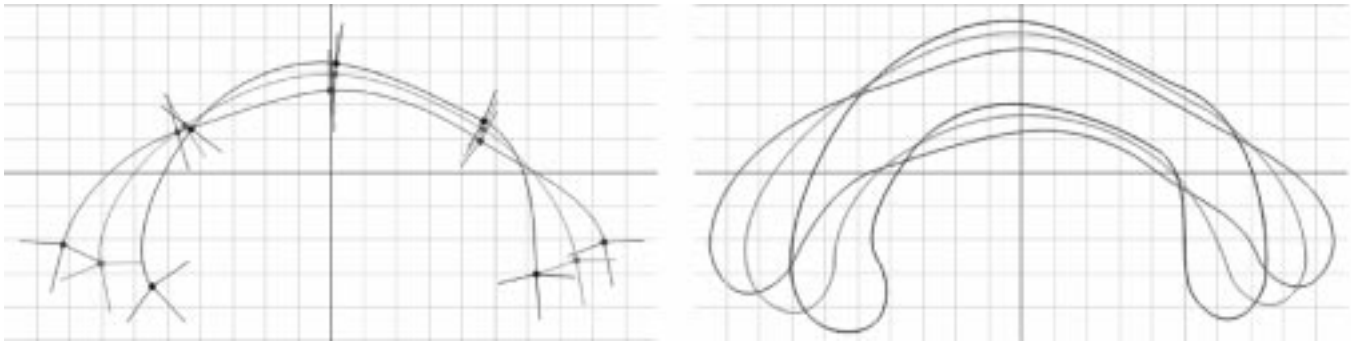

Fig. 11. First principal mode of variability computed using coarse features f^1 and f^2 which capture the bending of the corpus callosum in the population.

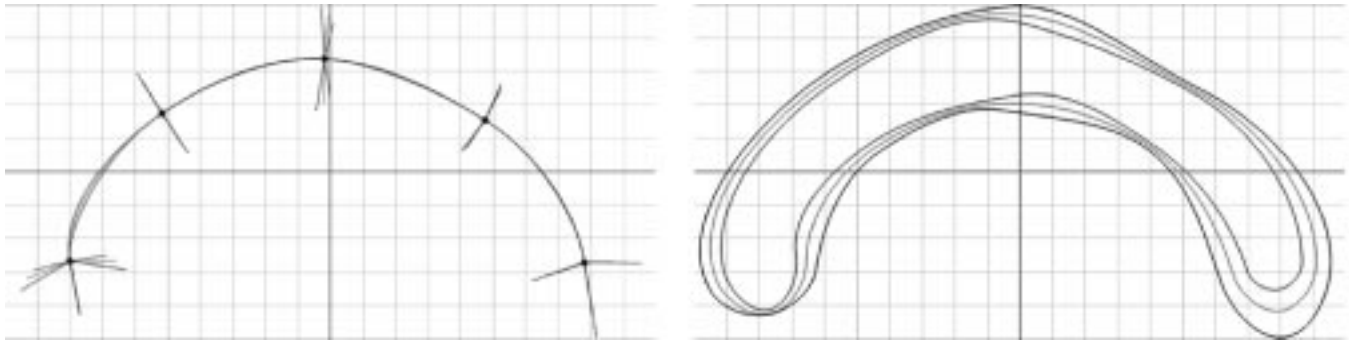

Fig. 12. First principal mode of variability computed using coarse features f^3, f^4, and f^5 which capture the growth of the corpus callosum in the population.

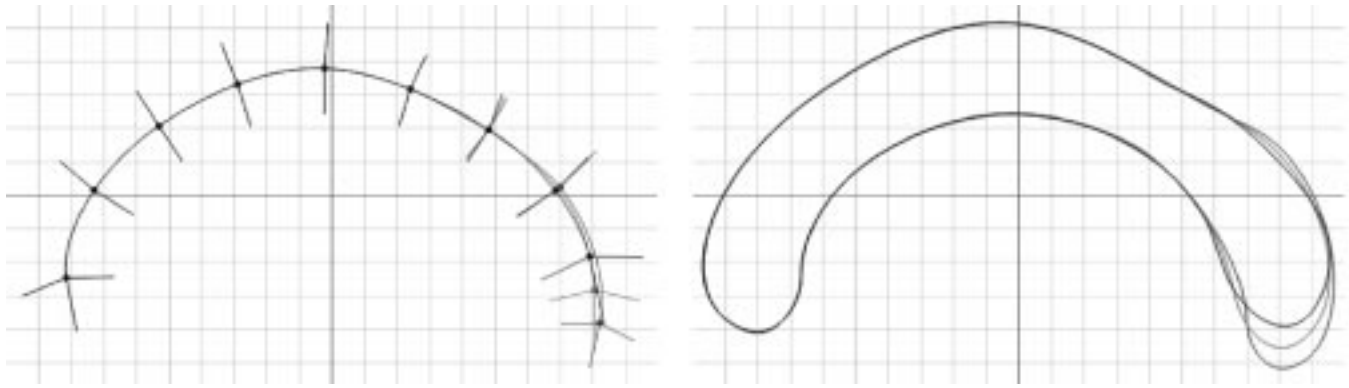

Fig. 13. First principal mode of variability computed using all fine features at the posterior of the organ which capture local fine-scale shape variability. The PCA is performed only on the features computed from the last two medial atoms. The mean on the population is used to reconstruct the remainder of the figure.

right panel shows the implied boundaries. Growth variability is captured by the features f^3, f^4, and f^5. Fig. 12 shows the first mode of variability associated with these coarse growth features. Notice that the primary component of growth variability is the uniform widening of the corpus callosum.

Fig. 13 shows the primary mode of variability resulting from applying the PCA to all the fine scale features associated with the last two medial atoms at the posterior of the corpus callosum. Here elongation is the main component of shape variability.

VII. Discussion and Future Work

It can be seen from the quantitative analysis of the segmentations that the accuracy of the automatic segmentation as measured via the average surface distance is on the order of the resolution of the imaging modality. All the results shown in this paper were generated using a Gaussian derivative template for the data-likelihood. Although these results show that our current methodology can segment structures in the abdomen such as the kidney with high level of accuracy, we expect that the results would be substantially improved by the use of a training image template in place of the Gaussian derivative template, thus allowing a spatially varying template to capture the different gray scale characteristics of the kidney boundaries. The model-to-image match would be further improved by a statistical model reflecting image-intensity variations across a population of subjects. We are currently extending this frame-work to the deformation of objects with multiple attached subfigures and multiple objects, with priors that reflect the knowledge of the associated relative typical geometry.

Another major contribution of this paper is the development of a shape analysis methodology that leverages the intuitive

and multiscale properties of the medial representation. Our statistical features have distributions that qualitatively appear Gaussian, although some features have outliers. Statistical methods that do not assume normality may further improve analysis.

In this paper, statistical methods have been applied to one level of detail at a time. Analysis on combined features appears promising because it can detect correlation between local variability at fine scale and global variability at coarse scale.

We are extending the statistical analysis method to three dimensions and multifigural complexes.

ACKNOWLEDGMENT

The authors would like to thank Dr. G. Gerig and M. Jomier for the use of their scoring tool for the comparison of segmentation as well as for the many insightful discussions and comments. They would also like to thank Dr. Z. Chen for the generating the table comparing the segmentations. Finally, they would like to thank Dr. E. Chaney for providing the kidney data sets and invaluable insights and Dr. Gerig and S. Ho for providing the corpus callosum data. They are grateful to Y. Fridman and G. Tracton for aid in m-rep segmentation. This research was carried out on computers donated by Intel.

REFERENCES

[1] I. Biederman, "Recognition-by-Components: A theory of human image understanding," *Psychological Rev.*, vol. 94, no. 2, pp. 115–147, 1987.

[2] R. Blanding, C. Brooking, M. Ganter, and D. Storti, "A skeletal-based solid editor," in *Proc. 5th Symp. Solid Modeling and Applications (SSMA-99)*, W. F. Bronsvoort, D. C. Anderson, and editors, Eds., New York, June 9–11, 1999, pp. 141–150.

[3] J. Bloomenthal and K. Shoemake, "Convolution surfaces," *Comput. Graph. (Proc. SIGGRAPH '91)*, vol. 25, no. 4, pp. 251–256, July 1991.

[4] H. Blum, "A transformation for extracting new descriptors of shape," in *Models for the Perception of Speech and Visual Form*. Cambridge, MA: MIT Press, 1967.

[5] A. C. Burbeck, S. M. Pizer, B. S. Morse, D. Ariely, G. Zauberman, and J. Rolland, "Linking object boundaries at scale: a common mechanism for size and shape judgments," Univ. North Carolina, Chapel Hill, Comput. Sci. Dept. Tech. Rep. TR94-041, 1996.

[6] B. M. Cameron and R. A. Robb, "An axial skeleton based surface deformation algorithm for patient specific anatomic modeling," in *Proc. Medicine Meets Virtual Reality 2000*, J. D. Westwood, Ed., 2000, pp. 53–58.

[7] T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models—Their training and application," *Comput. Vis., Graph. Image Processing: Image Understanding*, vol. 1, no. 61, pp. 38–59, 1994.

[8] T. F. Cootes, C. Beeston, G. J. Edwards, and C. J. Taylor, "A unified framework for atlas matching using active appearance models," in *Information Processing in Medical Imaging (IPMI)*. Berlin, Germany: Springer, 1999, vol. 1613, pp. 322–333.

[9] J. Csernansky, S. Joshi, L. Wang, J. Haller, M. Gado, J. Miller, U. Grenander, and M. Miller, "Hippocampal morphometry in schizophrenia via high dimensional brain mapping," in *Proc. National Academy of Sciences*, vol. 95, 1998, pp. 11 406–11 411.

[10] J. G. Csernansky, L. Wang, S. Joshi, J. P. Miller, M. Gado, D. Kido, D. McKeel, J. C. Morris, and M. I. Miller, "Early dat is distinguished from aging by high-dimensional mapping of the hippocampus," *Neurology*, vol. 55, pp. 1636–1643, December 2000.

[11] J. Damon, "Smoothness and Geometry of Boundaries Associated to Skeletal Structures," Dept. Math., Univ. North Carolina, Chapel Hill, Comput. Sci. Dept. Tech. Rep. TR94-041, 2001.

[12] C. Davatzikos, M. Vaillant, S. Resnick, J. Prince, S. Letovsky, and R. Bryan, "A computerized approach for morphological analysis of the corpus callosum," *J. Comput. Assist. Tomogr.*, vol. 20, pp. 207–222, 1995.

[13] I. Dryden and K. Mardia, *Statistical Shape Analysis*. New York: Wiley, 1998.

[14] R. Duda and P. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley, 1973.

[15] D. Fritsch, S. Pizer, L. Yu, V. Johnson, and E. Chaney, "Segmentation of medical image objects using deformable shape loci," in *Proc. Int. Conf. Information Processing in Medical Imaging*, 1997, pp. 127–140.

[16] U. Grenander, *General Pattern Theory*. Oxford, U.K.: Oxford Univ. Press, 1994.

[17] J. R. Highley, M. M. Esiri, B. McDonald, M. Cortina-Borja, B. M. Herron, and T. J. Crow, "The size and fiber composition of the corpus callosum with respect to gender and schizophrenia: A post-mortem study," *Brain*, vol. 122, pp. 99–110, 1999.

[18] M. Leyton, *Symmetry, Causality, Mind*: MIT Press, 1992.

[19] D. Marr and H. K. Nishihara, "Representation and recognition of the spatial organization of three-dimensional shapes," in *Proc. Roy. Soc. London Ser. B*, vol. 200, 1978, pp. 269–294.

[20] M. I. Miller, S. Joshi, D. R. Maffitt, J. G. McNally, and U. Grenander, "Mitochondria, membranes and amoebae: 1, 2 and 3 dimensional shape models," in *Statistics and Imaging*, K. Mardia, Ed. London, U.K.: Carfax, 1994, vol. II.

[21] L. R. Nackman, "Three-dimensional shape description using the symmetric axis transform," Ph.D. dissertation, University of North Carolina at Chapel Hill, Chapel Hill, NC, 1982.

[22] A. Pentland and S. Sclaroff, "Closed-form solutions for physically based shape modeling and recognition," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, pp. 715–729, July 1991.

[23] S. Pizer, D. Fritsch, P. Yushkevich, V. Johnson, and E. Chaney, "Segmentation, registration and measurement of shape variation via image object shape," *IEEE Trans. Med. Imag.*, vol. 18, pp. 851–865, Oct. 1999.

[24] A. Sherstyuk, "Shape design using convolution surfaces," presented at the Shape Modeling International '99, Aizu-Wakamatsu, Japan, Mar. 1999.

[25] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. W. Zucker, "The hamilton-jacobi skeleton," in *Proc. Computer Vision*, vol. 2, 1999, pp. 828–834.

[26] L. H. Staib and J. S. Duncan, "Model-based deformable surface finding for medical images," *IEEE Trans. Med. Imag.*, vol. 15, pp. 1–12, Oct. 1996.

[27] D. W. Storti, G. M. Turkiyyah, M. A. Ganter, C. T. Lim, and D. M. Stat, "Skeleton-based modeling operations on solids," in *SMA '97: Proc. 4th Symp. Solid Modeling and Applications*, Atlanta, GA, May 14–16, 1997, pp. 141–154.

[28] A. Yuille and P. Hallinan, *Active Vision*. Cambridge, MA: MIT Press, 1992.