

Deformable Solid Modeling using Sampled Medial Surfaces: A Multiscale Approach

Technical Report TR00-005

Andrew Thall
thall@cs.unc.edu

Stephen M. Pizer
smp@cs.unc.edu

P. Thomas Fletcher
fletcher@cs.unc.edu

Department of Computer Science
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-3175

Abstract

We introduce a multiscale approach to modeling and rendering deformable solid objects, replacing polygonal or spline-based B-reps with a medial representation (*M-rep*) based on hierarchically linked, discretely sampled medial figures. Each medial figure is a width-proportionally sampled medial surface which implies object boundaries within a width-proportional tolerance; boundary variations within this tolerance are specified separately by surface displacements or texture-maps. *M-reps* share advantages of other skeletal techniques for modeling deformable and articulated figures, but by decoupling object-shape from fine surface detail, they (a) eliminate skeletal instabilities due to small-scale boundary-perturbation, and (b) induce a natural hierarchy and level-of-detail for object-based deformations. This work discusses the theoretical foundation of *M-reps*, their data structures and implementation, their current use in medical-image analysis and volume rendering, and how model building and rendering are done via *M-reps*, using texture maps and subdivision surfaces for fine detail.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representation; Object hierarchies; Boundary representation; I.3.6 [Computer Graphics]: Methodology and Techniques—Graphics data structures and data types; I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Shape; Surface fitting

Keywords: CSG, Curves & Surfaces, Deformations, Geometric Modeling, Level of Detail Algorithms, Shape Blending, Solid Modeling, Spatial Data Structures, Volume Rendering

1 Introduction

In his SIGGRAPH '86 paper on raytracing deformed surfaces, Al Barr made the following characterization of then state-of-the-art geometric modeling:

The traditional geometric modeling primitives are well suited to modeling static shapes, such as rigid ma-




Figure 1: Guitar modeled by 3-figure *M-rep*

chine parts. They are not so ideal, however, for compactly representing objects whose shapes change as a function of time, such as the shapes of flexible materials, biological forms, and the blending of shapes. The simple primitives are either too limited (as are spheres) in that they cannot model many of the desired shapes, or else are too fine-grained (as are polygons and bicubic patches) in that they require large quantities of data that human beings find difficult to manipulate.[2]

Subsequent techniques have addressed this problem with some success. Free-form deformation has become a common method for object deformation, requiring no change in the underlying geometric primitives. Subdivision surfaces, building from early Catmull-Clark and Doo-Sabin surfaces, have been researched more recently by Loop[16], Hoppe[15], Zorin[27], and others, developing hierarchical mesh structures allowing multiscale editing of hierarchical subdivision meshes to control the geometry of their limit surfaces. There has also been an increasing use of medial and skeletal techniques in the creation of deformable models and articulated figures, beginning with 2D work by Burtnyk and Wein[8], and including Blinn's work on blobby models[4] and Wyvills' and McPheeter's soft objects[25], Bloomenthal and Shoemake's convolution surfaces[5], Markosian's [17] on implicit skin-based construction, and Gagvani's[14] on volume animation using skeletal trees.

Despite the quality of this work, modern techniques still suffer from major drawbacks. Boundary representations, while allowing excellent local surface control, fail to provide a global description

of an object to be preserved under deformation. They are also *over-precise*, in that there is no clear indication of how much precision is necessary to create a shape with a given morphology. It is this over-precision and lack of global morphological description that makes model simplification such a complex task. Medial and skeletal representations handle global properties well, but while blobby models tend to be good for modeling blobs, fine surface-detail presents problems: medial representations are unstable under surface perturbation, with noise or small changes to the boundary producing branching and bifurcation of associated medial regions, and require pruning of the axis to isolate relevant detail. This, too, is a result of overprecision and lack of scale-based information in the representation.

The M-reps we have developed combine strengths of both B-reps and medial methods, while overcoming the above drawbacks. They combine local and object-level deformability, and explicitly include object-scale, tolerance, and hierarchical level-of-detail in both modeling and rendering. They use a sampled medial axis with width-proportional boundary tolerances. This choice allows a medial axis that is *not* isomorphic to the boundary representation and decouples surface detail from medial structure. Thus, M-reps eliminate the brittleness/instability with regard to surface perturbation suffered by Voronoi-region-based medial models. Even for static models, tolerance-based M-reps eliminate the over-precision from which B-rep and CSG models suffer. Fig. 1 shows an example object created with M-reps.

Of existing methods, that of Markosian et al comes nearest our own, with surface detail modeled by subdivision meshes and not affecting the medial structure. Their use of polyhedral skeleta, however, is as *ad hoc* a choice of medial primitive as others' use of Gaussian blobs, piecewise-polygonal skeleta, or space-curves. A polyhedral skeleton has the same deficiencies for deformable modeling as polygon-based surface primitives. By drawing on scale-based 3D shape analysis techniques, we have developed a more general medial primitive which provides a necessary and sufficient basis for deformable 3D object shape.

M-reps have already found application in medical-image analysis[19]. This paper explains their use in computer graphics. In Section 2, we discuss theoretical aspects: the basics of medial axis theory, as we have applied them to 3D shape description, the needs for tolerance-based medial atoms and sampled meshes, and the aspects of local invariance which follow from our need for robustness and lead to width-proportionality of both tolerance and mesh sampling. In Section 3, we discuss data structures for sampled medial figures of various topologies, including the modeling/rendering hierarchy, large-scale corner/edge primitives, and medial interpolation problems. We give examples of models constructed from 3D medical data and discuss the use of M-reps to guide segmentation and volume rendering. In Section 4, we discuss the use of M-reps in modeling and rendering, including boundary definition by polygons, splines, and subdivision surfaces, model design, figure-subfigure joining, and surface texturing. In Section 5, we conclude with some longer term goals for and benefits offered by use of M-reps, including multiscale deformation, model simplification at rendering time, physically based dynamics, shape-based morphing, CSG construction, and image-based rendering.

2 Theoretical Background

2.1 Medial Axis Theory

The history of skeletal representation in 2D and 3D image analysis goes back to Blum[6] [7], followed by careful mathematical development in 3D by Nackman[18]. In this early work, as in much development in computer graphics and computational geometry, the skeleton is derived from a continuous, precise boundary, and it im-

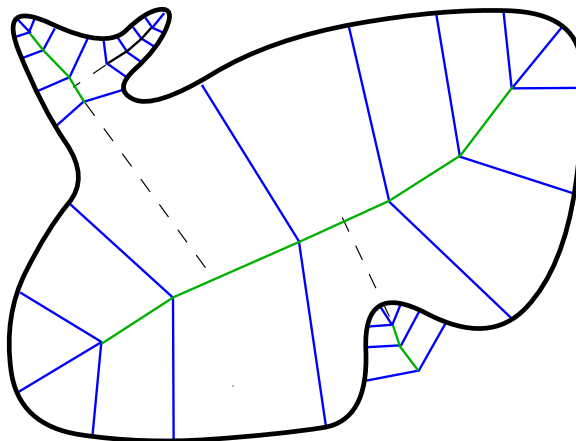


Figure 2: Simple DAG representation of a 2D M-rep

plies this boundary precisely. This property causes the medial axis to be brittle under boundary deformation, with small scale alterations in surface geometry causing large scale branchings and bifurcations of the medial structure, in accordance with the changes in the internal Voronoi regions. This also produces a nonintuitive division into *figures*, sections of the object with a nonbranching medial axis. Such skeletons are poorly suited for use as a basis for object synthesis, as in our work, by deformation of well-defined morphological figures.

In addition, the skeleton is seen in prior work as a continuous locus of atoms $[\underline{x}, r]$, where \underline{x} gives a skeletal position in 3D and r gives the radial width of a sphere bitangent to the boundary. This has led to the frequent use of piecewise linear or polygonal skeletal elements in graphics applications. In the interest of efficient representation, we want the medial locus to be sampled only as is necessary and sufficient to capture the desired 3D shape information.

These arguments have led us to a figural representation which consists of a mesh or chain of medial atoms and an object representation which is a tree or directed acyclic graph (DAG) of figures (see Fig. 2). They have also led to a medial atom that locally captures the geometry of a solid section of a figure and which implies (typically) two sections of boundary with tolerance, so that the precise variations, within this tolerance, can be captured by information at a finer level of detail.

2.2 Medial Atoms: Width-Proportional Boundary Tolerance

Our atom must capture the local relation between the medial position and section of boundary positions it implies. As illustrated in Fig 3, this atom can be visualized as a medial position and typically two arrows of equal length, at whose ends the implied boundary is to be orthogonal.¹ For reasons of geometric expressiveness and the resulting computational advantage, we choose to use an equivalent representation: $[\underline{x}, \mathbf{F}, r, \theta]$, where \mathbf{F} is a frame fitted to the local medial geometry, and thus carrying important local orientation information, and θ , the *object angle*, implies another intuitive quantity, the angle between the two implied boundaries at the end of the arrows.

\mathbf{F} is chosen with two vectors spanning the tangent plane to the implied continuous medial locus and one vector \vec{n} normal to that locus. One of the medial locus' tangent vectors, \vec{b} , is taken to be in the

¹These associated boundary points are each other's *medial involutes* associated with the medial location.




Figure 3: 2D Medial atom with boundary tolerance

direction of greatest local narrowing of the implied solid object, i.e., the direction bisecting the two implied boundaries’ tangent planes. It can be shown that if the two implied boundaries are not parallel, the two boundary pointing arrows can be determined by rotating \vec{b} by $\pm\theta$ in the $\{\vec{b}, \vec{n}\}$ plane. Moreover, \mathbf{F} and θ , as defined, give first-order derivative information for \underline{x} and r : $\nabla r = -\cos\theta \cdot \vec{b}$, i.e., the maximum change in the width r at \underline{x} has magnitude $\cos\theta$ in direction $-\vec{b}$; and $\{\vec{b}, \vec{b}^\perp\}$ gives the tangent plane (first-order derivatives) of \underline{x} , where \vec{b}^\perp is normal to the $\{\vec{b}, \vec{n}\}$ plane.

A medial figure, then, is a chain or mesh of such medial atoms. As well, we introduce a boundary tolerance τ for each atom and require the implied boundary to lie within this tolerance. Finer specification of the boundary position is determined by surface displacements. Fig. 4 illustrates these relationships.

2.3 Medial Meshes: Width-Proportional Sampling

Natural deformability requires the local geometry of the medial locus to be magnification invariant.² Local magnification invariance implies that the medial sampling, as well as the tolerance of the implied boundary, be r - (i.e. width)-proportional[19]. We apply this condition in our design tool, but somewhat loosely so that the representation need not change continuously as the object is deformed. Thus, the sampling of the mesh is roughly r -proportional, and there is an r -proportional tolerance on the distances between mesh-atoms.

In contrast to our r -proportional sampling, Amenta et al[1] have shown that *curvature*-proportional boundary sampling gives the necessary sampling frequency for reconstruction work. In our context this would imply that a broad, flat surface should require only sparse sampling, regardless of object thinness; our view, on the other hand, would require dense sampling for a thin, flat plate.

²A property is *magnification invariant* if the measurement of the property commutes with a scaling factor, i.e., $\mathbf{P}(sX) \equiv s\mathbf{P}(X)$, meaning that the measurement after object scaling is the same as the measurement before times the magnification.




Figure 4: Medial atom in 3D with boundary tolerance (cyan regions)

The reason for our choice is the need for an adequate sampling for width-proportional deformation; analogously, one can think of using a finite-element mesh for cloth simulation—certainly the cloth can be stretched flat, but that is not its expected state, and the sampling must be adequate for the variations desired. For situations where we can limit the frequency of the deformations, we can adjust the constants of proportionality appropriately to reduce the sampling rate, particularly in a LOD-hierarchy.

In summary, we have chosen a medial representation of objects that, as a consequence of its width-proportional boundary tolerances and spatial sampling, overcomes the over-precision in traditional graphics representations to which Barr referred. It also eliminates the instability of traditional medial skeletons under boundary perturbation. In the following, we give further detail on the data structures developed for medial figures and M-rep hierarchies, and discuss how model design and rendering is done using M-reps.

3 M-rep Structure

An M-rep is a compound object made up of multiple figures, typically a parent figure with child protrusion and indentation figures. This section discusses single-figure mesh connectivity and the special medial atoms needed at edges and corners, and then discusses the hierarchical modeling of M-reps from component protrusions and indentations, and giving examples of M-reps created to model anatomical organs.

3.1 Mesh Connectivity

Medial figures can be categorized according to dimensional symmetry as spheres (symmetric about a point), cylinders (symmetric about a 1D axis), or slabs (symmetric about a 2D sheet). To enable shape synthesis using these elements, we create meshes of medial atoms of different topology and connectivity to serve as primitive shape elements. For the initial work, these are

1. *spherefig*: the simplest medial figure, with a single medial primitive and with radius-proportional boundary tolerance;
2. *tubefig* and *ringfig*: 1D chains, cylindrically-symmetric about their axis, creating either a solid tube or a solid torus;
3. *quadfig*, *trifig*, and *slicefig*: 2D medial meshes of regular quadrilaterals or triangles, modeling slab-like symmetry;




Figure 5: Example medial mesh types

slicefigs are regular quad meshes with rows of medial atoms constrained to lie within planes;

4. cylinder sheet: a 2D quad mesh of medial atoms with cylindrical connectivity, creating a hollow tube with walls deformable proportionally to their thickness;
5. hollow sphere: modeling a sphere as a hollow shell with thickness $2 * r$, the medial radius, and allowing deformation of the shell proportional to its thickness.

See Fig. 5 for examples of these structures. Some additional comments:

- A spherefig is a generic object: unlike its analogue in Voronoi-based skeletal methods, a spherically-symmetric figure with a single medial atoms is stable under small boundary perturbations. Its surface may be parameterized by $[\phi, \theta]$ or by quaternion frames giving both position and rotation for subfigure placement. As with other figures, surface displacements are in the normal directions.
- Slicefigs have the drawback of modeling only the in-slice 2D medial structure and not the 3D medial structure of the object, thus limiting the types of figures that can be modeled. This has seldom been a problem in practice, and the following benefits make up for their lack of “medial purity.” (1) They can be superimposed on a corresponding plane of volume-image data, allowing one to visualize the match of M-rep to the underlying image. (2) They allow the design of objects formed by a generalized extrusion, in which a deforming medial-curve sweeps through space, possibly with twisting.
- Trifigs provide a triangular mesh for the medial sheet; triangles are used rather than a regular quad mesh to simplify work-in-progress on adaptive meshes, with interpolation of new medial atoms during interactive modeling and shape deformation.
- Mesh degeneracies are avoided due to the sparseness of the medial sampling.

3.2 Ends, Corners, and Creases

Because mesh sampling has its width-proportionality requirement, mesh sampling must be cut off at finite scale to avoid an infinite progression of finer and finer sampling at sharp discontinuities. This

implies an essentially blobby nature for M-reps in their pure form. To overcome this, we create special end-primitives for edges and corners of varying sharpness. These are medial atoms with additional parameters specifying the behavior of the boundary curvature in the edge or corner region, and allowing variation between flat, rounded, or sharply tapered behavior. In this way, we can specify variation between a blunt endcap and a first- or second-order discontinuity.

At smaller scale, edges and creases will be handled as surface deformations within given boundary tolerances, either by as displacement or bump maps, or by limiting subdivision for a subdivision surface fit to the implied boundary, as per techniques developed by DeRose[10] and Sederberg[21].

3.3 Figural Hierarchies

An M-rep is a geometric solid created by adding or subtracting protrusion or indentation figures from the composite object. In this way, M-reps can be built in CSG-fashion; though the modeling hierarchy can be smaller and simpler, due to the greater geometric richness of the medial primitives. A drawback is that object intersections are more complicated to compute, but the boundary tolerances simplify approximation and coarse-scale description of object structure and imply bounding volumes within the tolerances.

A subfigure may be attached to a parent figure by linking atoms from its medial mesh to boundary positions of the parent figure, described by associated medial locations. This requires interpolation of medial locations on the parent’s medial surface, since sampled atoms would only by chance have involutes at attachment points on the boundary. Interpolation of medial points from a sampled representation is a solved problem for 1D axes in 2D, utilizing properties of Pythagorean hodograph splines (see Farouki[12] and Yushkevich[26]). The problem of interpolating a 2D medial surface remains only partially solved and is difficult because it requires interpolation of frame fields and overconstrained object-curvature and width properties.

Because an M-rep is not isomorphic to its boundary representation, the medial structure is inherently a *lossy* representation for the 3D shape. We have already discussed the benefits of this in terms of stability, but there is a further advantage: because the tolerance-based medial structure for an object is not unique, there is inherent polymorphism in M-rep modeling. This can be seen clearly by considering that a doughnut shape may be created either by a single, toroidal figure, or by a circular, slab-figure pierced completely by an indentation figure. The choice of representation lies with the human modeler and can be made to allow particular variations in the object while avoiding discontinuous changes in the geometry. For the doughnut, one might wish on one hand to model changes in the hole, or on the other to describe warpings of the ring. As in CSG, the representation is not simply of geometry but of the process of construction.

3.4 M-rep Design in 3D Image Analysis

M-reps have already found application in medical-image 2D and 3D volume-data segmentation, registration, and volume rendering. The approach is to design an anatomic model in the form of an M-rep and then to deform it onto the image of a target patient using an objective function rewarding both image match and closeness to the model[19]. Displacement of the medially implied boundary within its tolerance collar at scan converted locations on that boundary has been shown an effective way to volume render the object while avoiding occlusion by intervening tissue[9].

This image analysis approach requires building M-rep anatomic models from a training image. Our approach is to interactively deform stock M-rep quad grids, figure by figure, into the image data




Figure 6: Human pelvis modeled by 5-figure M-rep

and then attach the grids into a hierarchy, forming a constructive solid model. An example of the result for a human pelvis is shown in Fig. 6. The M-rep for this object was made from 5 medial figures, representing three complex bones in the pelvic girdle and two bony ridges that surround the femur.

Roughly positioning the stock grid into the image data requires bending and twisting the successive rows of the grid relative to each other. We accomplish this with the M-rep deformation tools mentioned in Section 4.3, while displaying the result relative to a pre-segmented and rendered version of the anatomic object. The inter-M-figure links between an additive (protrusion) or subtractive (indentation) child M-figure and its parent M-figure are made between user-specified medial atoms that are at the open end of the child and the medial surface of the parent. The positions on the parent’s medial surface are computed by placing the open-end medial atoms of the child M-figure near the implied boundary of the parent and then projecting these atoms’ locations along the implied boundary normals of the parent.

Deforming a quad grid into image data requires visualization of the medial atoms relative to the image data. Since medical experts typically examine volume data in planar slices, we restrict the medial atoms in each row of the quad grid to being coplanar. We can then display the image plane defined by the row (see Fig. 7) and use 2D M-rep deformation to match medial structure to the image data.

4 M-reps in Computer Graphics

M-reps are rich, robust 3D-shape descriptors, well suited for use as computer graphics modeling and rendering primitives. They will find application anywhere that a solid, skeletal or blobby primitive is desirable, either for deformable models or for rigid models with shape variation between instances. Because they describe shape in a coarse-to-fine manner and are a sparse description at coarse scales, they sidestep the need for expensive model simplification algorithms. While being easily computed, the finer-scale object description is the costlier one to produce, as it should be.³

This section discusses M-reps as computer graphics primitives, in terms of boundary description, figure-subfigure blending, and model design, including coarse and fine scale modeling deformations and texturing.

³As one of our colleagues puts it, referring to traditional graphics primitives, “You want a *simpler* model!? Whoa! That’ll cost you!”




Figure 7: Matching a 2D medial figure to slice data

4.1 Modeling and Rendering: Single-figure M-rep Boundaries

A coarse polygonal tessellation is implied by the boundary-locations and connectivity of the medial mesh atoms. This gives a crude, large-scale approximation to the object shape, though failing to represent surface normals and curvatures accurately. It is nonetheless adequate for simple renderings, as in the early stages of geometric modeling, especially as the normals given by the medial atoms can be used for Phong or Gouraud shading. The guitar model in Fig. 8 has a simple polygonal tessellation with Gouraud shading.

For regular rectangular medial meshes, a finer approximation to the implied boundary can be produced by fitting Bezier patches to the surface (see Fig. 9). Control points are determined by the known boundary locations and normals/tangent-planes implied by the medial atoms. This represents both surface positions and normals where known, and gives a parametric description for them over the entire surface. Quad-meshes are easily manipulated and—whether slice-based or general—are often an ideal representation for organic or extruded objects. The creation of a parameterization of the boundary then allows us to create surface-property deformation maps, such as image-texture, bump, or displacement maps, to modify the surface within the given boundary tolerances (see Section 4.4).

Subdivision surfaces show the most promise as boundaries for M-reps. They can be generated for meshes of any of the desired topology/connectivities, and can give C^2 continuous surfaces away from extraordinary vertices. They allow for multiresolution editing—allowing deformations to be constrained to be within desired tolerance, and techniques exist for creating creases and corners of variable smoothness. Since texture coordinates can be propagated between subdivision levels, texture maps can be used as traditionally parameterized surfaces. Figure 10 shows a subdivision surface fit to the guitar body, and Figure 11 shows an image texture applied to the surface.

4.2 Modeling and Rendering: Figure-Subfigure Hierarchies

Managing the joining of figures and subfigures is a current area of research, with emphasis on using the tolerance information for the intersecting objects to guide the blend of their surfaces at the joint. Creation of a rendering interface which properly displays indentation figures is also a high priority.




Figure 8: Polygonally tessellated boundary with Gouraud shading




Figure 10: Decorated subdivision surface fit to the guitar body




Figure 9: Bezier spline fit to boundary region




Figure 11: Image texture applied to subdivision surface




Figure 12: M-rep with subdivision surface deformed by single parameter operation

4.3 Design Tools for M-rep Creation

Modeling tools for creating 3D M-reps are still in their early development stages but already show promise of the ease with which solid modeling can be done using meshes of medial atoms. The existence of a local frame with global information at each medial atom allows easy manipulation of the medial mesh and permits groupings of atoms for object-level deformations like bending and twisting. Slice-based M-reps are especially easy to work with in this way (see Fig. 12), and global deformations can be made by propagating simple rotations and scaling across the medial mesh.

Better tools will allow the user to select “fuzzy” regions of objects at various scales to extrude and deform, with the underlying medial representation invisible at the user-level. Techniques such as those developed by Rhoades[20] will allow bending and deformation operators to modify the medial meshes at variable scales, permitting intuitive, virtual sculpting operations. Tools are also required for creating the small scale deformations on the boundary surfaces, and to control blending between figure and subfigure.

For interactive modeling, we envision a dynamic mesh structure with atoms inserted or deleted to maintain the width-proportional sampling frequency. Due to the sparseness of the medial representation, this will be much less compute-intensive than typical adaptive mesh strategies as used in current model simplification schemes. The tolerance in the inter-atomic mesh links will also provide modeling hysteresis, reducing the need for atoms to be inserted into or deleted from the medial mesh and adding greater stability to the rep-

resentation under deformation.

4.4 Texturing

The fine-scale component of M-rep description requires the mapping of displacement and texture (image, procedural, etc.) onto boundary locations. The surface models we have explored for boundary interpolation all allow surface parameterization; subdivision surfaces may be parameterized by splitting and averaging texture coordinates, or by direct evaluations methods as per Stam[22].

5 Anticipated Uses in Graphics

M-reps are establishing themselves as powerful tools for object analysis, and it is our hope that they will prove equally useful to the graphics community. We here describe some areas of research that should bear fruit in the near future.

5.1 Coarse-to-Fine Design

M-reps are inherently coarse-to-fine. During modeling, we would like to interpolate the coarse level atoms to a finer grid by a successive subdivision approach. We would then specify the medial displacements at the finer grid via deformation techniques discussed in Sec. 4.3. This refinement effects a corresponding decrease in the tolerance of the implied boundary. Algorithms will needed to control the substitution of meshes of varying detail as appropriate to the scene, as per many other LOD methods.

5.2 Morphing and Other Deformation

As demonstrated by Eberly[11] in 2D, the morphing of an M-rep into another with the same figural topology is completely natural. Figure by figure, one interpolates between corresponding medial atoms in all of their dimensions of position, width, and orientations. Extensions of this to 3D hold much promise.

Stochastic deformations will allow objects to be instanced with variation from a base model according to probabilistic models. This is a simple reverse of the process in image-analysis whereby statistical models are created from M-reps fitted to a sample population.

5.3 Physically-based Dynamics

An M-rep defines a solid model, and each medial atom with its neighbors provides a shape-related solid region. The medial atoms therefore give a natural basis for finite element computations. M-reps share the advantages of other skeletal modeling techniques in their simplification of shape information, and bring their own notions of tolerance and level-of-detail to the problems of physical simulation. Moreover, M-reps provide an efficient basis for the computation of volumes, moments and other mechanical properties[24].

5.4 IBR applications

Image Based Rendering will handle occlusion effects better if it is based on a carrier for the images that generally captures the geometry of image objects in an efficient fashion. M-reps seem a natural carrier for IBR data, and their image-fitting properties will have use in IBR data extraction. Locating a specified shape, specified by an M-rep topology, in 3D image data such as range images will be enabled by the methods of Stetten [23]. M-rep techniques apply well to such reconstruction tasks, and are robust and remarkably insensitive to noise in the data.

6 Acknowledgments

The authors would like to thank Gregg Tracton for his continuing work on M-rep design tools for anatomic modeling, and Paul Yushkevich for his 2D medial modeling software. We are grateful for the help of Mary Whitton and Jeremy Ackerman in the preparation of this paper. This research has been supported by NSF SGER Grant CCR-9910419 and NIH P01 CA47982, as well as by a fellowship from the Link Foundation during 1998/99.

References

- [1] N. Amenta, M. Bern, M. Kamvysselis, A new Voronoi-based surface reconstruction algorithm. *Computer Graphics Proceedings, Annual Conference Series, Orlando, 1998, ACM SIGGRAPH*, pp. 415-422.
- [2] A. H. Barr, Ray tracing deformed surfaces. *SIGGRAPH '86 Conference Proceedings, in Computer Graphics* **20**(4): 287-296, 1986.
- [3] E. Bittar, N. Tsingos, M. Gascuel, Automatic reconstruction of unstructured 3D data: combining a medial axis and implicit surfaces. *Eurographics, conference proceedings*, 1995.
- [4] J. F. Blinn, A generalization of algebraic surface drawing. *ACM Trans. on Graphics* **1**(3): 235-256, July 1982.
- [5] J. Bloomenthal, K. Shoemake, Convolution surfaces. *SIGGRAPH '91 Conference Proceedings, in Computer Graphics*, **25**(4): 251-256, 1991.
- [6] H. Blum, A transformation for extracting new descriptor of shape. In W. Wathen-Dunn, ed., *Models for the Perception of Speech and Visual Form*, MIT Press, Cambridge, MA: 363-380, 1967.
- [7] H. Blum, R. N. Nagel, Shape description using weighted symmetric axis features. *Pattern Recognition*, **10**:167-180, 1978.
- [8] N. Burtnyk, M. Wein, Interactive skeleton techniques for enhancing motion dynamics in key frame animation. *Communications of the ACM*, **19**(10): 564-569, October 1976.
- [9] D. T. Chen, S. M. Pizer, J. T. Whitted, Using multiscale medial models to guide volume visualization. Tech report TR99-014, Dept. of CS, UNC-Chapel Hill, 1999. Available via WWW at <ftp://ftp.cs.unc.edu/pub/publications/techreports/FILE.html>.
- [10] T. DeRose, M. Kass, T. Truong, Subdivision Surfaces in Character Animation. *Computer Graphics Proceedings, Annual Conference Series, Orlando, 1998, ACM SIGGRAPH*, 85-94.
- [11] D. Eberly, Techniques for morphing two and three-dimensional objects. Internal report, Dept. of CS, UNC-Chapel Hill, 1994.
- [12] R. T. Farouki, C. A. Neff, Hermite interpolation by Pythagorean hodograph quintics. *Mathematics of Computation*, **64**(212): 1589-1609, October 1995.
- [13] P. T. Fletcher, A. Thall, D. S. Fritsch, Y. Fridman, A solid modeling program using slice-constrained medial primitives for modeling anatomical objects. Tech report TR99-035, Dept. of CS, UNC-Chapel Hill, 1999. Available via WWW at <ftp://ftp.cs.unc.edu/pub/publications/techreports/FILE.html>.
- [14] N. Gagvani, D. Kenchammana-Hosekote, D. Silver, Volume animation using the skeletal tree. *Proceedings ACM/IEEE Symposium on Volume Visualization '98*, RTP, North Carolina, pp. 47-53. 1998.
- [15] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, W. Stuetzle, Piecewise smooth surface reconstruction. *Computer Graphics Proceedings, Annual Conference Series, Orlando, 1994, ACM SIGGRAPH*, 295-302.
- [16] C. T. Loop, Smooth subdivision surfaces based on triangles. Master's Thesis, Dept. of Mathematics, Univ. of Utah, August 1987.
- [17] L. Markosian, J. M. Cohen, T. Crulli, J. Hughes, Skin: a constructive approach to modeling free-form shapes. *Computer Graphics Proceedings, Annual Conference Series, Los Angeles, 1999, ACM SIGGRAPH*, 393-400.
- [18] L. R. Nackman, S. M. Pizer, Three-dimensional shape description using the symmetric axis transform, I: Theory. *IEEE Trans. PAMI*, **7**(2): 187-202, 1985.
- [19] S. M. Pizer, D. S. Fritsch, P. A. Yushkevich, V. E. Johnson, E. L. Chaney, Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Trans. Med. Imaging*, **18**(10): 351-865, 1999.
- [20] J. Rhoades, Shaping curved surfaces. Doctoral Thesis, Dept. of CS, UNC-Chapel Hill, 1993. Available via WWW as Tech report TR93-066 at <ftp://ftp.cs.unc.edu/pub/publications/techreports/FILE.html>.
- [21] T. Sederberg, J. Zheng, D. Sewell, M. Sabin, Non-uniform recursive subdivision surfaces. *Computer Graphics Proceedings, Annual Conference Series, Orlando, 1998, ACM SIGGRAPH*, 387-394.
- [22] J. Stam, Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values. *Computer Graphics Proceedings, Annual Conference Series, Orlando, 1998, ACM SIGGRAPH*, 395-404.
- [23] G. D. Stetten, S. M. Pizer, Medial node models to identify and measure objects in real-time 3D echocardiography. *IEEE Trans. MI*, **18**(10): 1025-1034, 1999.
- [24] G. D. Stetten, Volume of arbitrary shapes from boundary curvature and medial scale. Tech report TR99-01, Dept. of CS, UNC-Chapel Hill, 1999. Available via WWW at <ftp://ftp.cs.unc.edu/pub/publications/techreports/FILE.html>.
- [25] G. Wyvill, C. McPheeters, B. Wyvill, Data structures for soft objects. *Visual Computer*, **2**(4): 227-234, 1986.
- [26] P. Yushkevich, S. M. Pizer, T. Culver, Statistical object shape via a medial representation. Tech report TR00-02, Dept. of CS, UNC-Chapel Hill, 1999. Available via WWW at <ftp://ftp.cs.unc.edu/pub/publications/techreports/FILE.html>.
- [27] D. Zorin, P. Schröder, W. Sweldens, Interpolating subdivision for meshes with arbitrary topology. *Computer Graphics Proceedings, Annual Conference Series, 1996, ACM SIGGRAPH*, pp. 189-192.