A Continuous 3-D Medial Shape Model With Branching

Timothy B. Terriberry, Guido Gerig

MICCAI 2006, Workshop on the Mathematical Foundations of Computational Anatomy, Springer Lecture Notes in Computer Science, X. Penned and S. Joshi, eds., October 2006.

Timothy Terriberry and Guido Gerig
UNC Chapel Hill

A Continuous 3-D Medial Shape Model With Branching

Timothy B. Terriberry[1] and Guido Gerig[1,2]

[1] Dept. of Computer Science
[2] Dept. of Psychiatry
Univ. of North Carolina, Chapel Hill, NC 27599, USA
{tterribe,gerig}@cs.unc.edu

Abstract. We present a new, continuously defined three-dimensional medial shape representation based on subdivision surfaces. The shape is modeled via its medial axis, and the associated boundary is computed directly from this axis at every point. Our model is parameterized over a fixed domain, so comparison among different shapes is possible. It is the first such model to support branch curves, which allows it to represent complex medial axes with more than one medial sheet.

1 Introduction

Medial shape models are powerful tools for shape analysis since they can represent complex biological shapes using a relatively small number of parameters. Defined as the locus of the centers and radii of the maximally inscribed balls of an object, they also have a natural connection with the cognitive processes of the human visual system [1]. They were originally introduced by Blum for 2-D shape analysis, where they can be organized into a tree structure with a finite number of singular points [2]. They are also well-suited for 2-D shape synthesis, as they can be modeled by a series of continuous curves with a finite number of constraints at their end-points [3]. A model of the centers and radii of the inscribed balls is sufficient to reconstruct the complete boundary of the shape.

The singular structure of the medial axis has been thoroughly analyzed in three dimensions [4], and the differential geometry of its associated boundary has undergone rigorous mathematical treatment in arbitrary dimension [5]. However, its adaptation to shape synthesis in three dimensions has proven difficult.

The primary challenge is that in 3-D, unlike in 2-D, the medial axis contains infinitely many singular points, e.g., along its edge. Each of these points requires a boundary condition to be satisfied in order for the associated boundary of the object to be closed. However, without a shape model designed explicitly to enforce these boundary conditions, a straightforward model design will only have a finite number of parameters to adjust to satisfy them.

Yushkevich introduced the first three-dimensional continuous medial model, which was based on B-spline curves, and supported only a single medial sheet [6]. Instead of trying to adjust control-point values to force the boundary conditions

to be true on the edge of the spline domain, he adjusted the radius function on the edge to force the boundary conditions to be true somewhere in the interior. Yushkevich then implicitly solved for the curve where this occurred, and defined this to be the edge of the axis.

Unfortunately, this meant that the edge of every shape model had a different location in the domain, and due to the unconstrained, irregular shape, there was no straightforward mapping between the interiors of different models. This makes the representation unsuitable for comparing different shapes or performing shape statistics across a population. It also makes it impossible to join multiple sheets, as there is no way to make the sheets intersect along a specific curve without solving high degree polynomial equations, much less obey boundary conditions.

Yushkevich et al. later describe an approach that produces an explicit domain by interpolating a potential function ρ across the medial axis instead of the radius [7]. The radius is recovered by numerically solving a differential equation with boundary conditions via finite element methods up to sufficient accuracy to get a finely sampled shape model. However, when connecting multiple sheets, using a potential function instead of an explicit representation of the radius does not leave enough free parameters to ensure both that the necessary boundary conditions are satisfied and that the radii are equal where the sheets meet. We are unaware of an existing solution to this problem.

We take a different approach. We use Catmull-Clark subdivision surfaces with what we shall term an *ordinary, corner-free boundary* to model the medial sheets. The boundary conditions are then enforced by modifying the patches at the edges of the sheet to use an interpolating spline. This allows us to reduce the degree of the polynomial equations for the boundary conditions from 12 to 2, yielding an explicit and efficient solution. The solutions of these equations are then used to construct a "control curve" which replaces the outer layer of control points, effectively providing the infinite number of free parameters needed to enforce the condition everywhere. These two key ideas—interpolating splines and control curves—are what makes this approach possible. We begin by giving a mathematical description of the medial axis and the necessary boundary conditions, and then outline our method of enforcing them.

2 Geometry of the Medial Axis

Damon formally describes the medial axis as a special kind of Whitney stratified set [5], a stratification into smooth manifold pieces of codimension one, with their boundaries and corners classified into different smooth strata of higher codimension. An associated radius function can be used to compute a multivalued *radial vector field*, S, over the axis. These radial vectors, also called *spokes*, point from a point on the axis to the corresponding points on the associated boundary, and are normal to the boundary at these points. The medial axis and its associated radial vector field are a special type of what Damon terms a *skeletal structure* [5].

For a generic genus zero object, the medial axis is composed of smooth, two-dimensional manifolds called *medial sheets*, whose boundaries form one-

dimensional curves called *edge curves*. Three sheets may be connected via a smooth one-dimensional curve called a *branch curve*. Branch curves end when they intersect an edge curve of one of the three sheets, at a *fin creation point*. Finally, four branch curves joining six medial sheets may intersect at a *six-junction point*. These are all of the singularities that may occur in three dimensions [4]. We will address them all in this work, except for six-junction points.

On a medial sheet's interior, there are exactly two values of the radial vector field, one pointing out from each side. Denoting the field S in terms of the radius and a unit radial vector field, $S = rU$, the two values of U at each smooth point and the corresponding points on the associated boundary are given by

$$U^{\pm} = -\nabla r \pm \sqrt{1 - \|\nabla r\|^2} \cdot \mathcal{N} \ . \qquad \mathcal{B} = m + rU \tag{1}$$

Here $\mathcal{N}$ is the unit normal vector for the medial sheet, and ∇r is the *Riemannian gradient* defined on the manifold. Let $m(u, v)$ be a local parameterization of the manifold and $\mathbf{I}_m$ its first fundamental form. Then ∇r is given by

$$\nabla r \triangleq \begin{bmatrix} m_u \ m_v \end{bmatrix} \mathbf{I}_m^{-1} \begin{bmatrix} r_u \\ r_v \end{bmatrix} \ , \qquad \mathbf{I}_m \triangleq \begin{bmatrix} E_m \ F_m \\ F_m \ G_m \end{bmatrix} \triangleq \begin{bmatrix} m_u \cdot m_u \ m_u \cdot m_v \\ m_v \cdot m_u \ m_v \cdot m_v \end{bmatrix} \ . \tag{2}$$

These radial vectors are arranged symmetrically about the tangent plane. One can see that the vector $U^+ + U^-$ points in the $-\nabla r$ direction, which is tangent to the medial sheet, and $U^+ - U^-$ points in the normal direction.

The values at singular points are obtained by smooth extension from the locally neighboring medial sheets. Along the edge of a medial sheet, the top spokes must meet the bottom spokes, or the surface will not be closed. As can be seen from the coefficient of $\mathcal{N}$ in (1), this occurs precisely when

$$\|\nabla r\| = 1 \ . \tag{3}$$

This "edge constraint" makes the normal component zero and the tangential components equal. Everywhere else on the medial axis, $\|\nabla r\| < 1$ must hold.

Along a branch curve, if one has three medial sheets $m^{(i)}$ oriented so that the bottom of one is adjacent to the top of the next, then the restriction becomes [3]

$$\nabla r^{(i \oplus 2)} - \nabla r^{(i \oplus 1)} = \mathcal{N}^{(i)} \cdot \sqrt{1 - \|\nabla r^{(i)}\|^2} \ , \tag{4}$$

where $\oplus$ denotes addition modulo 3. At fin creation points, $\|\nabla r^{(i)}\|$ goes to 1 for one of the sheets, causing the constraint in (4) to disappear.

A medial axis (m, r) which satisfies these boundary conditions will be closed, but there may be singularities on the reconstructed boundary itself, causing kinks and overfoldings. These occur in areas of concave curvature, where the radius extends "too far" into the concavity. Damon outlines a set of conditions to check which ensure no such illegalities occur [5]. They are based on the *radial shape operator*, S_{rad}, and a related *edge shape operator*, S_E, which measure the change in U for an infinitesimal step along a medial sheet. The radial shape operator is not a shape operator in the traditional differential geometric sense, but it *is* a

linear operator on the tangent space of m, although not necessarily self-adjoint. A matrix representation can be computed from second derivatives of m and r, and its eigenvalues, called the *principal radial curvatures*, can be used to check for overfolding of the boundary. When combined with a measure defined on the medial axis, it can also be used to compute *skeletal integrals* over the boundary or interior of the object in medial coordinates [8]. One of the advantages of a continuous representation is that this S_{rad} can be computed analytically.

3 Catmull-Clark Subdivision on the Medial Axis

Catmull-Clark subdivision surfaces are a generalization of B-spline knot insertion to meshes of arbitrary topology [9]. They were initially created as a graphics primitive to represent the boundary of closed objects. A continuous surface, called the *limit surface*, is created from an initial, discrete mesh by recursive subdivision. After the first level of subdivision, all of the faces are quadrilaterals, and every new vertex will have valence four.

Any vertex of valence four is thus called an *ordinary* vertex, while the remaining vertices are *extraordinary*. After the second level of subdivision, every face has at most one extraordinary vertex. Away from an extraordinary vertex, the limit surface behaves exactly like a B-spline patch, and thus the surface and its derivatives can be evaluated efficiently at arbitrary points. Stam showed how, with some one-time setup, the limit surface could also be efficiently evaluated near an extraordinary vertex [10].

Catmull-Clark surfaces are everywhere C^2 continuous, except at extraordinary vertices, where they are C^1 continuous. C^1 continuity ensures that ∇r, and thus the spoke field, are continuous everywhere. Since the spoke field is normal to the boundary, a continuous spoke field will generate a G^1 continuous boundary. C^2 continuity everywhere except on a set of measure zero ensures that we can compute S_{rad} almost everywhere, and thus can check for illegalities and compute medial integrals. Subdivision surfaces are also appealing because they offer *local control*. That is, a control point only influences a small local region of the surface surrounding it, instead of the whole surface. This becomes advantageous when trying to fit a model to an image, since the derivative of most points on the boundary with respect to one of the parameters is zero.

Originally restricted to closed surfaces, there have been several strategies for handling meshes with edges and creases. Initially, Hoppe et al. proposed a set of rules for triangular subdivision surfaces [11], which have a straightforward adaptation to quadrilateral subdivision surfaces—including Catmull-Clark surfaces—as described by, e.g., DeRose et al. [12] and Warren and Schaeffer [13]. The latter also describe a simple method of implementing these rules, however neither proved that the rules produced C^1 limit surfaces. A set of rules that are provably C^1 everywhere were proposed by Biermann et al., addressing problems with extraordinary edge vertices and convex and concave corners [14].

For simplicity, we require that the edge of the mesh contain only ordinary vertices (which will be valence three, not four). For reasons that will become

clear in the next section, we also disallow corners, i.e., no two adjacent edges on a crease or edge curve can belong to the same face. Besides being a practical requirement, eliminating corners is also desirable for many of the objects we wish to model. Under these restrictions, all of the previous rules for handling edges and creases are identical. They effectively cause the edge of the mesh to be interpolated as a one-dimensional B-spline curve, unaffected by other control points in the interior of the mesh. This makes joining three medial sheets along a branch curve straightforward. At a fin creation point, the fin still ends in a corner and a *dart vertex* is used to merge the crease smoothly into the surface.

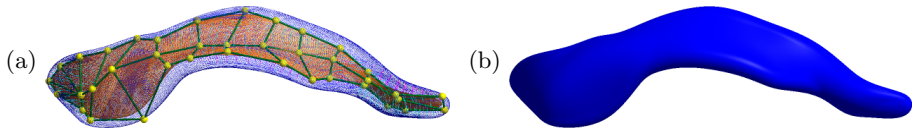

Fig. 1. A continuous medial model defined on a mesh with ordinary, corner-free boundary, depicting a healthy left lateral ventricle. (a) The mesh, medial axis, and radial vector field. (b) The enclosing boundary.

We call the result a mesh with ordinary, corner-free boundary. We perform two levels of subdivision, and then construct B-spline patches, except near extraordinary vertices, which are handled via Stam's evaluation method. The next two sections address patches adjacent to an edge or branch curve.

4 Edge Patches

First, we address patches along the edge of a medial sheet. Without loss of generality, we assume the edge lies along the line $u = 0$, as illustrated in Figure 2(a).

A full 4×4 matrix of B-spline control points $\mathbf{P}$ that interpolate the limit surface $\mathcal{M} \triangleq (m, r)$ of this patch for both m and r can be constructed by adding $P_{0,j}$ points out past the boundary, where[3] $P_{0,j} \triangleq 2P_{1,j} - P_{2,j}$. Then the local portion of the medial locus is defined by the equation

$$\mathcal{M}(u,v) \triangleq \begin{bmatrix} 1 & u & u^2 & u^3 \end{bmatrix} \mathbf{B}\mathbf{P}\mathbf{B}^T \begin{bmatrix} 1 & v & v^2 & v^3 \end{bmatrix}^T \quad , \quad \mathbf{B} \triangleq \frac{1}{6} \begin{bmatrix} 1 & 4 & 1 & 0 \\ -3 & 0 & 3 & 0 \\ 3 & -6 & 3 & 0 \\ -1 & 3 & -3 & 1 \end{bmatrix} \quad . \tag{5}$$

The medial sheet m is interpolated as usual, but in order to enforce the edge condition (3) along this edge, we interpolate r using a *control curve* $r_0(v)$ instead

[3] Using $P_{0,j} \triangleq \frac{1}{2}(P_{1,j} + P_{2,j})$ and replacing $P_{1,j}$ with $\tilde{P}_{1,j} \triangleq \frac{11}{8}P_{1,j} - \frac{3}{8}P_{2,j}$ helps prevent the right side of the patch from overfolding, at the cost of a small modification to the neighboring patch, while interpolating the same $u = 0$ edge.

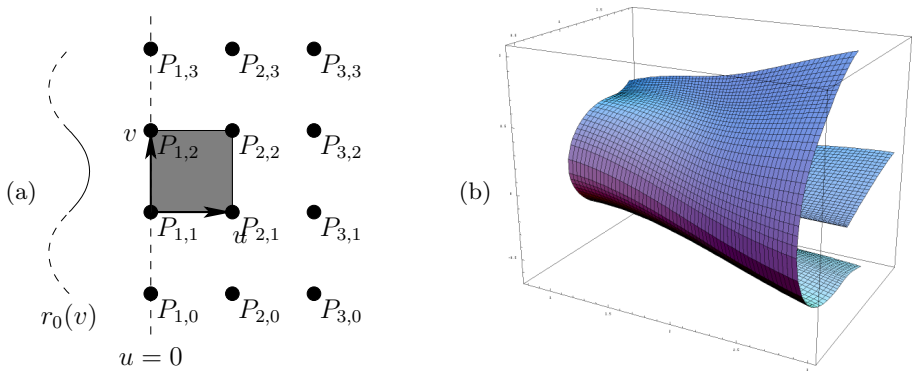

Fig. 2. (a) The control curve, $r_0(v)$, on an edge patch. (b) An example of interpolated edge patches and their associated boundary.

of a few isolated control points. This lets us ensure the condition holds at every point on the edge, which would not be possible with only a finite number of control point values as free parameters. It is important to emphasize that this curve does not represent a set of parameters explicitly defined by the modeler, but is implicitly defined by the boundary condition.

Unfortunately, the use of such a curve makes the interpolation fundamentally asymmetric with respect to the two variables u and v. We must first interpolate in the v direction, obtaining new control points for a spline in u. Then, we use a change of basis to convert the spline in u to an *interpolating spline*, which passes through its control points. Finally, we replace the control point on the left with the point from our control curve, and perform interpolation in the u direction. Because of this asymmetry, we cannot, for example, add a control curve to both the $u = 0$ and $v = 0$ edges of the same patch, which is why we disallow corners on the edge of the mesh.

The particular interpolating spline we use is the C^1 Catmull-Rom spline. We perform the change of basis by modifying the spline defined above. Replacing $\mathbf{BP}$ in equation (5) by $\mathbf{CP'}$, we use a new set of control points, $\mathbf{P'}$, given by

$$\mathbf{P'} \triangleq \mathbf{C}^{-1}\mathbf{BP} \ , \qquad \mathbf{C} \triangleq \frac{1}{2}\begin{bmatrix} 0 & 2 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ 2 & -5 & 4 & -1 \\ -1 & 3 & -3 & 1 \end{bmatrix} . \tag{6}$$

Now, we can replace the row $P'_{0,j}$ with our control curve, and the result will still pass through the limit surface at either end of the patch. Furthermore, note that $P'_{0,j}$ has no influence on the derivative with respect to u on the right edge of the patch. Hence, whatever we use for a control curve, we will retain C^1 continuity along the right edge. We now derive an equation for the control curve, and show that we also retain C^1 continuity along the top and bottom edges.

4.1 Solving for $r_u(0, v)$

In order to ensure that the top and bottom spokes meet, we proceed to enforce the constraint in (3). With some algebra, we can decompose ∇r into a component in the m_v direction and a component in the $m_v \times (m_u \times m_v)$ direction:

$$\nabla_r^{(v)} = \frac{r_v}{\sqrt{G_m}} , \qquad \nabla_r^{(\perp v)} = \frac{r_u G_m - r_v F_m}{\sqrt{G_m(E_m G_m - F_m^2)}} . \tag{7}$$

Now, we hold $m_u(0, v)$, $m_v(0, v)$, and $r_v(0, v)$ fixed, and solve for $r_u(0, v)$. The key observation is that using a spline interpolative in u means that $r(0, v)$, and hence $r_v(0, v)$, is completely determined by the control points $P'_{1,j}$ along the edge. In particular, they do *not* depend on the values of our unknown control curve. Then substituting (7) into (3) and solving for r_u produces

$$r_u = \frac{1}{G_m} \left(r_v F_m \pm \sqrt{(G_m - r_v^2)(E_m G_m - F_m^2)} \right) . \tag{8}$$

There are two possible solutions, one of which corresponds to spokes along the crest pointing outward from the medial sheet, and the other corresponding to spokes pointing inwards. The latter is clearly illegal for a Blum medial axis, but as we can see by substituting (8) into (7), the sign of the component perpendicular to m_v is determined entirely by the sign of the square root term. The plus solution always corresponds to outward-pointing spokes.

4.2 The Complete Control Curve

Now, given $r_u(0, v)$, it is a simple matter to solve for the value of the control curve $r_0(v)$ that produces this derivative. Let $r_1(v)$, $r_2(v)$, and $r_3(v)$ be the interpolation of the three lines of ordinary control points in the v direction:

$$r_i(v) \triangleq \left[P'_{i,0}\ P'_{i,1}\ P'_{i,2}\ P'_{i,3}\right] \mathbf{B}^T \left[1\ v\ v^2\ v^3\right]^T , \qquad i = 1 \ldots 3 . \tag{9}$$

Then solving for $r_0(v)$ in terms of $r_i(v)$ yields

$$r_0(v) = \frac{1}{C_{1,0}} \left(r_u(0, v) - C_{1,1} r_1(v) - C_{1,2} r_2(v) - C_{1,3} r_3(v) \right) , \tag{10}$$

where $C_{i,j}$ is the (0-indexed) i, jth element of $\mathbf{C}$. It is important to emphasize that this formulation only works for interpolating splines, since otherwise $r_u(0, v)$ would not be independent of $r_0(v)$ in the right-hand side of (10).

Note that $r_0(v)$ is a function of $m_u(0, v)$, $m_v(0, v)$, and $r_v(0, v)$ only. As these are all first derivatives of functions obtained from B-spline interpolation, they are C^1 across patches, and hence so is $r_0(v)$. This ensures that $r(u, v)$ is C^1 with neighboring patches along the top and bottom edges, as well as the right edge. In practice, we can regain C^2 continuity on the right by using a fourth-order C^2 interpolating cubic spline (omitted for brevity). Such a spline even contains enough free parameters to enforce a second boundary condition on the same edge, although we do not make use of that in this work.

5 Branch Curve Patches

Although (4) is a succinct expression of the boundary conditions along a branch curve, it is not obvious how they should be enforced. We motivate the solution with a geometric approach. Without loss of generality, we assume our three patches are oriented so that they meet at the $u = 0$ curve in each. Then r_v and G_m are the same in each patch, and hence so is $\nabla r^{(v)}$.

Now, since all three $\mathcal{N}^{(i)}$s are perpendicular to m_v, they lie in the same plane, and hence so do the endpoints of $\nabla r^{(i)}$ and $U^{(i)\pm}$. We project everything into this plane, as illustrated in Figure 3(a). As one can see, the sums of the angles $\theta^{(i)}$ must be π, and the in-plane components of the tangent vectors $m_u^{(i)}$ must bisect these angles. This provides a full set of geometric constraints.

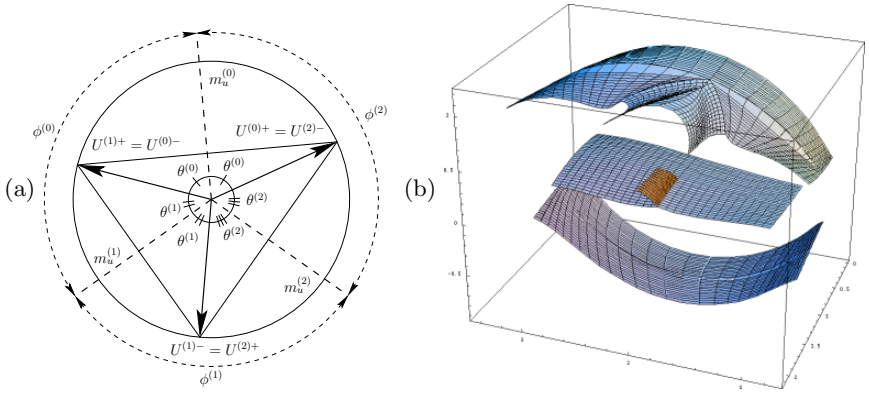

Fig. 3. (a) The cross-section of a branch curve, with m_v pointing out of the page. (b) An example of the interpolated patches near a fin creation point.

5.1 Away From Fin Creation Points

Away from a fin creation point, we can enforce these constraints by noting that $\nabla r^{(\perp v,i)} \propto \cos\theta^{(i)}$. Let $\phi^{(i)}$ be the counter-clockwise angle between the in-plane components of the tangent vectors $m_u^{(i)}$ and $m_u^{(i+1)}$, so $\phi^{(i)} = \theta^{(i)} + \theta^{(i+1)}$. Then solving for $\theta^{(i)}$ and $\nabla r^{(\perp v,i)}$ gives

$$\theta^{(i)} = \frac{1}{2}(\phi^{(i)} + \phi^{(i\oplus 2)} - \phi^{(i\oplus 1)}) \ , \qquad \nabla r^{(\perp v,i)} = -\sqrt{1 - \frac{r_v^2}{G_m}}\cos\theta^{(i)} \ . \tag{11}$$

These constraints can be enforced with a control curve exactly as in Section 4.

5.2 At a Fin Creation Point

Catmull-Clark subdivision does not produce B-spline patches around a dart vertex, but after the first two levels of subdivision, we approximate the medial surface with B-spline patches that interpolate the same branch curve. Without loss of generality, we will assume $m^{(0)}$ is the fin patch. At the fin point itself, we adjust $r_u^{(0)}(0,0)$ while holding $r^{(0)}(0,v)$ fixed so that $\|\nabla r^{(0)}\| = 1$ at $(0,0)$. Since this is a constraint at just one point, this is easily done by adjusting control points. We also note that we can rotate $m_u^{(i)}$ around m_v arbitrarily without affecting $\|\nabla r^{(i)}\|$. Thus, we rotate $m_u^{(0)}(0,0)$ around m_v until $U^{(0)-}$ and $U^{(1)+}$ coincide. This can also be done without affecting $m^{(0)}(0,v)$ by adjusting control points. We now enforce the $\|\nabla r^{(0)}\| = 1$ condition on the entire $v = 0$ edge as normal. However, this prevents us from using a control curve along the other edge. Hence, we hold the values of $r_u^{(0)}(0,v)$ and $m_u^{(0)}(0,v)$ fixed.

This fixes $\theta^{(0)}$, so we adjust $\theta^{(1)}$ and $\theta^{(2)}$. Let $\theta^{(1)} = \phi^{(0)} - \theta^{(0)} - \alpha$ and $\theta^{(2)} = \phi^{(2)} - \theta^{(0)} - \beta$. We require $\theta^{(0)} + \theta^{(1)} + \theta^{(2)} = \pi$ and for a second constraint choose $\alpha\phi^{(2)} = \beta\phi^{(0)}$, yielding

$$\alpha = \phi^{(0)} \cdot \frac{\phi^{(0)} + \phi^{(2)} - \pi - \theta^{(0)}}{\phi^{(0)} + \phi^{(2)}} \ , \qquad \beta = \phi^{(2)} \cdot \frac{\phi^{(0)} + \phi^{(2)} - \pi - \theta^{(0)}}{\phi^{(0)} + \phi^{(2)}} \ . \tag{12}$$

To make the tangent planes form the correct angles between them, we rotate $m_u^{(1)}$ by α and $m_u^{(2)}$ by $-\beta$ around m_v. This can be accomplished with a control curve in a way exactly analogous to the procedure in Section 4.2 for r_u.

At the fin creation point, $\phi^{(0)}(0) + \phi^{(2)}(0) = \pi$ and $\theta^{(0)}(0) = 0$, so $\alpha(0) = \beta(0) = 0$. However, to ensure C^1 continuity, we require $\alpha_v(0) = \beta_v(0) = 0$. This occurs precisely when $\phi_v^{(0)}(0) + \phi_v^{(2)}(0) + \theta_v^{(0)}(0) = 0$. Since $\phi^{(0)} + \phi^{(2)} = \pi$ past the fin creation point, and these are functions of the unrotated $m^{(i)}$s, which join the rest of the surface with C^2 continuity at $v = 0$, this implies $\phi_v^{(0)}(0) + \phi_v^{(2)}(0) = 0$, and hence $\theta_v^{(0)}(0) = 0$ must hold. This can be enforced by adjusting control points to modify $r_{uu}^{(0)}(0,0)$ so that $\frac{\partial \|\nabla r\|^2}{\partial v} = 0$ at the fin creation point, while holding $r_u^{(0)}(0,0)$ and $r^{(0)}(0,v)$ fixed.

To transition between the procedure at a fin creation point and the procedure elsewhere on the branch curve, we smoothly blend between the two strategies in an adjacent patch using the weight function $\omega(v) = 1 - 3v^2 + 2v^3$. Note that $\omega(0) = 1$, $\omega(1) = 0$, and $\omega_v(0) = \omega_v(1) = 0$. Then in the transition region we use

$$\theta^{(i)} = \omega \cdot \theta^{(i)}_{\text{fin}} + (1 - \omega) \cdot \theta^{(i)}_{\text{branch}} \ , \qquad \alpha = \omega \cdot \alpha_{\text{fin}} \ , \qquad \beta = \omega \cdot \beta_{\text{fin}} \ . \tag{13}$$

6 Conclusion

We have presented a generative shape model using the medial axis in three dimensions. It is based on cubic splines defined on subdivision surfaces, and so can be implemented efficiently. It is the first such model to support branching. In future work, we also plan to incorporate six-junction points.

Acknowledgments

The research was supported by the NIH NIBIB grant P01 EB002779. The authors would also like to thank Paul Yushkevich and Jim Damon for several enlightening discussions, and the reviewers for their helpful comments.

References

1. Burbeck, C.A., Pizer, S.M., Morse, B.S., Ariely, D., Zauberman, G.S., Rolland, J.P.: Linking object boundaries at scale: a common mechanism for size and shape judgements. Vision Research **36** (1996) 361–372
2. Blum, H.: A transformation for extracting new descriptors of shape. In Wathen-Dunn, W., ed.: Models for the Perception of Speech and Visual Form, Cambridge, MA, MIT Press (1967) 362–380
3. Yushkevich, P.A., Fletcher, P.T., Joshi, S.C., Thall, A., Pizer, S.M.: Continuous medial representations for geometric object modeling in 2D and 3D. Image and Vision Computing **21** (2003) 17–28
4. Giblin, P.J., Kimia, B.B.: A formal classification of 3D medial axis points and their local geometry. IEEE Transactions on Pattern Analysis and Machine Intelligence **26** (2004) 238–251
5. Damon, J.N.: Smoothness and geometry of boundaries associated to skeletal structures I: Sufficient conditions for smoothness. Annales de l'Institut Fourier **53** (2003) 1941–1985
6. Yushkevich, P.A.: Statistical shape characterization using the medial representation. PhD thesis, University of North Carolina at Chapel Hill Department of Computer Science (2003)
7. Yushkevich, P.A., Zhang, H., Gee, J.C.: Parametric medial shape representation in 3-D via the Poisson partial differential equation with non-linear boundary conditions. In Christensen, G.E., Sonka, M., eds.: Information Processing in Medical Imaging (IPMI'05). (2005) 162–173
8. Damon, J.N.: Global geometry of regions and boundaries via skeletal and medial integrals. Preprint: http://www.math.unc.edu/Faculty/jndamon/Skel.Str.IV.v5.pdf (2005)
9. Catmull, E.E., Clark, J.H.: Recursively generated B-spline surfaces on arbitrary topological meshes. Computer Aided Design **10** (1978) 350–355
10. Stam, J.: Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values. Computer Graphics **32** (1998) 395–404
11. Hoppe, H., DeRose, T., Duchamp, T., Halstead, M., Jin, H., McDonald, J., Scheitzer, J., Stuetzle, W.: Piecewise smooth surface reconstruction. Technical Report TR 94-01-01, University of Washington Department of Computer Science and Engineering (1994)
12. DeRose, T., Kass, M., Truong, T.: Subdivision surfaces in character animation. In: Proceedings of the 25th Annual SIGGRAPH Confererence on Computer Graphics and Interactive Techniques. (1998) 85–94
13. Warren, J., Schaefer, S.: A factored approach to subdivision surfaces. IEEE Computer Graphics and Applications **24** (2004) 74–81
14. Biermann, H., Levin, A., Zorin, D.: Peicewise smooth subdivision surfaces with normal control. In: Proceedings of the 27th Annual SIGGRAPH Confererence on Computer Graphics and Interactive Techniques. (2000) 113–120