Chapter 8: Conclusions

The benefits of solid modeling and, by extension, of geometric modeling are usually considered in terms of the simplicity of evaluation [Requicha and Voelcker 1982], the compactness of representation [Mäntylä 1988], the ease of blending [Rockwood 1989], [Woodwark 1986], [Blinn 1982], and the interactivity of design. In this dissertation we have developed design techniques that, we believe, satisfy in some measure each of these criteria.

We have based our design techniques on the use of the skeleton, or inner structure, of an object. In considering branched objects, we concluded that implicit surfaces were simpler to implement and worked well with skeletons. We emphasized the concrete representation afforded by polygonization of the implicit surface.

The biological leaf caused considerable concern by its requirement for a volumetric and a surface representation. We have unified these into a single, non-manifold representation, which, in its abstract form, extends the conventional implicit surface treatment of regions of space. It permits, for example, a simple expression and evaluation for parametric surfaces combined with and trimmed against implicit volumes. The power of this extension and the degree to which a designer can easily modify complex shapes remain to be determined.

Although non-manifold polygonization requires a substantial implementation, several benefits result from the concrete representation produced. Unlike conventional implicit surfaces, non-manifold surfaces may have borders and intersections. Those parts of surfaces contained within implicitly defined volumes are discarded. Volume/surface intersections, although approximated, are represented uniquely, which can avoid visualization artifacts. The implementation has its limitations, however, and future work might accommodate arbitrarily complex intersections of surface and polygonizing cell, improved boundary accuracy, and a treatment of parametric surfaces more efficient than the distance surface presented in chapters 4 and 7.

The use of convolution surfaces within a design environment enables the creation of complex, well-behaved shapes when given the position, orientation, and scale of individual segment and polygonal skeletal elements. Although the designer loses detailed control over blending, he or she gains freedom in combining one and two-dimensional skeletal elements to produce smooth surfaces. The results in chapter 6 suggest that polygonal elements can produce seamless, bulge-free shapes.

In our discussions of convolution surfaces we emphasized those integration properties that have a beneficial geometrical impact. In particular, our choice of ½ as an iso-contour value requires the convolution surface to interpolate the edges of an isolated polygonal element. Because it offers a near circular cross-section, we find the Wyvill kernel particularly advantageous for polygonal skeletal elements. It has a closed form integral, finite support, and is cubic in d^2, avoiding a square root to compute distance. Unfortunately it is not separable, but we have seen it treated as such with reasonable results. The application of convolution to other elements, such as three-dimensional volumes or bundles of strands, merits additional investigation.

We have argued that natural forms are elegant, interesting, and inspiring. They also provide strenuous challenges to our methods of representation and visualization, and have prompted us to study a range of geometric methods. Along the way, we mentioned several ancillary techniques; spatial partitioning, for example, is of use in rendering implicit surfaces by ray-tracing, contour-lines, or polygonization, and in calculating distance to a polygonal mesh and in culling uninfluential implicit primitives.

Although we provided examples of simple articulation and examples of complex forms, we did not provide examples of shape transformation or metamorphosis.

We did not animate any of our complex models, although this would have been a useful means to determine the adequacy of the models. Many of our examples were incomplete. The hand did not have creases or fingernails, nor did it prevent unwanted blending. The leaf was not rendered with texture.

All of the forms we explored were represented by skeletons. We hope the variety of these shapes is evidence that the skeleton is a useful mechanism for the design of three dimensional forms. Without a skeleton, these and other shapes are difficult to animate, procedurally define, developmentally generate, or replicate with alterations. A skeleton by itself allows articulation, but provides insufficient surface detail. Covering a skeleton provides both articulation and realism. We hope the comparative ease in working with skeletons bodes well for their future use within implicit surface modeling systems. Although this dissertation is for use in computer science and has examined several mathematical relations, it has sought to argue on behalf of the designer.

It has only been in the past few years that implicit surface modeling has gained significant acceptance, and it remains to be seen what system features will, in the future, be regarded as conventional. As computer graphics becomes increasingly accessible to a wider set of users and designers, we should expect an increasing demand for ever more varied geometric models.

In this dissertation we did not fully observe the common exhortation to focus narrowly and examine deeply. At times, we have taken a broad view of design methods. This view is necessarily speculative until a design system is fully implemented and enhanced during a period of use. Our results suggest that the design methodology described herein is of value. In this dissertation we have introduced and promoted this design methodology, and we hope it will find its way into mainstream geometric modeling. With exposure to more users and access to greater computational power, we will learn the significance of these techniques.