

M-Reps: A New Object Representation for Graphics

Stephen M. Pizer, Andrew L. Thall, David T. Chen

Department of Computer Science

University of North Carolina, Chapel Hill

Abstract

M-reps are a multiscale approach to the modeling and rendering of 3D solid geometry. Traditional geometric models, whether b-reps or CSG, are represented at infinitesimal spatial scale and then require simplification to meet needs requiring coarser scale or smaller data sets. We have developed a model that is designed at successively smaller scales and supports a coarse-to-fine hierarchy in design, rendering, physical deformation, and other graphics operations. We base our representation on *figural models*, defined at coarse scale by a hierarchy of figures – protrusions, indentations, corners, neighboring figures, and included figures – which simultaneously represent solid regions and their boundaries. To capture local shape at scale and thus local zoom-invariance, the figural components imply a fuzzy, i.e., probabilistically described boundary position with a width- and scale-proportional tolerance. At small scale these figures are made precise by displacement maps, geometric textures, or image textures. While these models can exist in 2D, we focus on models of 3D objects.

We note the needs of a variety of graphics operations and thus motivate the following definition of an m-rep. A model for a single figure is made from a mesh of medial atoms (hence the name *m-reps*), each atom describing not only a position and width, but also a local figural frame implying figural directions, and an object angle between opposing, corresponding positions on the implied boundary. In addition, width proportionality constants indicate mesh link length, boundary tolerance, boundary curvature limits, and, for 3D image visualization, an interrogation aperture. Thus, a figural mesh defines the figural boundary to within a width-proportional tolerance and provides a width-proportional sampling of the figure’s implied medial surface. A figural model, then, consists of a directed acyclic graph of figure meshes, with interfigural, intramesh, and interscale links capturing information about differences in position (and thus figural length and subfigural offset), figural width, narrowing rate, boundary curvature, and figural orientation. The leaves of the figural graph also contain the displacement maps, geometric textures, or image textures that are used to give infinitesimal scale to the model where fine detail is needed.

We describe methods for building m-rep models by a design tool, from b-rep or CSG representations, or from a 3D-intensity dataset. We present a method for deforming m-reps into image data, allowing model-directed visualization of objects in volume data. We describe algorithms for rendering m-reps, and show two uses in graphics for them – rendering for 3D visualization and computer-aided design – with preliminary results. Algorithms for other graphics objectives are sketched.

1. The Need for a New Model

Computer graphics modeling in 3D rests on five basic techniques and their variations:

- boundary representations (*b-reps*) by polygons, higher-order patches, or subdivision meshes;
- constructive solid geometry (CSG), which uses Boolean operations on simple solid primitives such as spheres, cylinders, blocks, and tori;
- implicit surfaces, which define object boundaries as isosurfaces of density functions;
- skeletal surfaces, which define object boundaries as offset or convolution surfaces; and
- image-based rendering (IBR), where geometry is reduced to a depth- or disparity-map over a pixel image.

We will see that most current techniques of computer graphics share the following problems: lack of multilocal¹ and multiscale² information with the consequence of inefficiency or inability in handling an object or a significant component as a unit; over-precision with the consequence of inflexibility, ill conditioned computation, and serious inefficiency; and inability to handle objects simultaneously as surfaces, which is needed for rendering, and as solids, which is needed for physically based modeling and is convenient for design. We motivate m-reps by examining these weaknesses of current graphics techniques regarding (1) object design, (2) rendering, (3) physically based deformation, and (4) fusion of real-world visualizations.

Object design. With b-reps, nonlocal properties are often hard to specify – for example, that a portion of a surface is (and perhaps should remain) a section of a cone. While subdivision surfaces allow multiscale modeling at different subdivision levels, they still lack such multilocal information. On the other hand, object design via CSG gives too little flexibility in the vocabulary of primitive shapes. Implicit models have problems of local control and unpredictable blending. The restricted, projective geometry of IBR and its discrete pixel-sampling make it ill-suited to object design. Skeletal methods provide better local and multilocal control over implicit object shape but have problems with pruning and bifurcations of their medial structures.

Rendering. The rendering of nearby objects requires detailed boundary information; while b-reps provide this directly, CSG models require considerable computation. Efficient rendering of distant objects by b-reps or by implicit representations requires a simplification that is computationally complex. The source of the complexity is the unavailability of multilocal information, e.g., as to the way in which an object is formed by a hierarchy of protrusions and indentations. While IBR techniques were created for fast, $O(1)$ rendering, they don't easily allow dynamic scenes because the multilocal properties of occlusion are not represented and their pixel-based nature creates sampling problems for reconstruction and display. Skeletally-based objects are generally rendered as implicit, blobby models or by tessellation of an implied-boundary isosurface, which allows level-of-detail (LOD) control of the blobs or tessellation; however, they typically do not carry tolerance information that would support an optimal choice of this LOD.

Physically based operations. These operations depend on having a mechanical model on a solid representation, which CSG provides directly but b-reps, implicit representations, and IBR do not. Moreover, even when a solid

¹ *Multilocal* is defined as “involving multiple separated positions.”

² *Multiscale* is defined as “at multiple levels of detail.”

representation is derived from b-reps, it is dealt with mechanically only at small scale, e.g., by finite element techniques, resulting in a large computational load. An ability to work in coarser units seems necessary, but the locality of the boundary patches does not yield such mechanically sensible units without lengthy computation. Skeletal methods have advantages in this regard and have found common usage in animation of articulated figures.

Fusion of images. The fusion of images from nearby viewpoints in image-based rendering depends on object correspondence determinations, given the typical inability to measure the viewpoint from each image with adequate precision. The information in b-reps, implicit representations, or pixel maps is too local to determine the correct way to fuse scenes where changes in the inter-object occlusion from one image to the next reveal previously hidden image regions or occlude previously visible ones.

To overcome these drawbacks, we need a representation that is simultaneously easy to view as boundaries or as solid, that is easy to view at multiple levels of scale, and that can handle the multilocal aspects of objects and the solid figures that make them up. The very precision of standard object representations is an impediment. B-rep and CSG models create a Boolean division of modeling space: for every real position, one can determine unequivocally whether the position is inside or outside the represented object. Even if the representation is an approximation and thus not necessarily accurate – as with an object boundary represented by a polygonal mesh or a spline – the representation itself is seen as precise. As a result, computing Boolean operations on objects is a difficult, often ill-conditioned problem. Also, when an object is to be viewed or manipulated in a situation requiring only large-scale information, either the calculation is quite inefficient or simplification must be carried out to reduce the object to a coarser approximation. Examples of such applications include perspective rendering of objects very distant from the viewer, occlusion culling in complex models, and cases in mechanical modeling where properties must be based on the whole object but need only be grossly correct.


An alternative approach is to associate explicit boundary tolerances with various stages and locations of the representation and to represent the object successively from large scale to small, with a larger tolerance associated with the larger scale components. But for the representation to be locally magnification invariant – and thus a valid representation for shape not only globally but also locally – the acceptable tolerance for a component of an object must be proportional to its scale. This scale may vary across a component as well.

A useful and longstanding measure of the object size at a position of a figure is the figural width, an aspect of a multilocal, solid representation. This measure is based on the understanding that the medial relationship between points on opposite sides of a figure is an important factor in the object's shape description (see Fig. 1). Biederman [1987], Marr [1978], Burbeck [1996], Leyton [1992], Lee [1995], and others have produced psychophysical and neurophysiological evidence for the importance of medial relationships (in 2D projection) in human vision. The relation has also been explored in 3D by Nackman [1985] and Vermeer [1994], and medial axis and modeling techniques have been applied by many researchers, including Bloomenthal [1991], Wyvill [1986], Singh [1998], Amenta [1998], Bittar [1995], Igarashi [1999] and Markosian [1999]. Of these: Bloomenthal and Wyvill provided skeletal-based soft-objects; Singh provided skeletally defined tubes; Amenta and Bittar worked on medially based reconstruction; Igarashi used a medial spine in 2D to generate 3D surfaces from sketched outlines; and Markosian used implicit surfaces generated by skeletal polyhedra.

Our idea is to expand the notion of medial relations not only to the notion of a medial atom implying boundaries but also to the inclusion of a natural (width-proportional) tolerance, and to use a natural (width-proportional)

sampling of the medial surface in place of a continuous medial manifold. The advantages, relative to the ideas of medial axis descended from Blum [1967], are in representational and computational efficiency and in stability with respect to boundary perturbation. Associating a tolerance with the boundary position provides advantages for (a) efficient operations in design, rendering, morphing, etc., on object representations and (b) simplified conceptualization of shape for interactive use. Stages of the representation with large tolerance can ignore detail and focus on gross shape, and in these large-tolerance stages discrete samplings 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 more local aspects. Operations can limit the scales of the representation used to only those needed and can do the grosser work via very efficient parts of the representation.

M-reps provide a means of figural object representation wherein the atoms have associated widths and many levels of scale. Along with each scale is a width-multiplying factor producing a boundary tolerance associated with the boundary position implied by the atom. The levels of scale are from largest to smallest, the whole object, the hierarchy of figures, a coarse-to-fine chains of meshes of medial atoms representing each figure, a coarse-to-fine chain of boundary displacements, either deterministic or random (textural), and finally image-based texture maps. The hierarchy of figures (see Fig. 1) consists of the main figure(s), these figures' subfigures, and those subfigures' hierarchies in turn.



2. Figural Representations by Medial Atoms

Intuitively a figure is a main component of an object: a protrusion, an indentation, a hole, or an associated nearby or internally contained object. In [Pizer 1998] we carefully define a figure, making it clear that the notion is centered on the association of opposing points on the figure called by Blum [1967] "medial involutes" (Fig. 2). Whereas Blum conceived of starting from a boundary representation and deriving the medial involutes, our idea is to start with a representation giving medial information and thus widths, and imply sections of figure

bounded by involutorial regions. As illustrated in Fig. 3, in order for a medial atom $\underline{m}$ by itself to imply two opposing sections of boundary, we define it to consist of

- 1) a position, $\underline{x}$, the skeletal position;
- 2) a width, r , the distance from the skeletal position to two or more implied boundary positions;
- 3) a frame $\underline{F}=(\vec{n}, \vec{b}, \vec{b}^\perp)$, implying the tangent plane to the skeleton (via its normal $\vec{n}$) and $\vec{b}$, the particular unit vector in that tangent plane that is along the direction of fastest narrowing/widening between the implied boundary sections;
- 4) an “object angle” θ that determines the angulation of the implied sections of boundary relative to $\vec{b}$. $\vec{b}$ is rotated by $\pm\theta$ towards $\vec{n}$ to produce normals to the implied boundary.





Fig. 3. A medial atom shown in plane of the primitive ($\vec{b}^\perp$ is perpendicular to the page). For a slab-like section of figure, the two vectors $rR(\theta)\vec{b}$ and $rR(-\theta)\vec{b}$ provide links between medial point and the implied boundary, giving approximations to both its position and its normal. The tolerance of the position of the implied boundary is indicated by the shading, intended to imply a Gaussian tolerance. The implied boundary is slab-like and centered on the head of the bold arrow, i.e., it is extended in the $\vec{b}^\perp$ direction just as it is illustrated to do in the direction perpendicular to the two links.

Also implied by the atom are

- 1) a tolerance $\tau = \alpha r$ of boundary position normal to the boundary,
- 2) the length of links to other primitives, approximately of length βr , with β a significant fraction or 1.0, and
- 3) a constraint δr on the radius of curvature of the boundary.




Fig 4. Segment types and their atoms

Upper left) slablike net; upper right) implied surface, lower left) tubelike net, lower right) implied surface

Whereas, we will ultimately be defining a hierarchy of figures, for the time being we focus on a single figure. A figure is represented by a net (or sequence) of medial atoms. Fig 4 illustrates such a net, in which there is a medial atom at each node of the net. There are three basic types of medially defined figural segments:


- 1) slab-like segments, in which the net is two-dimensional and internal nodes of the net have a pair of boundary-pointing vectors, $\mathbf{x} + R(\theta)\vec{b}$ and $\mathbf{x} + R(-\theta)\vec{b}$, where $R(\theta)\vec{b}$ denotes rotation of $\vec{b}$ by θ in the plane defined by $\vec{b}$ and the normal to the medial tangent plane.

- 2) tube-like segments, where the net is a 1-D sequence of medial atoms, and the boundary segments pointed to by each medial atom are circles whose axes are tangent to the medial track. $R(\theta)\vec{b}$, in turn, is a right circular cone of vectors pointing from x to the implied tubular boundary and normal to that boundary.
- 3) spherical segments, where the net consists of a single atom whose position is at the center of boundary primitives arranged in a sphere at distance r from x .

We see that for each type of figural segment, the set of boundary pointing vectors has a different dimensionality. Hybrids of the figural segment types are possible, such as a slab that rounds out into a tube and then perhaps flattens out back into a slab.

We call a figure represented via m-reps an *m-figure*. Within each m-figure type, the net of medial atoms contains internal nodes and most often end nodes, as well. The end nodes for a slab are linked together to form the boundary of the net. For an object made from a single figure, the end nodes need to capture how the boundary of the slab or tube is closed by a crest. For example, a pancake is closed at its sides along what is called a *crest* in differential geometry [Koenderink 1990]. Whereas the internal nodes for a slab-like segment have two boundary-pointing vectors, end nodes for slab-like segments have three boundary pointing vectors, with the additional vector pointing in the $\vec{b}$ direction to the end. Whereas internal nodes for tubes have a circle of boundary pointing vectors, end nodes for tube-like segments have a circularly symmetric end about and closing in the $\vec{b}$ direction, beginning from the circle symmetric about $x + R(\theta)\vec{b}$ where the boundary makes an angle θ with the vectors from x to that circle.

As illustrated in Fig. 5, segment closed ends may be rounded or they may be pointed, i.e., take the form of corners. Rounded atoms can be parameterized by a constant γ giving their elongation: the vertex is taken at $x + \gamma\vec{b}$. Corner atoms are needed to avoid the inefficiencies of very fine medial sampling where the width becomes very small as one approaches the corner. They are orthogonal to the boundary-pointing vectors that do not point to the vertex and have a corner at $x + (1/\cos(\theta))\vec{b}$. For slabs a sequence of corner atoms forms a curve of a crest or a curve of a corner. For tubes a corner end node implies a conical boundary segment.



As illustrated in Fig. 1, figures combine into objects in a hierarchical fashion, with the same Boolean operators as with CSG models, but here recognizing the tolerance of the figures. Fig. 6 shows how the corresponding




Fig. 6. Medial nets for sailboat. The main boat, the mast, and the rudder, are each protrusion figures. The cockpit (green mesh) is an indentation figure. Top left: Meshes of medial atom positions (at the vertices of the mesh). Top right: Shown with medial atoms and boundary tolerances. Bottom: Rendering of the full multifigure boat.

medial nets are related. A figure may be separated from all other figures, or it may be the parent of protrusion and indentation figures. A protrusion on a slab has a segment of its medial net's end nodes that are at the open end of the figure, where it attaches to its parent. The remaining end nodes form the closure of that figure. A protrusion on a slab or tube may be tube or a slab. If the protrusion is a tube, it has a single open-end atom, where the tube is attached to its parent, and a closed end atom at the other end.

Indentation figures are similar to protrusion figures, except that their axes are inside the indentation rather than within the mass of the object and their closed ends are inside the parent figure.

The successive refinement, coarse-to-fine, of a medial mesh is seen as a correction from medial atoms at the finer spacing predicted from those at coarser spacing. This refinement brings with it a decrease in tolerance of the implied boundary and radius of curvature constraint. From the mesh of medial atoms defining a particular m-figure one can interpolate the medial atoms at any finer mesh [Yushkevich 1999]. A medial atom in this finer scale mesh (for which the constants of proportionality, α , β , and δ , must also be proportionately smaller) can be considered as a correction of the interpolated medial atom. That is, if $\underline{m} = (\underline{x}, r, \underline{E}, \theta)$ and $\underline{m}_{\text{interp}} = (\underline{x}_{\text{interp}}, r_{\text{interp}}, \underline{E}_{\text{interp}}, \theta_{\text{interp}})$, the correction is $\underline{x} - \underline{x}_{\text{interp}}, r - r_{\text{interp}}, \underline{R}, \theta - \theta_{\text{interp}}$, where $\underline{R}$ is the rotation matrix rotating $\underline{E}_{\text{interp}}$ into $\underline{E}$.

3. Geometric Implications of M-rep Models

3.1 Implied boundaries

While a net of medial atoms imply a boundary with tolerance, it is useful for the mean boundary to be efficiently computable from adjacent discrete atoms and for the tolerance to be taken in the direction normal to that mean boundary. To satisfy requirements of local zoom invariance, the implied mean boundary at a point must have a local radius of curvature upper bound proportional to the r -value of the corresponding medial atom. Assuming that the boundary is defined in a piecewise fashion, i.e., as a spline, the curvature limitations must apply at the knots as well.

Because the implied boundary will have tolerance, the precise location of the boundary mean is not critical. Therefore, a variety of schemes for determining the boundary are possible. In the early stages of our work, surfaces have been defined and displayed by coarse polygonal tessellation, which doesn't preserve boundary tangency, or by Hermite spline interpolation, which doesn't preserve boundary curvature from the medial representation. To improve over these, we have developed a subdivision scheme that constrains surface-mesh subdivision based on the medially-implied tangency and curvature conditions. The method is a generalization of the Loop [1987] subdivision method. In the subdivision stage of this generalized method it is not the boundary position and surface normal that is interpolated from adjacent node. Rather the full medial atom is interpolated. That is, the atom's position, the full local frame $\underline{E}$, the radius and thus the implied boundaries' local tolerance, and the object angle are all interpolated. Moreover, the interpolation's weighting of each of the adjacent nodes is related to its associated tolerance.

We are also exploring implicit-surface methods, whereby the implied mean boundary will be an isosurface of a boundary function.

How is the boundary formed at the seam between a figure and its subfigure? Each figure has its own mean boundary locus, its boundary normal-field, and its boundary tolerance-field. The normals and the tolerances can guide the blending of surface 1 with normals 1 at scale 1 with surface 2 with normals 2 at scale 2 to produce a blended surface with its own normal vector- and tolerance-field. Again the tolerance limits the curvature of the blend, and the precise formula is not critical as a result of the existence of tolerance. Implicit-surface techniques offer the most straightforward approaches to blending of figure-subfigure joints, as well as for modeling indentation figures.

3.2 Design and deformation of m-rep models

We anticipate that models via m-reps will come from two sources: design by a human and deformation of such a design into 3D image data. In each case we see the model being instantiated at one level of coarseness and refined at successively finer levels. Thus a finer level might a) add one or more subfigures at their primary level of coarseness, or b) refine a particular net into a finer net. If the net links were fraction f of their length in the coarser net, its boundary tolerances would be f of their value in the coarser representation. In its original form, the new net would have some of its nodes identical to the coarser one, and others would be interpolated between the nodes at the coarser scale [Culver 1999, Yushkevich 1999].

Deforming a figure is naturally done at a scale proportional to its radius, since the links are βr apart. That is, a narrower object will naturally be easier to put small wrinkles in than a wider object. For the wider object, we can use the hierarchical properties of the representation. That is, the net with radius-proportional links is taken to imply interpolated medial atoms at finer spacing, and a smaller scale component of the representation with these shorter links gives deviations of the medial atoms from those interpolated from the larger scale. Alternatively, these small-scale wrinkles can be represented by boundary displacements.

The dynamic nature of the medial-mesh is a critical factor in object-deformations. An m-rep may be stretched in a direction parallel to the medial surface by simply moving the medial atoms apart while keeping the radii fixed; when the inter-atomic links exceed the tolerance of the width-proportional sampling, a new medial atom will be interpolated between the separated ones. Similarly, compressions will necessitate the removal of intermediate atoms and the “absorption” of their medial properties by their neighbors. Because of the tolerance inherent in the width-based inter-atomic link distances, there will be no degenerate regions of the medial surface, and because the implied medial surface does not bifurcate (this in contrast to the Blum medial axis), the topology of the implicit medial surface remains planar and simple.

We have created an object sculpting and deformation tool, in which stock medial nets are molded into the desired m-figure, and these figures are composited as protrusions and indentations into a total object. The medial nets are tiled into a net of quadrilaterals or into triangles. During modification, selected subsets of medial atoms in an m-figure – typically, chains of linked medial atoms – are simultaneously geometrically transformed, sometimes in a common world frame, but frequently each in its own frame. For example, in Fig. 7, the medial atoms in each row across the centreboard formed a hinge about which the medial atoms below were rotated. Then each row was rotated in its own central atom’s coordinate system to produce the twist.

In contrast to the situation with the Blum medial axis, with m-reps the representation of an object in terms of figures is not simply a property of the object. Instead, many objects yield a variety of possibilities, so in design, one has some choice of the m-rep used to represent an object. For example, as illustrated in Fig. 8, a doughnut

can be modeled as a closed tube of dough (a single figure) or as a bun with a hole in it (two figures). One would choose the representation so that desired deformations are easy to specify, avoiding a weakness of the Blum medial axis that deformation leads to discontinuous, qualitative changes in the representation. In the example, one would use the tube representation if one wished principally to vary the width or center track of the tube locally, whereas one would use the the bun-with-a-hole representation if one wished principally to vary the shape of the hole or the local height of the bun.








Fig. 7. The centreboard on the sailboat after it ran into some rocks. Left: The boat before the accident. Middle: The boat after the accident. Right: The medial atoms of the centreboard after the accident.




Central cross section




Protrusion locus



Central cross section



Indentation locus



Central cross section




Fig. 8. Two figural decompositions of a doughnut:

Top) as cyclic bent cylinder,

Below) as ovoid with hole

Figs. 6 and 9 illustrate complex objects that can straightforwardly be defined with m-reps. The renderings there are the simplest, using quadrilateral boundary patches and Gouraud shading.




Fig. 9. Electric guitar modeled via m-rep

3.3 Display-resolution sensitive rendering from m-reps

While it is possible to construct a traditional surface representation for the medially implied boundary, it appears preferable to define the rendered surface only within the level-of-detail required by screen-resolution or desired precision. Therefore, the scale level of the representation can be determined by pixel projection – that is, if the pixel is projected to the medially implied boundary along a visual ray, the size of the pixel relative to the medial sampling can be determined. This idea could be implemented in either an adaptive subdivision rendering scheme or in an implicit function rendering scheme.

3.3.1 Subdivision schemes

There are two main approaches. (1) Using a subdivision-surface algorithm, such as Loop subdivision or the modified Butterfly algorithm [Dyn 1990] [Zorin 1996] for tri-meshes, we generate higher density meshes of medial atoms, interpolating orientation, radius, and object-angle information along with atom positions. These new meshes are used to generate finer surface tessellations, according to the detail necessary. (2) The coarse

medial mesh is used to generate a coarse boundary tessellation, which is then refined using a subdivision-surface generating algorithm. These two approaches can be combined, using subdivision both for medial mesh and boundary definition. Because of the tolerances in mesh and boundary location, interpolating subdivision schemes might be eschewed for approximating ones. Boundary subdivision may be performed adaptively, sending vertices to their limit points when a sufficient tessellation has been achieved, and subdividing to differing levels-of-detail at different points on the boundary.

3.3.2 Implicit function schemes

We fit implicit patches to a coarsely tessellated boundary based on the boundary locations and normals specified by the medial primitives. These implicit patches can be similar to those proposed in [Bajaj & Ihm, 1992], based on a spacecurve mesh with defined normals. During rendering, relative pixel size at the boundary can be used, as in beam- or cone-tracing, as an interrogation aperture to evaluate the implicit surface functions; thus, the surface may be determined only when needed and only to the LOD desired. If a more efficient rendering is desired, a wider interrogation aperture can be used and its result blurred over multiple pixels; alternatively, an appropriately coarser version from the coarse-to-fine m-figure hierarchy could be selected and rendered.

Coarse-to-fine rendering for speeding scan conversion might use a coarse scale version of the object and then refine (correct) the surface position and normal, with the correction generated from the finer scale representation of the object. Shading for a scan converted pixel can then be computed from the gradient of the implicit function representing the medially implied boundary (possibly with a displacement) with tolerance or, alternately, by interpolation, with the position and normal of the implied boundary interpolated from the net of medial atoms. Work is in progress both on a ray-traced solution to rendering the m-rep boundary as an implicit surface and on subdivision methods both for boundary determination and for medial surface refinement.

3.4 Building a model from a b-rep or CSG representation

It is possible to mimic CSG primitives with m-reps, allowing a straightforward conversion from a CSG model to an m-rep. This would not take advantage of deformability of m-figures, which in some cases might allow a complex solid to be modeled with a sparser graph than an equivalent CSG-tree.

Conversion from b-reps to m-reps can be done by a manual ‘sculpting’ approach; this is simply an application of our object design tool, described in section 3.2, except that the tool is augmented by the display of the b-reps represented object.

Alternatively, we can fill the solids implied by the b-rep and then treat the resulting characteristic function³ as an image. A semiautomatic skeletonization technique can then be used to get an approximate medial surface, which, with manual pruning, could serve as a basis for creating the m-figure mesh. The boundary would then be specified as a displacement surface on the medially implied boundary. Semiautomatic techniques for fitting continuous loci of medial atoms or other skeletons to volume-data objects already exist [Furst 1999, Bloomenthal 1998]. These could be sampled to provide the m-rep. We have also built a tool that allows the fitting of m-reps organized into planar slices to greyscale images [Fletcher 1999].

³ A characteristic function corresponding to an object is an image that has the value 1 inside the object and 0 outside.

4. Graphics Operations via M-reps

4.1 CAD via m-reps

Modeling with m-reps will share many similarities with CSG approaches to CAD, but it has the advantage that the basic figures are more flexible, bringing with it a need to specify them. At the first stage, objects are constructed of directed acyclic graphs of figures and subfigures in a manner analogous to CSG-trees. The deformability of the individual figures adds considerably to the pallet of stock shapes representable, and an m-rep DAG will typically require many fewer nodes than an equivalent CSG-tree. Compositing operations analogous to the Boolean operations on binary CSG primitives are being developed, using the tolerances on the implied boundaries and their displacements to limit and determine intersections between primitives. Fig. 9 shows a result.

M-rep modeling also has similarities with blobby modeling techniques, as first developed by Blinn [1982] – each medial atom is a volume primitive, and an m-rep mesh is the union of the blobby medial atoms that comprise it, along with the boundary curvature information inherent in the atoms' object angles and in the links between the atoms. Specific knowledge that the curvature of a section is zero, i.e., the implied boundary must be flat can be used.

Three-dimensional sculpting techniques provide an ideal paradigm for creating m-figures. The easy deformability of an m-rep lends itself to interactive control of the medial mesh and the boundary, creating protrusions and indentations according to spatial intuitions about solid objects. Present-day CG-design interfaces are readily applicable, though not necessarily ideal. A better approach is to allow object-level shaping deformations, as done for surfaces by Rhoades [1992]. Virtual interfaces, especially with tactile force-feedback, would find excellent application in medially based morphological structure of the object, giving a distinct advantage over the space-warping deformation techniques often used in solid modeling. M-reps provide object-level deformation, allowing both global and local control. Moreover, because they are solid, deformation could be made to follow physical models, themselves implemented in a coarse-to-fine fashion derived from the underlying m-reps.

4.2 Visualization via m-reps

To compute a representation of an object in a volume image, an m-rep modeling the expected object can be deformed into the volume image data, via model-based segmentation techniques described in [Pizer 1999]. Because the model can be very sparse and can proceed coarse to fine, this step could be accomplished in some seconds. As described in [Chen 1998], the resulting m-rep implied object boundary with its tolerance can be used to scan convert the display pixels to boundary positions in 3D and then, for each such boundary position, to restrict the search for the boundary needed to render the object. This restriction

- a) avoids occlusion by other objects,
- b) provides efficiency by restricting the search area and allowing a 1D search region for each boundary position,
- c) provides directionality for the intensity derivatives needed to locate the boundary,
- d) provides the radius-proportional scale of boundary sampling needed, the scale of boundary search step needed, and the scale of the image-interrogating aperture.

Figures 10 and 12 show examples of this approach. In the kidney example, the m-reps involved ~ 50 medial atoms for each kidney. Two figures, with the medially implied boundaries seen in Fig. 11 – one main figure and one indentation (in red) – represent each of the two kidneys. Used in this way, M-reps provide a volume-primitive for rendering purposes which (unlike simple polyhedral volume-primitives such as those of Grzeszczuk [1998]) utilizes object tolerance and is not limited to linear space-warp distortions.




Fig. 10. M-rep based rendering of starfruit, avoiding occlusion from other fruits.
a) A ray traced image from a 3D CT image of a starfruit surrounded by 4 gourds.
b) Splatted rendering of the starfruit selected by the m-rep model-guided renderer.




Fig. 11. Medially implied boundaries from four kidney figures. The figures in red are indentations.




Fig. 12. M-rep based rendering of kidneys from CT volume image, avoiding occlusion.
a) A slice from the CT scan of an abdomen, with the kidneys marked with x's.
b) The kidneys rendered by our system.

4.3 Other graphics operations via m-reps

There are three topics worth special consideration: simplification, morphing, and physically based dynamics.

4.3.1 Simplification

M-reps automatically provide a method for simplification since the representation is built coarse to fine. There are two aspects to simplification:

- 1) cutting off the hierarchy of coarser/finer meshes with associated subfigures dependent on image scale – this involves the culling of figures/subfigures based on image-scale, in a way similar to techniques created for polygonalized models;
- 2) coarseness level culling – stopping the refinement of an m-figure at a given level of coarseness. Notice that this allows a particular figure still to narrow arbitrarily, e.g., to a corner, without rounding. Alternatively, the radius (and thus tolerance) information available in each medial atom also allows a culling on the basis of spatial scale of the feature. By this mean the evaluation of the implicit function of the medial figure can be done at the correct image scale, using a ray-casting, splatting, or displacement-texture approach to rendering.

The second method is the more exciting and innovative. M-figures are inherently minimal – the width-based sampling is an attempt use only the atoms necessary and sufficient to represent a given solid. Because of the polymorphism of an m-figure representation, there is no guarantee of minimalism, but the representation is nonetheless a sparse one. This makes m-figures apt for image-space rendering techniques similar to those devised for subdivision surfaces and spline-patches, while allowing superior control of object morphology. Volume-graphics techniques such as polygon-assisted ray-casting (PARC), as developed by Avila et al [1992] for the *Volvis* system, would find easy application as well, given the explicit boundary tolerances on the medially based figures.

4.3.2 Deformation / Morphing

Morphing of one m-rep-based object to another will take place on a figure-by-figure basis. Axis positions and distances (in radius units) will be mapped from starting to final representation, and then the m-rep parameters will be interpolated, with new mesh-atoms being inserted or deleted automatically to preserve width-proportionality. Eberly [1994] has shown this idea to be effective in 2D, using a Blum medial axis, and Gagvani [1998] has demonstrated the effectiveness of 3D volume-animation based on the medial skeleton. Morphing based on m-rep meshes will be equally straightforward, with the width-proportional sampling allowing the object to be deformed simply by modifying the existing mesh diatoms, letting the object maintain its morphological ‘identity’ with no additional work.

4.3.3 Physically-based object dynamics

M-figures offer a volume-primitive for creation of dynamic models. Based on work by Stetten [1998] using Blum-defined medial axes, we believe it will be possible to develop volume-preserving deformations for m-rep

models. Stetten also showed that flow along tubes could be simulated by particle-flow parallel to the medial axis, making m-reps idea for such work. M-reps also provide a width-sampled skeleton for mechanical deformation and articulation, and for computation of center-of-gravity and higher moments of inertia. Single-figure models should work well, at the least, though multifigure models are more problematic, especially with indentation figures. There are similar problems with CSG; however, m-rep hierarchies may be much sparser than CSG trees, due to information density in each medial-atom, so computation of moments for the created solids may be simpler.

5. Discussion and Conclusions

We have argued that m-reps have numerous advantages over both b-reps and other solid-modeling representations. We have supported these arguments by some prototype examples, showing that design and visualization of 3D images do indeed benefit from the efficiency and width-proportional tolerance properties of m-reps. We have demonstrated a basic design and rendering tool for m-rep models. We have shown a tested method for visualization from 3D images using m-reps.

However, much needs to be done in the way of implementations and comparisons with traditional methods before the full potential of m-reps can be realized. We require further development of the modified subdivision schemes – both approximating (Loop) and interpolating (modified Butterfly) – used in rendering and boundary computation, to better reflect the rich geometric structure of medial loci. Improvements in the design tool to provide sculpting, bending, and other natural operations at scale are needed. Applications in CAD, morphing for animation, physically based modeling, image based rendering, 3D-image visualization and a variety of other areas of graphics are anticipated, but the actual usefulness of m-reps in these areas awaits trial.

Acknowledgments

We appreciate advice and/or programming help from Gary Bishop, Edward Chaney, Dinesh Manocha, Daniel Fritsch, KC Low, Valen Johnson, J. Turner Whitted, and Alyson Wilson. This work was done under the partial support of NIH grant P01 CA47982, NSF SGER Grant CCR-9910419, and a fellowship from the Link Foundation. A gift from Intel Corp. provided computers on which some of this research was carried out.

References

- Amenta, N., Bern, M., Kamvysselis, M., (1998). A New Voronoi-Based Surface Reconstruction Algorithm. *Computer Graphics Proceedings, Annual Conference Series*, 1998, ACM SIGGRAPH, 415-422.
- Avila, RS, LM Sobierajski, AE Kaufman (1992). Towards a Comprehensive Volume Visualization System. *Proceedings of IEEE Visualization '92*, 13-20.
- Bajaj, CL, I Ihm (1992). Algebraic surface design with Hermite interpolation. *ACM Transactions on Graphics* **11**(1), Jan 1992, 61-91.
- Biederman, I (1987). Recognition by components: a theory of human image understanding. *Psychol Rev* **94**(2): 115-147.
- Bittar, E, N Tsingos, M Gascuel (1995). Automatic Reconstruction of Unstructured 3D Data: Combining a Medial Axis and Implicit Surfaces. *Eurographics 1995*.
- Blinn, J. F. (1982). A Generalization of Algebraic Surface Drawing. *ACM Transactions on Graphics* **1**(3), July 1982, 235-256.

- Bloomenthal, J, K Shoemake (1991). Convolution Surfaces. *SIGGRAPH '91 Conference Proceedings*, in *Computer Graphics*, **25**(4), 251-256.
- Bloomenthal, J (1998). Presentation at Tutorial at SIGGRAPH '98.
- Blum, H (1967). A transformation for extracting new descriptors of shape. In W. Wathen-Dunn, ed., *Models for the Perception of Speech and Visual Form*. MIT Press, Cambridge MA: 363-380.
- Chen, DT, SM Pizer, JT Whitted (1999). Using Multiscale Medial Models to Guide Volume Visualization. Tech report TR99-014, Dept. of Comp. Sci., Univ. of NC at Chapel Hill.
- Culver, T (1999). Intrinsic scale for boundary representation of two-dimensional objects, Tech Report TR99-013, Dept of Comp. Sci., Univ. of NC at Chapel Hill.
- Cootes, TF., A Hill, CJ Taylor, J Haslam (1993). The use of active shape models for locating structures in medical images. *Information Processing in Medical Imaging*, HH Barrett & AF Gmitro, eds., Lecture Notes in Computer Science **687**: 33-47, Springer-Verlag, Berlin.
- Dyn, N, D Levin, J Gregory (1990). A butterfly subdivision scheme for surface interpolation with tension control. *ACM Transactions on Graphics*, **9**(2):160-169.
- Eberly, D, (1994). Techniques for morphing two and three-dimensional objects. Internal report, Dept. of Comp. Sci., Univ. of NC at Chapel Hill.
- Fletcher P, A Thall, DS Fritsch, Y Fridman (1999). A solid modeling program using slice-constrained medial primitives for modeling anatomical objects. Internal report, Dept of Comp. Sci., Univ. of NC at Chapel Hill.
- Fritsch, DS, SMPizer, L Yu, V Johnson, EL Chaney (1997) . Localization and Segmentation of Medical Image Objects using Deformable Shape Loci. *Information Processing in Medical Imaging (IPMI)* , J Duncan & G Gindi, eds., Lecture Notes in Computer Science **1230**: 127-140, Springer-Verlag, Berlin.
- Furst, JD, RS Keller, JE Miller, SM Pizer (1997). Image loci are ridges in geometric Spaces. *Scale-Space Theory in Computer Vision: Proceedings of First International Conference, Scale-Space '97*, BM ter Haar Romeny, ed. Lecture Notes in Computer Science **1252**:176-187, Springer-Verlag, Berlin.
- Furst, JD (1999). Height Ridges of Oriented Medialness. Ph.D. Thesis, Univ. of NC at Chapel Hill, 1999.
- Gagvani, N, D Kenchammana-Hosekote, D Silver (1998). Volume animation using the skeletal tree. *Proceedings ACM/IEEE Symposium on Volume Visualization '98*, RTP, NC: 47-53.
- Grzeszczuk, R, C Henn, R Yagel (1998). Advanced Geometric Techniques for Raycasting Volumes, Tutorial Notes, *SIGGRAPH '98*, July 1998.
- Igarashi, T, S Matsuoka, H Tanaka (1999). Teddy: a sketching interface for 3D freeform design. *Computer Graphics Proceedings, Annual Conference Series*, 1999, ACM SIGGRAPH: 409-416.
- Koenderink, JJ (1990). *Solid Shape*. MIT Press, Cambridge, MA.
- Lee, TS, D Mumford, PH Schiller (1995). Neuronal correlates of boundary and medial axis representations in primate striate cortex. *Investigative Ophthalmology and Visual Science Annual Meeting*, abstract #2205.
- Leyton, M (1992). *Symmetry, Causality, Mind*. MIT Press, Cambridge, MA.
- Loop, CT (1987). Smooth subdivision surfaces based on triangles. MS Thesis, Dept. of Math., Univ of Utah.
- Marr, D, HK Nishihara (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proc Royal Soc, Series B*, **200**: 269-294.
- Markosian, L, JM Cohen, T Crulli, J Hughes (1999). Skin: a constructive approach to modeling free-form shapes. *Computer Graphics Proceedings, Annual Conference Series*, 1999, ACM SIGGRAPH: 393-400.
- Nackman, LR, SM Pizer (1985). Three-dimensional shape description using the symmetric axis transform, I: Theory. *IEEE Trans. PAMI*, **7**(2): 187-202.
- Pizer, S.M, DS Fritsch, P Yushkevich, V Johnson, EL Chaney (1996). Segmentation, registration and measurement of shape variation via image object shape. University of North Carolina Computer Science

Department technical report TR96-020. Tutorial notes, *Visualization in Biomedical Computing '96*.
Submitted for journal publication.

- Pizer, SM, D Eberly, BS Morse, DS Fritsch (1996). Zoom-Invariant Vision of Figural Shape: The Mathematics of Cores. *Computer Vision and Image Understanding*, **69**:55-71, 1998.
- Pizer, SM, Fritsch, DS, Low, KC, Furst, JD (1998). 2D & 3D figural models of anatomic objects from medical images. *Mathematical Morphology and Its Applications to Image Processing*, HJAM Heijmans, JBTM Roerdink, eds. (invited paper, Proc. ISMM '98), Kluwer Computational Imaging and Vision Series: 139-150.
- Rhoades, JS, (1993). *Shaping Curved Surfaces*. Ph.D. Thesis, Univ. of NC at Chapel Hill.
- Singh, K., Fiume, E., (1998). Wires: a Geometric Deformation Technique. *Computer Graphics Proceedings*, Annual Conference Series, 1998, ACM SIGGRAPH, 405-414.
- Stetten, G. (1998) Volume of arbitrary shapes from boundary curvature and medial scale. Tech report TR98-037, Dept. of Comp. Sci., Univ. of NC at Chapel Hill.
- Vermeer, P.J., (1994). Medial Axis Transform to Boundary Representation Conversion, Ph.D. Thesis, Purdue University, May 1994.
- Yushkevich, P, SM Pizer, T Culver (1999). Statistical object shape via a medial representation. Internal report, , Dept. of Comp. Sci., Univ. of NC at Chapel Hill.
- Wyvill, G., McPheeters, C., and Wyvill, B. (1986). Data Structures for Soft Objects, *Visual Computer* **2**(4): 227-234.
- Zorin, D, P Schroeder, W Sweldens (1996). Interpolating subdivision for meshes with arbitrary topology. *Computer Graphics Proceedings*, Annual Conference Series, 1996, ACM SIGGRAPH, 189-192.