



Adaptive Implicit Modeling using Subdivision Curves and Surfaces as Skeletons

Alexis Angelidis, Marie-Paule Cani

► To cite this version:

Alexis Angelidis, Marie-Paule Cani. Adaptive Implicit Modeling using Subdivision Curves and Surfaces as Skeletons. The seventh Symposium on Solid Modeling and Applications (SMA'02), ACM, Jun 2002, Saarbrücken, Germany. pp.45 - 52, 10.1145/566282.566292 . inria-00537494

HAL Id: inria-00537494

<https://inria.hal.science/inria-00537494>

Submitted on 18 Nov 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Adaptive Implicit Modeling using Subdivision Curves and Surfaces as Skeletons

Alexis Angelidis

Marie-Paule Cani

iMAGIS[†]-GRAVIR, INRIA Rhône-Alpes
655 avenue de l'Europe, 38330 Montbonnot, France

ABSTRACT

Recent work has shown that implicit modeling using levels of details (LODs) is possible thanks to the use of subdivision-curves as skeletons. However, the geometric skeleton of a 3D shape is, in the general case, a graph of interconnected curve segments and surface patches, the exclusive use of curve skeletons is therefore not sufficient.

We present a model that uses a graph of interconnected subdivision curves and surfaces as a skeleton, on which a varying radius can be specified in order to control surface thickness. The subdivision levels of the skeleton provide levels of detail for the field function that defines the implicit surface.

Its visualization is achieved by generating a coarse mesh that surrounds the skeleton. At high valence skeleton vertices, triangulation topology issues are managed by locally overlapping the iso-surface triangulations. The mesh is then adaptively refined in order to sample the current LOD of the implicit surface within an error tolerance.

The last contribution is a new solution to the unwanted blending problem. It avoids blending between parts of the surface that do not correspond to neighboring skeletal elements, and ensures C^1 continuity everywhere. All these methods are integrated into an interactive modeling system, where the user can create, view and edit complex shapes at different levels of detail.

Categories and Subject Descriptors

I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—*Curve, surface, solid, and object representations*

General Terms

Algorithms

[†]iMAGIS is a joint project of CNRS, INRIA, Institut National Polytechnique de Grenoble and Université Joseph Fourier.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SM'02, June 17-21, 2002, Saarbrücken, Germany.

Copyright 2002 ACM 1-58113-506-8/02/0006 ...\$5.00.

Keywords

Blends, sweeps, offsets & deformations, Heterogeneous models, Multi resolution models

1. INTRODUCTION

An implicit surface is defined as the set of points p verifying an equation $f(p) = c$, where f is called *field function*. Modeling shapes with implicit surfaces is an elegant solution to represent smooth objects with branchings [4]. However, their use has been restricted, as they can't easily be adapted to interactive modeling. The characteristic that made popular other models such as subdivision surfaces is the possibility to render them at different levels of detail (LOD).

Therefore, recent work has provided LODs for implicit surfaces. These LODs can be defined on both the field function [6] and on surface sampling. However, in this previous work, only implicit surfaces generated by curve skeletons were considered, which was severely restricting the range of shapes being modeled.

A second feature needed for interactive modeling with implicit surfaces is the control of the locality of the blend: it helps controlling the topology of the shape being modeled; local blending allows moreover a local control of the shape being modeled. The latter does not need to be entirely recomputed when an element is modified. Previous solutions in this area aren't fully satisfying: in the case of curve skeletons for instance, the skeleton needs to do large folds to prevent the surface from blending back with itself [6].

In this paper, a general solution to the interactive modeling and display of implicit surfaces at different levels of detail is presented. It relies on interconnected subdivision curves and surfaces for generating a field function at different levels of refinement. The resulting subdivision implicit surface is displayed using a set of locally adaptive closed meshes that exploit temporal coherence during refinement and editing operations, and adapt their resolution to the local underlying curvature, thus providing a second control on levels of detail. The convenience of the resulting modeling system is illustrated by the interactive modeling of complex organic shapes. Their quality is enhanced by the use of a new solution for avoiding unwanted blending, which ensures C^1 continuity everywhere, and can be used even when a skeletal element directly folds back onto itself.

Section 2 discusses related works. Our first contributions, described in Section 3, is the definition of an implicit primitive generated by subdivision surface, which generalizes the subdivision-curve implicit primitive [6]. The problem of

defining a varying surface radius along these primitives is addressed. Section 4 describes how to visualize with a piecewise triangulation the implicit shape. Section 5 introduces a new solution to the unwanted blending issue. Section 6 explains how these methods are integrated into an interactive modeling system, and presents results. Section 7 concludes and discusses the possible extensions of this work.

2. RELATED WORKS

2.1 Skeleton Defined Surfaces

Among the various categories of implicit surfaces (algebraic [18, 2], sampled field [11, 10, 15], skeleton defined [13]), the latter category provides many advantages. Firstly, in a modeling context, combining many simple primitives is an easy and intuitive way of modeling complex shapes. Secondly, this representation is the most adapted to animation and deformation: the skeleton gives a first intuitive representation of the iso-surface; it is used as an underlying control primitive for the object. Lastly, without considering surface visualization issues, its memory storage is cheap. Intuitively, the surface generated by a skeleton corresponds to the skeleton’s coating. The coating’s global thickness is controlled with both the iso-value and the speed at which the field decreases.

However, using the skeleton as a topology graph to control blending has seldom been exploited.

2.2 Blending

Smooth blending is one the main advantages of implicit surfaces. It can be performed by simply summing the field functions. When two skeletons of fields f_i and f_j stand close enough, the generated iso-surfaces blend, resulting into a smooth and organic looking shape.

There are two problems with blending: unwanted blending, and unwanted bulges at junctions.

2.2.1 Unwanted Blending

To model complex shapes, blending must be controlled in order to prevent all iso-surfaces from blending with each other. Previous solutions to the unwanted blending problem [14, 7], did not maintain C^1 continuity of the shape everywhere, since the field at a given point was defined as the maximal contribution from groups of skeletons that blend together. This solution results in a union of volumes, possibly creating tangent discontinuities in regions where blending properties change. The problem has only been partially solved [6], as the skeleton needs to do large folds not to blend again when parts come close to each other again.

2.2.2 Unwanted Bulges

The problem of unwanted bulges appeared with the generalization of simple punctual skeletons [3] to more complex skeletons. This generalization can be done in two ways: the first defines the field generated by a skeleton S at a point p as the punctual field of the closest point of S to p . In this case, every time the skeleton is subdivided, new punctual fields are added, thus creating bulges in areas where several primitives join, as shown in Figure 1.

A second way of generalizing punctual skeletons consist in summing all the punctual fields of the points of the skeleton S . This is done by integrating a punctual field over S [5]. Then, splitting a skeleton segment into two parts does not




Figure 1: Bulge resulting from the subdivision of a distance surface skeleton into two pieces.

introduce any bulge (Figure 2), and moving the newly inserted point smoothly deforms the iso-surface.

This property of convolution surfaces is of great interest to define the subdivision implicit primitives.




Figure 2: Conservation of the volume and shape while subdividing the skeleton of a convolution surface.

2.3 Subdivision-curve implicit primitive

Previous work [6] has proposed a LOD paradigm for the implicit field function, using subdivision curve primitives in combination with convolution surfaces. Subdivision is an iterative process that refines the curve at each step, and makes the polyline converge geometrically to a smooth curve. These polylines are used as skeletons for generating an implicit surface, thus providing the LODs on the field function.

The subdivision scheme can be defined with masks, describing how an initial vertex is moved, and how additional vertices are inserted (see for instance [17]). The limit curve can either approximate or interpolate the initial control points.

A second kind of LOD, defined on the surface’s sampling has been also suggested [6]. It has recently been extended to a locally adaptive iso-surface sampling [1]. This is done in three steps: the skeleton is cut into pieces that can be easily triangulated; a coarse triangulation is assigned to each piece; then the triangulations refines according to an error criterion.

Despite of their interest for modeling organic shapes, implicit surfaces generated by subdivision curves can only be used in a restricted range of applications; indeed, modeling surfaces with planar regions requires the use of a surfacic skeletal element. Thus, our first contribution presented next is the extension of the model to subdivision-surface skeletons.

3. THE SUBDIVISION-SURFACE IMPLICIT PRIMITIVE

Subdivision-surfaces offer the same advantages as subdivision-curves for defining skeletons that are used to generate convolution surfaces at different LODs. Subdivision schemes have been developed for polygonal surfaces made either of quadrilateral or triangle meshes. In this paper, we focus on the latter, since the only known closed-form solution for convolution along surface elements holds for triangles [16].

Subdivision schemes are more complex for triangulated surfaces than for curves. Indeed, the number of connected edges at vertices may vary from one vertex to another. Moreover, edges can either stand on surface's border or be embedded in the surface. The surface border is refined using a curve subdivision scheme. For surface embedded vertices, two types of vertices are usually distinguished: regular vertices, which are the vertices of valence 6, and irregular vertices. As for the curves, the subdivision scheme can either interpolate the initial vertices or approximate them.

3.1 Approximating Subdivision Scheme

In Loop's approximation scheme [12] the border of the surface is subdivided using the cubic uniform B-spline scheme. On the surface, this scheme produces a C^2 continuous surface, except at extraordinary vertices where surface is C^1 . Regular and irregular vertices are treated using the same mask, shown on Figure 3, where $\beta = \frac{1}{n}(\frac{5}{8} - (\frac{3}{8} + \frac{1}{4} \cos \frac{2\pi}{n})^2)$.




Figure 3: Loop subdivision masks

3.2 Interpolating Subdivision Scheme

The modified butterfly scheme [8] interpolates the initial vertices. New vertices on the border of the surface are inserted using the DLG mask [9]. For an edge between two regular vertices, the new vertex is inserted using the mask depicted in Figure 4 (a). For an edge between a regular and an irregular vertex, mask Figure 4 (b) is used. Else, the average of the values computed by the mask of Figure 4 (b) for each endpoint is taken. Coefficients are: if $k = 3$, $s_0 = \frac{5}{12}$, $s_{1,2} = -\frac{1}{12}$, if $k = 4$, $s_0 = \frac{3}{8}$, $s_{1,3} = 0$, $s_2 = -\frac{1}{8}$, if $k > 5$, $s_i = \frac{1}{k}(\frac{1}{4} + \cos \frac{2i\pi}{k} + \frac{1}{2} \cos \frac{4i\pi}{k})$.




Figure 4: Modified butterfly subdivision masks

3.3 Varying Thickness Implicit Surfaces

Let us consider the implicit surface generated at different LODs using a subdivision surface skeleton. The assignment of a varying radius along the skeleton allows to model a larger set of objects. As in previous works [6], a varying radius is assigned at each vertex, and computed vertices are

being assigned a radius value using the subdivision masks. Considering a single primitive, the field function generated by a segment or a triangle is not only a function of the skeleton vertices positions, but also a function of radii.

The proposed solution by Cani and Hornus [6], which consisted in interpolating linearly the radius along the segment, and multiplying the field at a point p by the radius of the projection of p on the skeleton is not satisfying, as it does not allow to control precisely the decreasing of the field, as shown in Figure 5.

We propose a simple solution that converges toward a varying radius field as quick as the implicit surface's skeleton converge towards a smooth skeleton through the subdivision process. Also, it allows the use of any non-varying radius field. The solution consists in using a constant radius, computed as the average of the vertices radii along each individual segment or triangle:




Figure 5: Iso-surface produced by the segment field proposed in [6]. The left vertex radius is the highest, the middle vertex radius is almost null. A bulge appears on the right segment, while the iso-surface on the left side takes off of the segment.

$$f_{segment}(p, v_0, v_1, r_0, r_1) = \frac{r_0 + r_1}{2} f_{segment}(p, v_0, v_1)$$




Figure 6: Varying radius fields obtained with skeleton subdivision.

$$f_{triangle}(p, v_0, v_1, v_2, r_0, r_1, r_2) =$$

$$\frac{r_0 + r_1 + r_2}{3} f_{triangle}(p, v_0, v_1, v_2)$$

Figure 6 shows the convergence of the resulting implicit surface toward a convolution surface with a smoothly varying radius, as the skeleton curve or surface refines. Thus

the subdivision process refines the curvature as well as the thickness of the implicit shape. This has however a cost, as each subdivision of the skeleton increases by 4 the number of implicit primitives, making the modeling 4 times slower.

Our modeling proposes to model implicit surfaces using a skeleton defined by a graph of branching curve and surface pieces. Underlying the control skeleton is a refined skeleton whose level of subdivision provides LODs on the implicit function, that affect both curvature and thickness of the implicit shape. In order to avoid the appearance of unwanted bulges when refining the skeleton, convolution surfaces are used; the subdivision is defined such that non-varying radius fields can be used.

Before describing how these primitives blend, we present our method for ensuring an interactive display during modeling sessions.

4. INTERACTIVE DISPLAY

Previous works have proposed a solution to visualize interactively the iso-surface generated by subdivision-curve primitives [6, 1].

In this paper, we adapt to subdivision surfaces the adaptive triangulation method defined in [1] for subdivision curve primitives. We still use the three steps summarized in Sub-Section 2.3: the skeleton is cut into pieces that can be easily polygonised; each piece of skeleton is being assigned a rough triangulation; the coarse triangulation is then locally refined using an error criterion.

4.1 Rough Local mesh

The implicit surfaces we consider are generated from graphs of branching curves and surfaces. Let us consider the coarsest version of the skeleton graph. The idea is to associate a local closed mesh to each maximal connected polyline or piece of surface in this coarse graph. Each vertex of the mesh will slide toward the iso-surface along an axis defined in the skeleton's coordinate set. See Figure 7.




Figure 7: A coarse skeleton made of two surface pieces and five curve pieces.

More precisely, the skeleton pieces are defined as:

- a maximal set of control segments connected by vertices of valence 2 (curve piece).
- a maximal set of control triangles sharing an edge (surface piece).

A coarse closed mesh is associated with each skeleton piece as shown in Figures 8 and 9. Mesh vertices are attached to fixed axes relative to their parent primitive in the skeleton. For vertices that “belong” to two joining line-segments, or to two neighboring triangles (as in Figures 8 and 9 right), the axes are defined in the bisecting plane between the two primitives, or in a plane perpendicular to the primitive if there is no neighbor. Additional axes are defined for the segment (resp. the triangle) at the extremities (resp. the edges and vertices).




Figure 8: Construction of a closed mesh around a maximal curve-piece in the coarse skeleton.




Figure 9: Construction of a closed mesh around a maximal surface-piece in the coarse skeleton. (a) Attach axes for a single triangle. (b) Single triangle's coarse mesh. (c) Two triangle coarse mesh. (d) Three triangle coarse mesh.

Once the axes are defined, the iso-surface sampling points converges along their axis toward the iso-surface. This is done using a Newton-Raphson minimization on the function:

$$F(t) = (f(s_i + t * d_i) - iso)^2$$

where f is the field function, and s_i and d_i are the starting point and direction of axis i . There can be more than one solution to this equation. During the modeling process however, time coherence is of great help to keep the right

solution when the skeleton is being deformed. The same equation holds to find the solution at initialization, as long as the newly inserted primitive does not stand too close to the skeleton to which it has been connected.

4.2 Refinement

The rough mesh is attached in the local coordinate set of the coarse skeleton, so as the newly inserted iso-surface triangles when iso-surface sampling is refined. This enables to refine either the skeleton, or the iso-surface sampling while using temporal coherence.

The iso-surface sampling adaptive refinement is performed in the following way: a triangle edge is subdivided if the difference between the field value at its midpoint and the iso-value iso is greater than a constant η (which may be view-dependent to refine areas of interest). See Figure 11.

$$\epsilon = \left| f\left(\frac{v_0 + v_1}{2}\right) - iso \right|$$

Several iso-surface samplings of a spaceship model are shown in Figure 10. It is to note that the triangulation is higher in curved areas.




Figure 10: Spaceship model refined with decreasing values of η




Figure 11: Iso-surface refinement. (a) At least one edge does not satisfy the criterion. (b) The triangle is splitted. (c) Edges that do not need refinement are snapped, other sample the iso-surface.

5. AVOIDING UNWANTED BLENDING

A first reason for controlling the blending of the fields generated by different primitives is to control the topology of the final shape. A second motivation is to control the locality of the implicit surface skeleton edition in an interactive modeling context, in order to restrict updates to local, well defined deformed regions.

A solution would be to consider the portion of the skeleton at a constant distance from reference points associated to the iso-surface sample points; this method has already been applied to curve primitives ([1]), but is not however obvious to apply to skeleton constituted of both curves and surfaces.

The solution we have developed is based on the idea that only a predefined number of neighbors in the coarse blending graph should be allowed to blend their field contributions. When a skeleton is refined using curve and surface subdivision techniques, the blending properties are transmitted from parent to child primitives. In consequence, a primitive only blends with others that have either the same parent in the coarse skeleton graph, or whose parents are allowed to blend in this graph. To maintain surface smoothness using such a restricted blending graph, the definition of field values has to be modified. The idea is to ensure that only primitives with locally null contributions will be added (respectively suppressed) to (from) the group of active primitives when we traverse the implicit surface.

We implemented this idea in the case of a blending graph where each primitive only blends with its immediate neighbors in the coarse skeleton structure. Let P be a query point, S_i be the skeleton primitive in the coarse skeleton that contributes the most to $f(P)$ (i.e. P 's "parent skeleton"; if P is a mesh sample point constructed as described in Section 4), then P_i be the point of S_i which is the closest to P . To maintain surface smoothness, the contribution of S_i 's neighbors in the coarse graph has to vanish smoothly when P_i moves to the opposite extremity of S_i (since the neighbor will not be considered anymore if P 's projection crosses S_i 's border and moves to another skeleton element).

To do this, the contribution $f_j(P)$ of a neighboring primitive S_j is multiplied by a function $\alpha(u)$, where $u \in [0, 1]$ is the parameter of P_i on S_i if S_i is a segment skeleton and the parameter of P_i on the segment that goes from S_j to the other extremity of S_i if the latter is a triangle skeleton. See Figure 12. The function α is set to vanish smoothly when P_i is far from S_j , which ensures the same level of continuity for the resulting implicit surface.

This new solution has several benefits:

- first of all, it allows more locality in shape description, since each skeleton primitive only affects a restricted, predefined portion of the surface. This solves the case when neighboring branches should not blend together,




Figure 12: To maintain the implicit surface smoothness while using a restricted blending graph, field contributions from neighboring primitives are multiplied by a coefficient that smoothly vanishes when we move away from these primitive. (a) The point’s parent primitive S_i is a segment. (b) S_i is a triangle, and the neighbors are attached to a vertex of S_i . (c) S_i is a triangle, and the neighbors are attached to an edge of S_i . The neighboring primitives contribution is full at v_1 , null at v_0 , and interpolated between.

such as in Figure 13 (c) and (d). It also provides a good shape consistency if the skeleton is animated. Note that the object’s thickness may be smaller with our new method due to the influence of α (see Figure 15).

- secondly, it increases efficiency, since the portions of the meshes to be recomputed after a skeleton edition are known in advance, and since each field query at a mesh point only requires the computation of a restricted number of field contributions. Table 14 compares the efficiency of the two blending methods for the two surfaces depicted on Figure 15.

However, the method still has drawbacks: it is sensitive to variations in size and shape of the coarse skeleton’s elements, as illustrated in Figure 16. These parameters influence the extent of space of the blending regions, thus affecting the resulting surface radius. Also, such a defined field is C^1 along the skeleton. In order to have a C^1 iso-surface sampling, the directions of the sampling axes of the piecewise triangulation need also to be continuous. We are thus still looking for a better solution.




Figure 13: Controlled blending: comparison between the solution in [6] (b) and our new solution based on nearest neighbor blending with a refined skeleton (c).

objects	number of triangles	number of segments	blending of [6]	new blending
bretzel	0	24	22.77 s	4.80 s
alien	15	102	504.9 s	33.77 s

Figure 14: Computing times on a 180MHz O2 station.

6. RESULTS

The techniques we have presented have been implemented in an interactive modeling system, which interfaces is shown in Figure 17.

Modeling is done as follows: the user inserts skeleton vertices, then edges and/or triangles between the vertices. The implicit surface is automatically generated around it. The shape can be deformed by translating the vertices and/or increasing (resp. decreasing) the thickness of the surface at each vertex. A scene status bar allows to control the parameters of the scene or of the selected object (iso-value, fields functions, iso-surface sampling precision, unwanted blending type).

Characters have been modeled with this system: in Figure 18, the refinement of Nessie’s neck skeleton illustrates the subdivision-curve LOD on the field. The LOD provided by the subdivision-surface primitive is illustrated in Figure 19 by the refinement of the bat’s wings.

Both figures illustrates our unwanted blending solution, as the two fins of Figure 18, or ear and wing of Figure 19 do not blend with each other as they would with previous method. In practice, modeling is performed with the coarse skeleton, or after one subdivision, and iso-surface sampling error is 0.33 for iso-value 1.0.

7. CONCLUSION

In this paper, we have extended the subdivision-curve implicit primitive representation with the subdivision-surface implicit primitive, allowing to model shapes such as in Figures 18, 19 and 7. Our second contributions which is a solution to unwanted blending allows to constrain the surface to the skeleton’s topology. However, the unwanted blending method still has a drawback: it is sensitive to variations in size of the coarse skeleton’s elements. Also, it is dependent of the iso-surface triangulation method, as sample points need an associated reference point on the skeleton.

These contributions have been illustrated by their integration in an interactive modeling system, where the user models complex shapes by progressively refining and deforming their skeleton. Extending these techniques to multi-resolution subdivision schemes, where details would be in-




Figure 19: Bat character, with three levels of skeleton subdivision




Figure 15: Blending using previous method [6] (a, c) and our method (b, d).

incorporated to the skeleton structure during subdivisions, would be a very useful extension to this approach.

Finally, the approximation and interpolating schemes respectively slims and fattens the implicit surface after each subdivision, as neither of these schemes does take into account the implicit volume, or, more generally, desirable characteristics of the implicit shape (implicit surface area, implicit volume, size of the shape) that are to be kept while subdividing. A subdivision scheme that would take into account these properties would be of interest.




Figure 16: (a) Two skeleton elements. (b) The insertion of a third small segment produces a smooth crease, since blending is not allowed anymore between the two largest segments.

8. REFERENCES

- [1] Alexis Angelidis, Pauline Jepp, and Marie-Paule Cani. Implicit modeling with skeleton curves: Controlled blending in contact situations. In *Shape Modeling International*. ACM, IEEE Computer Society Press, 2002. Banff, Alberta, Canada.
- [2] Alan H. Barr. Superquadrics and angle-preserving transforms. *IEEE Computer Graphics and Applications*, 1(1):11–23, 1981.
- [3] J. Blinn. A generalization of algebraic surface drawing. *ACM Transactions on Graphics*, 1(3):235–256, July 1982.
- [4] Jules Bloomenthal, Chandrajit Bajaj, Jim Blinn, Marie-Paule Cani-Gascuel, Alyn Rockwood, Brian Wyvill, and Geoff Wyvill. *Introduction to Implicit Surfaces*. Morgan Kaufmann, jul 1997.
- [5] Jules Bloomenthal and Ken Shoemake. Convolution surfaces. *Computer Graphics*, 25(4):251–256, jul 1991. Proceedings of SIGGRAPH'91 (Las Vegas, Nevada, July 1991).

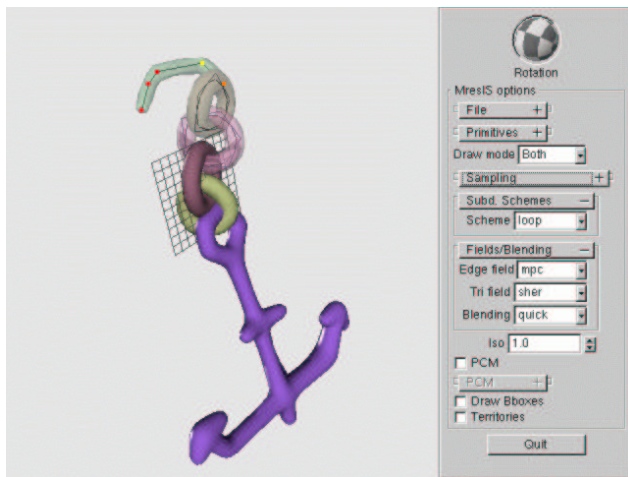


Figure 17: Modeling tool.

- [6] Marie-Paule Cani and Samuel Hornus. Subdivision curve primitives: a new solution for interactive implicit modeling. In *Shape Modelling International*, Italy, May 2001.
- [7] Marie-Paule Cani-Gascuel and Mathieu Desbrun. Animation of deformable models using implicit surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 3(1):39–50, mar 1997.
- [8] D. Doo and M. Sabin. Analysis of the behaviour of recursive division surfaces near extraordinary points, 1978.
- [9] N. Dyn, D. Levin, and J. Gregory. A four-point interpolatory subdivision scheme for curve design, 1987.
- [10] Marie-Paule Cani Eric Ferley and Jean-Dominique Gascuel. Practical volumetric sculpting. In *Implicit Surfaces'99—Eurographics and ACM-Siggraph Workshop*, Bordeaux, France, sep 1999.
- [11] Tinsley A. Galyean and John F. Hughes. Sculpting: An interactive volumetric modeling technique. *Computer Graphics*, 25(4):267–274, 1991.
- [12] C. Loop. Smooth subdivision surfaces based on triangles, 1987.
- [13] J. McCormack and A. Sherstyuk. Creating and rendering convolution surfaces, 1998.
- [14] Agata Opalach and Steve Maddock. Implicit surfaces: Appearance, blending and consistency. In *Fourth Eurographics Workshop on Animation and Simulation*, pages 233–245, Barcelona, Spain, sep 1993.
- [15] V. V. Savchenko, A. A. Pasko, O. G. Okunev, and Tosiyasu L. Kunii. Function representation of solids reconstructed from scattered surface points and contours. *Computer Graphics Forum*, 14(4):181–188, 1995.
- [16] Andrei Sherstyuk. Kernel functions in convolution surfaces: a comparative analysis. *The Visual Computer*, 15(4), 1999.
- [17] E. Stollnitz, T. Deroose, and D. Salesin. *Wavelets for Computer Graphics*. Morgan Kaufmann, San Francisco, California, 1996.
- [18] Brian Wyvill and Geoff Wyvill. Field functions for implicit surfaces. *The Visual Computer*, 5:75–82,




Figure 18: Nessie character, with three levels of skeleton subdivision

December 1989.