# Shape Modeling and Image Visualization in 3D with M-rep Object Models 

P. Thomas Fletcher, Stephen M. Pizer, Andrew Thall, A. Graham Gash<br>Medical Image Display \& Analysis Group<br>University of North Carolina, Chapel Hill<br>\{fletcher, smp, thall, gash\}@cs.unc.edu


#### Abstract

We present a framework and geometry for object modeling in medical image analysis and display based on a sampled medial representation called an m-rep. This approach provides a coordinate system that is relative to the shape of the model-important because it allows the definition of more accurate visualization of 3D image data relative to a model. It also provides a means for describing the correspondence between points of two spaces or shapes. The methods described in this paper have been used successfully in 3D image segmentation by deformable m-reps.


## 1 Introduction

A common problem in medical image display and analysis is the visual comprehension of image intensity data relative to anatomic objects of interest. The difficulty of visualization is that the image slices required for greyscale presentation are not naturally aligned with the object.

Another problem in medical image display and analysis, arising in deformable model based segmentation and in image registration, is the correspondence of places in the image or on the object to places in a model of the situation, e.g., an atlas. Comparing a deformed model with the original model requires the correspondence of places on the object; comparing intensities between reference and target images in elastic image registration or between training and target images in deformable model based segmentation requires the correspondence of places in the image.

The solution to both of these problems is to define an object-relative geometry. The principal points of this paper are that

1. A useful geometry for these purposes can be obtained by (a) viewing an object as a set of linked figures (Fig. 1), (b) seeing that figures are defined and well comprehended in medial terms, and (c) realizing that the magnification invariance required by the objectives of local shape description leads to distances being relative to the local figural width-the medial radius or modeling aperture.
2. This geometry can be induced by a sampled medial representation called an m-rep that can conveniently be built based on a training image. Moreover, this model building is enabled by the capabilities of the object-relative geometry for providing useful visualizations.

As will be detailed below, a medial representation provides the following attractive properties in describing figural geometry:

1. It provides local information about the cross-figural direction, and it distinguishes directions along the figure according to the behavior of the figural width. It thereby provides a local, figurally-related compass.
2. The medial width provides a local ruler for the figure, inducing a magnificationinvariant geometry that provides helpful local shape characterizations.

## 2 Shape Representation with M-rep Models

The medial axis transform was first introduced in 2D and 3D by Blum [2], and its mathematical properties were developed in 2D by Blum \& Nagel [3] and in 3D by Nackman [7]. A continuous medial axis, defined precisely from the boundary of an object, is complex and difficult to represent or manipulate exactly. Also, small boundary perturbations can result in changes in the topology of the medial axis. M-rep models instead describe shape through a sampled medial description from which a boundary is derived rather than vice versa. If the medial sampling is maintained at an appropriately coarse scale, the medial model will imply a smooth boundary. Boundary detail may then be achieved by defining offsets to boundary points along the normal.

A simple figure, one that has a single, non-branching medial surface, can be considered as a manifold of medial atoms and can be represented as a mesh of medial atoms. The continuous manifold is then reconstructed from the samples in the mesh. More complex objects are built as a collection of connected figures.

### 2.1 Medial Atoms

From Blum's point of view, each point on the medial axis of a 3D object represents the center of a largest inscribing sphere to that object. From our point of view, in which the m-rep implies the boundary, a medial atom serves as a description of the relation between that sphere's center and the position and normal at the two boundary points (medial involutes) at which the sphere touches. It thus encapsulates information about the solid locally about a point on the medial axis. Except for points along the edge of the medial surface and points where the medial surface branches, an atom can be thought of as a position with two equal length vectors. These vectors are incident and normal to the surface at their ends, as shown in Fig. 2.


Fig. 1. A 2D schematic of an object consisting of a tree of figures (protrusions and indentations). Hinge atoms on the subfigures are circled, their implicit connections to the parent shown as dotted lines


Fig. 2. A medial atom with its implied boundary

More formally, a medial atom $\mathbf{m}$ of the space $\mathbb{M}=\left(\mathbb{E}^{3} \times S O(3) \times \mathbb{R} \times \mathbb{R}\right)^{1}$ can be defined as

$$
\begin{equation*}
\mathbf{m}=\{\mathbf{x}, \mathbf{F}, r, \theta\} \tag{1}
\end{equation*}
$$

where $\mathbf{x}$ is a position on the medial axis, $r$ is the radius of the inscribing sphere, $\theta$ is one half of the angle between the two vectors to the boundary, and $\mathbf{F}$ is a 3 D frame located at $\mathbf{x}$ and fitted to the geometry of the medial axis. The frame is a rotation of the standard Euclidean basis and thus an element of $S O(3)$. It can be thought of as being composed of three orthonormal vectors,

$$
\begin{equation*}
\mathbf{F}=\left\{\vec{b}, \vec{b}^{\perp}, \vec{n}\right\} \tag{2}
\end{equation*}
$$

These vectors define the local compass for the medial surface at the point $\mathbf{x}$, where $\left\{\vec{b}, \vec{b}^{\perp}\right\}$ span the tangent plane at $\mathbf{x}$, and $\vec{n}$ is its normal. Moreover, the vector $\vec{b}$ is chosen as the bisector of the two vectors pointing to the boundary. This is the direction of greatest narrowing of the implied solid, given by the relation

$$
\begin{equation*}
\nabla r=-\vec{b} \cdot \cos \theta \tag{3}
\end{equation*}
$$

Therefore, the vector $\vec{b}^{\perp}$ is in the direction of constant $r$, i.e., no widening or narrowing. The two vectors pointing to the boundary can be defined as

$$
\begin{equation*}
(\vec{s}, \vec{p})=(r(\vec{b} \cdot \cos \theta+\vec{n} \cdot \sin \theta), \quad r(\vec{b} \cdot \cos \theta-\vec{n} \cdot \sin \theta)) \tag{4}
\end{equation*}
$$

From Blum's point of view, on the boundary of the medial manifold the maximal inscribing sphere will no longer touch the boundary of the object at

[^0]exactly two points but rather touches at one point of multiplicity 2 . At that point of contact, a crest position on the implied boundary, the corresponding medial atom will have an angle $\theta$ equal to zero. Such an atom is ill suited to image analysis due to the instability of its dependence on a single image point. We therefore define an end atom that implies a crest position at
\[

$$
\begin{equation*}
\mathbf{p}_{e}=\mathbf{x}+\eta r \vec{b}, \tag{5}
\end{equation*}
$$

\]

with $\vec{b}$ as surface normal. The parameter $\eta$ defines the pointiness, i.e, elongation of the crest region. This elongation factor is restricted to lie within $[1,1 / \cos \theta]$. When $\eta=1$, the end atom implies a circular cross-section; $\eta=1 / \cos \theta$ yields a sharp corner (Fig. 3).


Fig. 3. Cross-sections of three end atoms. The atom on the left has an elongation of $\eta=1$, the atom in the middle has maximal elongation, $\eta=\frac{1}{\cos \theta}$, while the rightmost atom has an elongation in between

In the case of a branching figure, from Blum's point of view the inscribing sphere is incident to three points on the boundary. From our point of view, the connecting section of medial axis between a parent figure's medial position and the place on its boundary where the subfigure attaches is redundant and not helpful. We therefore begin a protrusion with atoms, called hinge atoms, on the surface of the parent figure (Fig. 1). This usage also frees us to describe indentation subfigures in a similar way.

Typically, one figure can be identified as the main figure of an object. This main figure may have several protrusion or indentation subfigures attached to it. In turn each subfigure may have it's own protrusion and indentation subfigures, and so forth, down to the desired scale. This suggests a natural data structure for an object: a tree of figures, where the main figure is the root. Since some objects may have several non-overlapping main figures (for example, the brain might be modeled as two hemispheres), m-rep objects are a forest of trees of figures, allowing several main figures to be defined.

### 2.2 Reconstructing the Medial Manifold from Meshes of Medial Atoms

For the generic case of a slab figure, the locus of medial atoms forms a continuous manifold with a bounding 1D manifold of end-atoms. In an m-rep, only a discrete

2D sampling of the medial manifold is defined, and the end-atom manifold is sampled as well.

We need to produce a sampling that is locally magnification-invariant; this implies that the sampling intervals of the medial mesh should be proportional to the radius $r$ in the sampled neighborhoods. While there are no theoretical restrictions on the connectivity of the discrete mesh, a simple quadrilateral mesh (quadmesh) allows convenient sampling of the medial manifold and reconstruction from the samples. The quadmesh has $M$ columns and $N$ rows, with $M$ and $N$ chosen in proportion to the ratio of the two dimensions of the figure along the medial manifold. The basic mesh structure is indexed by $u, v$, where column $u$, row $v$, and the medial normal direction $\vec{n}$ have a right-handed order. Quadmeshes allow simple indexing and straightforward parameterization and allow for splineor subdivision-based interpolation. They have proven adequate for the modeling tasks and classes of deformation used thus far.

Given the discrete, integral $(u, v)$ parameterization of the mesh, we begin reconstruction of a continuous manifold of medial atoms by fitting piecewise $\mathbf{C}^{2}$ Bézier spline patches $\mathbf{x}(u, v)$ for each quadrilateral of the grid which are adjusted to be $\mathbf{C}^{1}$ along boundaries. These patches are based on the known medial positions and surface normals at the patches' corners. The parameterization $(u, v)$ of the patch covers the unit square bounded by the integer indices of the quadrilateral. This gives the $\mathbf{x}$ and $\vec{n}$ components of interpolated medial atoms.

The next step of the reconstruction is the interpolation of $r$ and $\nabla r$ on each patch. This is accomplished by bicubic interpolation of $r$ on the patch's $(u, v)$ parameter space, using a pullback of $\nabla r$ to the parameter space for derivative information (and assuming that the twist $d u d v=0$ ). Following interpolation, the new $\nabla r$ is pushed forward onto the surface $\mathbf{x}(u, v)$. The pullbacks and pushes are allowed by the $\mathbf{C}^{1}$ continuity within and between the patches. The frame $\mathbf{F}$ and the object angle $\theta$ on the patch are derived by the relation (3) and the fact that the frame element $\vec{n}$ is normal to $\mathbf{x}(u, v)$.

## 3 Figural Coordinates

The ability to reference points in space is essential to visualization and modeling applications. Most commonly this is accomplished by representing points in a coordinate system with fixed axes. Although this method may be familiar, a much more natural method would be to reference space in a manner which is somehow relative to the shape in question. For example when discussing a model of a human hand, the reference to "the point at 3.5 cm . in $x, 6.3 \mathrm{~cm}$. in $y$, and 8.2 cm . in $z$ " is less intuitive than the description of "the point on the back of the index finger halfway between the second and third knuckles." Moreover, if the index finger is bent, the former description becomes invalidated, while the latter still specifies a locatable position on the model.

A figurally defined coordinate system will define a correspondence between variations of a shape. That correspondence will be locally magnification-invariant
if it uses a medially defined coordinate system and distances which are proportional to the width of the object.

### 3.1 Definition of the Coordinate System

We define a coordinate system for referencing the space around a figure in a manner that is relative to the figure itself. Consider a single figure with medial manifold $\mathcal{M} \subset \mathbb{M}$. We use the coordinate system $(u, v)$ defined by a quadmesh in Sect. 2.2 to provide a locally magnification-invariant parameterization $\mathbf{f}_{\mathcal{M}}: U \rightarrow$ $\mathcal{M}$, where $U \subset \mathbb{E}^{2}$. We can then reference space around the figure by associating each point in space with a point on the medial manifold.

We first define a mapping from points on the medial surface to corresponding points on the boundary of the figure. Since a medial atom implies two boundary points, we need an extra parameter $t \in\{-1,1\}$ to distinguish between the two points. The mapping a : $(\mathcal{M} \times\{-1,1\}) \rightarrow \mathcal{B}$, where $\mathcal{B} \subset \mathbb{E}^{3}$ is the boundary of the figure, is defined as

$$
\begin{equation*}
\mathbf{a}((\mathbf{x}, \mathbf{F}, r, \theta), t)=\mathbf{x}+r(\cos (\theta) \cdot \vec{b}+t \sin (\theta) \cdot \vec{n}) \tag{6}
\end{equation*}
$$

Also, the surface normals of the boundary positions are easily derived from the medial atoms. They are found using the mapping $\vec{q}:(\mathcal{M} \times\{-1,1\}) \rightarrow \mathbb{R}^{3}$, which is defined by

$$
\begin{equation*}
\vec{q}((\mathbf{x}, \mathbf{F}, r, \theta), t)=\cos (\theta) \cdot \vec{b}+t \sin (\theta) \cdot \vec{n} . \tag{7}
\end{equation*}
$$

Using the mapping $\mathbf{b}$, the boundary may be referenced in medial terms. A parameterization of the boundary $\mathbf{f}_{\mathcal{B}}:(U \times\{-1,1\}) \rightarrow \mathcal{B}$, which coincides with the given medial parameterization, now arises as

$$
\begin{equation*}
\mathbf{f}_{\mathcal{B}}(u, v, t)=\mathbf{a}\left(\mathbf{f}_{\mathcal{M}}(u, v), t\right) \tag{8}
\end{equation*}
$$

Now, for a point $\mathbf{x} \in \mathbb{E}^{3}$ there is a point $\mathbf{y}$ on the boundary which is closest to it. That is, the Euclidean distance is minimal: $d(\mathbf{x}, \mathbf{y}) \leq d(\mathbf{x}, \mathbf{z})$ for all $\mathbf{z} \in \mathcal{B}$. Although the point $\mathbf{y}$ may not be unique, there exists a region $V \subset \mathbb{E}^{3}$ containing $\mathcal{B}$ where all points do have a unique closest point on the boundary. If $\mathbf{g}: V \rightarrow \mathcal{B}$ is the mapping that carries each point of $V$ to its nearest point on the boundary, then a correspondence to the medial surface is $c: V \rightarrow \mathcal{M}$, where

$$
\begin{equation*}
\mathbf{c}=\pi_{\mathcal{M}} \circ \mathbf{a}^{-1} \circ \mathbf{g} \tag{9}
\end{equation*}
$$

where $\pi_{\mathcal{M}}:(\mathcal{M} \times\{-1,1\}) \rightarrow \mathcal{M}$ is the projection mapping.
In the figural coordinate system a point is referenced by the parameterization to the boundary, $\mathbf{f}_{\mathcal{B}}$ plus a signed distance from that boundary point, with negative distances for the interior of the figure and positive distances for the exterior. However, since we want a coordinate system that is locally magnificationinvariant, this distance must be defined in terms of the scale of the figure, that is, the radius of the corresponding medial atom. Thus, we define a distance
$\hat{d}_{\mathcal{B}}:(\mathcal{B} \times V) \rightarrow \mathbb{R}$ that is simply the Euclidean distance divided by the corresponding medial radius:

$$
\begin{equation*}
\hat{d}_{\mathcal{B}}(\mathbf{x}, \mathbf{y})=\frac{ \pm d(\mathbf{x}, \mathbf{y})}{\mathbf{c}_{r}(\mathbf{y})}, \tag{10}
\end{equation*}
$$

where $\mathbf{c}_{r}$ is just the $r$ component of the mapping $\mathbf{c}$, and the sign of $d$ is positive outside the figure and negative inside the figure.

Now we are ready to define our figural coordinate system for points in $V$ which coincides with the given medial parameterization and is magnificationinvariant. We produce $\mathbf{f}_{V}:(U \times\{-1,1\} \times \mathbb{R}) \rightarrow V$ by the equation

$$
\begin{equation*}
\mathbf{f}_{V}(u, v, t, \hat{d})=\mathbf{f}_{\mathcal{B}}(u, v, t)+\hat{d} \cdot r(u, v) \cdot \vec{q}(u, v, t) \tag{11}
\end{equation*}
$$

where $r(u, v)$ is the $r$ component of the given medial parameterization $\mathbf{f}_{\mathcal{M}}$.
Also, for a point in $V$ we can find its figural coordinates by the inverse mapping $\mathbf{f}_{V}^{-1}: V \rightarrow(U \times\{-1,1\} \times \mathbb{R})$, which is

$$
\begin{equation*}
\mathbf{f}_{V}^{-1}(\mathbf{v})=\left(\mathbf{f}_{\mathcal{B}}{ }^{-1}(\mathbf{v}), \hat{d}_{\mathcal{B}}(\mathbf{g}(\mathbf{v}), \mathbf{v})\right) . \tag{12}
\end{equation*}
$$

### 3.2 Implementation of the Figural Coordinate System

The figural coordinate system defined above is difficult to implement exactly. The primary reason for this is that the distance from a point to an arbitrary surface, needed to find $\hat{d}$, is difficult to compute exactly. Thus, a discrete approach is taken in the implementation. First, an axis-aligned bounding box is found for the volume of interest around an object. This space is divided into uniform voxels. Next, the boundary of the object is reconstructed using the interpolation method described in Sect. 2.2, and each voxel that contains a boundary point is labeled with the $(u, v, t)$ coordinate of that boundary point. Now we can compute the width-proportional distance $\hat{d}$ from the boundary for each voxel. The Euclidean distance from the boundary is approximated for each voxel through the Danielsson distance map [4]. For each voxel this map gives both the distance to the boundary and also which boundary voxel is closest. Since each boundary voxel is labeled with a $(u, v, t)$ coordinate, we know the medial atom to which it corresponds. We can then divide the Euclidean distance by that medial atom's radius to get the $\hat{d}$ value for the voxel. We now label the voxel with its full ( $u, v, t, \hat{d}$ ) coordinate.

### 3.3 Blending Multiple Figures

When objects consisting of multiple figures are used, the figural coordinates at a point correspond to the closest figure to that point. To produce smooth transitions at the seams between figures (Fig. 4), we take an implicit approach by blending the distance function of a subfigure with the distance function of its parent. The distance functions of the two figures are $r$-normalized, giving a final blended distance function that is also proportional to the radius.


Fig. 4. A 3D m-rep object with a protrusion and indentation figure; the sampled medial representation with hinge atoms highlighted (left), and the implied boundary without blending (middle) and with blending (right)

Coordinates for the blended regions are given by the parameters $\left(u^{\prime}, w, t\right)$, where $u^{\prime} \in[0,1]$ parameterizes the curve of hinge atoms, $w \in[-1,1]$ goes between the two figures in the blend, and $t \in\{-1,1\}$ again differentiates between the two sides of the figure.

To obtain the distance-to-boundary for the union of a parent figure $\mathcal{F}_{1}$ and its intersecting subfigure $\mathcal{F}_{2}$, a blending region is first defined via an extent parameter $T$. This parameter gives the width-proportional distance along each figure's boundary over which the blend is applied. Thus, in the implementation the blending region is taken as all voxels for which the two figures' widthproportional distance functions $\hat{d}_{1}$ and $\hat{d}_{2}$ satisfy the relation

$$
\begin{equation*}
\left|\hat{d}_{1}-\hat{d}_{2}\right|<T \tag{13}
\end{equation*}
$$

For each voxel in the blending region the value $w \in[-1,1]$, defined as

$$
w= \begin{cases}\left(\hat{d}_{1}-\hat{d}_{2}\right) / T & \text { for } \mathcal{F}_{2} \text { a protrusion subfigure }  \tag{14}\\ \left(-\hat{d}_{1}-\hat{d}_{2}\right) / T & \text { for } \mathcal{F}_{2} \text { an indentation subfigure }\end{cases}
$$

tracks across the blend region. Outside of the blend region the distance is defined to be the minimum of the two figures' distances. Inside the blend region this distance to the closest figural boundary is diminished by a correction function $C(|w|)$ that depends upon a second parameter, $s \in[0,1)$, controlling the smoothness of the blend. The correction $C$ is a cubic defined by the following constraints:

1. At $|w|=1$, where the correction stops, $C=0$ and $C^{\prime}=0$.
2. At $|w|=0$, i.e., at the $r$-normalized bisector between the two surfaces, $C$ has the value determined by the smoothness and extent parameters and $C^{\prime}$ has the value making $\hat{d}$ first-order continuous there.

This yields the function

$$
\begin{equation*}
C(|w|)=\frac{T}{2}\left[s\left(1-3|w|^{2}+2|w|^{3}\right)-|w|+2|w|^{2}-|w|^{3}\right] \tag{15}
\end{equation*}
$$

giving a distance value for the voxels in the blend region as

$$
\begin{equation*}
\hat{d}=\min \left(\hat{d}_{1}, \hat{d}_{2}\right)-C(|w|) . \tag{16}
\end{equation*}
$$

Thus, the figural coordinates inside of a blend region are the $\left(i, j, u^{\prime}, w, t, \hat{d}\right)$, where $\hat{d}$ is the blended distance value, and $i$ and $j$ are the figure indices. Outside of blend regions, the figural coordinates become $(i, u, v, t, \hat{d})$, where the parameters are those from the single figure case (Sect. 3.1) relative to figure $\mathcal{F}_{i}$.

This blending strategy gives a useful estimation to the true boundary distances in the blend regions. The implicit blending technique is simple to compute, and the storage requirement is similar to the single figure case. Only an extra figure index is required; otherwise both coordinate systems have the same number of parameters.

A common problem with implicit blending techniques, which this method also suffers from, is unwanted blending. That is, when two connected figures come within close proximity of each other away from the hinge, they will be blended.

## 4 Building Models

There are two main categories of shape operations that one would like to possess in a modeling system: global and local transformations. It is important that a shape representation allow these transformations in a manner that satisfies the following requirements:

1. Operations are intuitive and easy to perform from a user's perspective.
2. Mathematically, operations are simple enough to implement in real time.

M-rep models fulfill both of these requirements for both global and local shape transformations.

### 4.1 Global Shape Transformations

For most applications the most essential global transformations are the similarity transformations: translation, scaling, and rotation. Each of these can be implemented for m-rep objects as follows:

1. Translation of an m-rep object by a vector $\vec{v}$ is performed by translating each medial atom's position by $\vec{v}$.
2. Scaling an object about the origin is performed by scaling the vectors to each medial position along with scaling the radius of each medial atom.
3. Rotation of an object around an axis $\vec{v}$ by an angle $\theta$ is performed by rotating each medial atom's position and frame around $\vec{v}$ by $\theta$.

More formally, a generic similarity transform $S=(\alpha, \mathbf{O}, \vec{T}) \in\left(\mathbb{R} \times S O(3) \ltimes \mathbb{R}^{3}\right)$ transforms each medial atom $\mathbf{m}=\{\mathbf{x}, \mathbf{F}, r, \theta\}$ by the rule

$$
\begin{equation*}
S \mathbf{m}=(\alpha \mathbf{O} \mathbf{x}+\vec{T}, \mathbf{O} \circ \mathbf{F}, \alpha r, \theta) \tag{17}
\end{equation*}
$$

Most representations allow similarity transformations. The advantage of using m-rep objects becomes apparent with nonlinear deformations. M-reps provide a natural means for defining tapering, bending, twisting, and thickness scaling (fattening or thinning). Such deformations have previously been available through methods which warp the space in which an object is embedded $[1,9]$. However, space warping algorithms do not provide object-aligned deformations; the warping is done along a fixed axis or with respect to user-defined control points. A medial representation, on the other hand, provides an object axis along which the deformation can be defined. Not only does this provide a more correct orientation for deformations, it is also more intuitive from a user's perspective. Predicting the effects a space warp will have on an object can be difficult. It is much easier, especially for complex objects, to envision deformations that follow the object's axis.


Fig. 5. Nonlinear deformations on an m-rep slab figure (top left). M-reps allow easy definitions of bending (top middle), tapering (top right), twisting (bottom left), fattening (bottom middle), and thinning (bottom right)

Nonlinear deformations on m-reps are typically defined by simple operations that are propagated along a direction on the medial surface. A direction is defined by a family of non-intersecting curves, $\mathbf{m}_{i}:[0,1] \rightarrow \mathcal{M}$, that lie on the medial surface. When using quadrilateral meshes, it is easiest to define these direction curves as following the columns or the rows of sample atoms, although arbitrary directions are also possible. Each operation is defined continuously but is, of course, only applied to the sample medial points in practice.

1. Tapering. Along each direction curve, the radius is scaled by a monotonic decreasing function. The object angle, $\theta$, and the frame, $\mathbf{F}$, of each medial atom must also be adjusted to preserve the $\nabla r$ relationship of (3).
2. Bending. At each position on the direction curve, the frame $F$ is rotated. All medial positions lying further along on the direction curve are also rotated with the frame.
3. Twisting. A curve that is perpendicular to each direction curve is chosen, call it $\mathbf{s}:[0,1] \rightarrow \mathcal{M}$. Also, a monotonic function $\phi:[0,1] \rightarrow \mathbb{R}$ is used to describe the amount of twist along the curve $s$. Given a direction curve $\mathbf{m}_{i}$, it intersects $\mathbf{s}$ at some point $\mathbf{s}(t)$. Both the frame and position of each medial atom in $\mathbf{m}_{i}$ are rotated by the angle $\phi(t)$ about the tangent vector of $\mathbf{s}$ at $t$.
4. Thickness scaling. Each medial atom's radius is scaled by a parameter $\alpha \in$ $\mathbb{R}^{+}$, but distances between atoms remains the same. This results in an object which is fatter $(\alpha>1)$ or thinner $(\alpha<1)$.

### 4.2 Local Shape Transformations

Taking advantage of the sampling of the medial axis, local operations on m-rep models are simple extensions of the global operations defined above. Each medial atom influences a local region of the solid by the same algorithm as its global counterpart. Thus, a medial atom can be thought of as a control point for it's portion of the model. Translating an atom results in shifting a local piece of the model. Scaling an atom's radius creates a local bulge or dip in the surface. Also, a set of contiguous atoms may be grouped together, with operations affecting the group. This allows deformations of larger portions of the solid.

### 4.3 The Model Building Process

M-rep models are built within a graphical modeling program. Modeling operations are performed through easy-to-use mouse interactions. There are two guides to aid in the building of an object from a 3D medical image. The image itself can be used as a reference (using the visualization method described in Sect. 5). Also, a polygonal boundary model of the pre-segmented object may be loaded and viewed against the m-rep model. The general steps of constructing an object proceed as follows:

1. Divide the object into main figures and subfigures. For intuitiveness to users of medical imaging, figures should correspond to identifiable anatomy.
2. Choose the mesh size for each figure. The number of rows and columns in the quadmesh of medial atoms must be chosen so that the sampling may be width-proportional. That is, medial sampling should be coarser in thick regions and finer in thin regions.
3. Roughly shape the individual figures. Using the global transformations of Sect. 4.1, the figure is roughly placed and its general shape is modeled. A polygonal boundary model of the object is a useful guide in this stage.
4. Attach subfigures. Choose the sequence of end atoms along the edge of the subfigure's medial surface that form the hinge. After locating the subfigure approximately relative to the parent, these hinge atoms are automatically repositioned to their nearest points on the boundary of the parent figure.

The $r$-normalized Danielsson distance map described in Sect. 3.2 is used to determine the closest boundary points. Then the subfigure is added to the tree data structure as a child node of its parent.
5. Fine tune each figure. With the local modeling operations described in Sect. 4.2 , the final shape of a figure is sculpted. When building objects to model image data, single atoms are fine tuned with the aid of the visualization method detailed in Sect. 5. The segmentation method described in [6], with the hand-built model as the prior, can be used to deform the model further.

These steps may be repeated and interchanged. For instance, figures may be added at any time and the fine tuning may be done before attaching a subfigure.

## 5 Visualization of Images by Medially-Defined Cut Planes

When working with models in medical images, one often needs to visualize the model within the image data. This visualization may be necessary for surgical or radiation treatment planning, for building a model, or for evaluating or manually editing a segmentation. It is important to have a visualization method that accurately and intuitively displays the spatial relationships of the model and the image data.

Visualization of 3D image data traditionally entails viewing the original 2D slices of the image or slices orthogonal to them. Such a visualization can be misleading when considering objects that are not aligned with the image planes or that curve and twist through space. Spatial relationships are distorted because the image planes intersect surfaces at shear angles. We create image planes that are oriented relative to the object being viewed, thus providing a more accurate visualization of the image data.

Medial atoms provide a local frame, defined in (2), that provides an ideal orientation for local object-relative visualization. Given a medial atom, we define four planes for visualizing image data. These planes can be displayed both in 3 D and 2D. The 3D view displays the image cut planes with the m-rep model to give a perspective of how they are related in size and orientation. In 2D the image data through the plane is displayed with the atom and the intersection of the boundary with the plane.

As illustrated in Fig. 6, the first of these planes is the atom plane, which is the plane spanned by the atom's frame vectors $\vec{b}$ and $\vec{n}$. Also contained in this plane are the two vectors incident to the boundary, $\vec{p}$ and $\vec{s}$. Since, the vectors $\vec{p}$ and $\vec{s}$ are normal to the boundary of the object, this plane gives an accurate view of the image data and model at these points. For atoms on the edge of the medial surface, the end point $\mathbf{p}_{e}$, see (5), with $\vec{b}$ as the surface normal at that point, is also viewed accurately in this plane.

The next two planes are the port and starboard sail planes, named after the two sides of a ship since they land on either side of the medial atom. The port sail plane is spanned by the vectors $\vec{p}$ and $\vec{b}^{\perp}$, while the starboard sail plane is


Fig. 6. Medially-defined image cut planes in 3D on an m-rep kidney model showing the atom plane and crest plane (left), atom plane and sail planes (right)


Fig. 7. These cut planes in 2D are the atom plane (left top), crest plane (left bottom), and sail planes (right)
spanned by the vectors $\vec{s}$ and $\vec{b}^{\perp}$. These two planes give a view that also includes the points of contact with the boundary for the atom. They are perpendicular to the atom plane, and thus together give enough information to determine whether the surface coincides with the image data locally.

One more plane is needed for atoms along the edge of the medial sheet. The crest plane is the plane spanned by the vectors $\vec{b}$ and $\vec{b}^{\perp}$. This plane gives a perpendicular view to the atom plane of the crest of the figure.

In each of these image cut planes, the plane is perpendicular to the boundary of the model. This means that the medial cut planes will accurately portray local spatial relationships with the image since distances are viewed perpendicular to the surface. Since we can interpolate atoms at any point on the medial surface (Sect. 2.2), this visualization is available anywhere on the model.

## 6 Correspondence via Figural Coordinates

Representation via m-reps provides locally magnification-invariant correspondences between locations relative to two objects with the same m-rep topology but differing geometry. It accomplishes this by associating locations with the same figural coordinates (Fig. 8). While these correspondences can be leveraged in a variety of areas of medical image analysis, including registration and in comparing gold standard segmentations to those provided by a computer based method under study, we give here examples of the usefulness of this approach in deformable model segmentation and statistical discrimination of classes via object shape.

In [6] and [8] we describe means of segmenting an object from by deforming a reference m-reps model, i.e., an atlas. Model-based segmentation requires the measurement of a prior given by the geometric deviation of the deformed model


Fig. 8. Figural coordinates make correspondence of points between two models possible in regions outside the objects (A), inside the objects (B), and on the boundary of the objects (C)
from the original model. Figural coordinates provide both correspondences and $r$-proportional distances that can be used to give an effective measure for this deviation. The boundary points for the original model and their figural coordinates are precomputed, so that correspondence can be computed efficiently.

Model-based segmentation also requires the measurement of the match of the target image relative to the object to a template defined relative to the original model. For example, the template might be derived from one or more training images registered via the model $[6,10]$, or it might be defined in terms of Gaussian derivatives located, oriented, and scaled relative to the model [6]. Whether the image match measure is based on correlation, mutual information, or PCA-based statistical deviation, figural coordinates provide correspondences between template and target image intensities that give an effective image match. Precomputation of the voxelized figural coordinate map leads to an efficient implementation.

The success of these correspondences for both the prior and image match terms can be judged from the segmentations reported in [6], together with the fact that the 3D segmentations reported there could be produced typically in 5 minutes of computing on a Pentium II PC.

Statistical classification of objects based on their shape requires the matching of features among training objects by object-relative location. It also requires a consistent description of these features on target objects. Such correspondences have proven successful in early studies of multiscale object classification [11].

## 7 Discussion and Conclusions

We have proposed and implemented an m-reps based geometry that allows effective object-relative image visualizations and provides object-relative spatial correspondences to be made that are useful in a variety of image analysis tasks requiring such correspondence between the same object in two different elastic forms. A number of concerns and desired improvements remain.

The spatial correspondences become less and less credible, the farther one gets from the object. Moreover, exterior to nonconvex boundary positions there is guaranteed to be a distance beyond which the figural coordinates are discontinuous. Means of handling these issues remain to be discovered. In addition, the inclusion of additional homology information, such as landmark correspondence, would be valuable.

The implementation of the figural coordinate system involves voxel maps of distances and implicit blending on those distances. The resulting distances in the blend region do not necessarily lead to a consistent metric. Despite the computational costs one might try to reconstruct blended medial surfaces first by interpolating the branching medial surfaces between figures.

Hexagonal simplex meshes [5], with their related triangle meshes, could be used in place of the quadmeshes used to sample the medial manifold of a figure. This would allow local refinement of the mesh by subdivision and would ease calculating curvature properties of the interpolated surface.

Medially-defined image cut planes give a local view of the image data. We are considering methods to produce larger slices that cut through the full object. This will require either a compromise cut plane or nonplanar image slices.

## 8 Acknowledgments

We are grateful to Sarang Joshi for his many contributions, Edward Chaney, Paul Yushkevich, and Guido Gerig for useful advice, Yonatan Fridman for programming contributions, and Gregg Tracton for testing and documenting our method. This work was done under the partial support of NIH grant P01 CA47982. A gift from Intel Corporation provided computers on which this research was done.

## References

1. Barr, A.H.: Global and Local Deformations of Solid Primitives. Computer Graphics, 18 (1984) 21-30.
2. Blum, H.: A Transformation for Extracting New Descriptors of Shape. In: WathenDunn, (ed.): Models for the Perception of Speech and Visual Form. MIT Press, Cambridge MA (1967) 363-380.
3. Blum, H., Nagel, R.N.: Shape Description Using Weighted Symmetric Axis Features. Pattern Recognition, 10 (1978) 167-180.
4. Danielsson, P.E.: Euclidean Distance Mapping. Computer Graphics and Image Processing, 14 (1980) 227-248.
5. Delingette, H.: General Object Reconstruction based on Simplex Meshes. Int. J. Computer Vision, 32 (1999) 111-146.
6. Joshi, S., Pizer, S.M., Fletcher P.T., Thall, A., Tracton, G.: Multi-scale 3-D Deformable Model Segmentation Based on Medial Description. To be submitted to: IPMI 2001
7. Nackman, L.R., Pizer, S.M.: Three-dimensional Shape Description Using the Symmetric Axis Transform, I: Theory. IEEE Trans. PAMI, 7 (1985) 187-202.
8. Pizer, S.M., Fletcher, P.T., Fridman, Y., Fritsch, D.S., Gash, A.G., Glotzer, J.M., Joshi, S., Thall, A., Tracton, G., Yushkevich, P., Chaney, E.L.: Deformable M-reps for 3D Medical Image Segmentation. Submitted to: MedIA
9. Sederberg, T.W., Parry, S.R.: Free-form Deformation of Solid Geometric Models. SIGGRAPH '86 Conference Proceedings, in Computer Graphics, 20 (1986) 151-160.
10. Styner, M., Gerig G.: Medial Models Incorporating Object Variability for 3D Shape Analysis. To be submitted to: IPMI 2001
11. Yushkevich, P., Pizer, S.M., Joshi S., Marron J.S.: Intuitive, Localized Analysis of Shape Variability. To be submitted to: IPMI 2001

[^0]:    ${ }^{1}$ We will distinguish between points and vectors, where a point $\mathbf{x}$ is a member of n-dimensional Euclidean space, $\mathbb{E}^{n}$, and a vector $\vec{v}$ is a member of the n-dimensional real vector space, $\mathbb{R}^{n}$.

