RAPID APPROXIMATE SILHOUETTE RENDERING OF IMPLICIT SURFACES

David J. Bremer and John F. Hughes

Department of Computer Science
Box 1910
Brown University
Providence, RI 02912
{djb,jfh}@cs.brown.edu

ABSTRACT

We describe a method for rapidly producing a non-photorealistic rendering of an implicit surface. The rendering includes silhouettes and some shading near silhouettes to help indicate curvature. The method identifies silhouettes probabilistically, but we include strategies to make it likely that we find all silhouette curves, especially in multiple-frame sequences. The method is approximate, in the sense that the silhouette curves may suffer some position error; the degree of approximation is determined in part by an adjustable parameter to a curve-tracing method, allowing a tradeoff of accuracy against speed.

1. SILHOUETTE RENDERING

Inspired by non-photorealistic rendering (NPR) techniques [9, 10, 3, 14, 11, 13, 7], especially interactive NPR for rapid display of complex polyhedral objects [7], we applied an alternative style of rendering to implicit surfaces.

Our basic approach concentrates computational resources on drawing the surface's silhouette quickly and accurately from a given viewpoint, with hidden lines removed. Additionally, a little curvature information is displayed near the silhouette and on the interior of the rendering. Example results of the algorithm, showing off the available features, are shown in Figure 1.

While standard rendering techniques require that at least the entire surface, if not a whole 3D volume, be traversed to make a rendering, our silhouette edge-drawing approach needs only to trace the silhouette curve or curves, allowing the rendering to take place much faster, and/or to be done more carefully than, say, a rough polygonization. With this algorithm, good rendering can often occur at interactive or near-interactive rates.

This new approach makes good use of computational resources. Although we draw only a few lines on the screen to represent a whole surface, we are
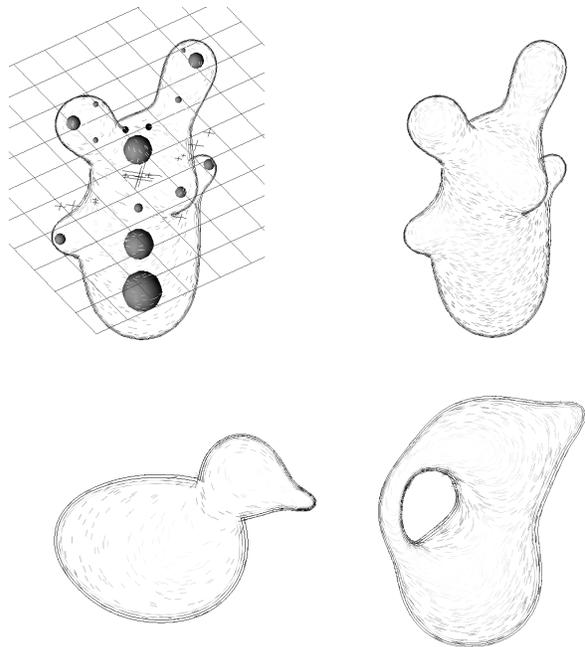

Figure 1: Several models rendered with our non-photorealistic renderer. (a) A bunny, in the modeler. (b) The bunny rendered without the modeling aids. (c) A duck. (d) A "vase."
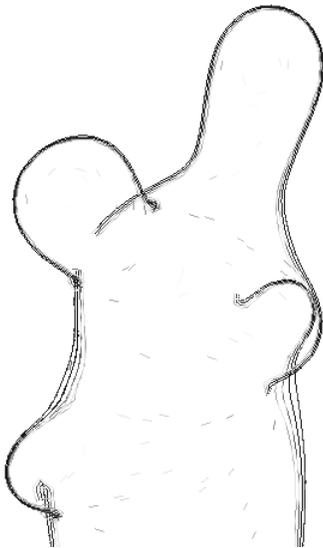
Figure 2: The effects of failing to removing occluded silhouettes on our bunny model.

giving the eye the visual cues which seem most important to it in perceiving shape. These cues are the position and size of each part of the object relative to others, as well as a bit of curvature information near the silhouettes and across the whole surface. The silhouette lines provide a way to distinguish the surface from the background and to show a little of the surface's shape, with relatively little computational effort. In addition, occluded portions of silhouette edges are culled, because occlusion is among the most important cues our visual system gets regarding the position of parts of the surface relative to one another. It seems that everyone realizes from an early age that if one part of an object is in front of another, the closer part occludes the farther [6, 5], so we felt it absolutely necessary to cull out occluded parts of the silhouette. See Figure 2. In addition, we show a little information about the surface curvature near the silhouette and on the interior of the rendering. It does give a little shape from shading, which is filled in slowly instead of immediately.

For small models, our algorithm's rendering speed is comparable to that of a coarse polygonizer or point renderer, such as the one shown by Stander and Hart [12]. For large models the algorithm should prove to be asymptotically faster, although no longer real-time. For similar reasons, our algorithm would also be asymptotically faster than a volume renderer, which must traverse a whole 3D volume, or at least the whole surface, before producing an image. Because of the asymptotic difference in rendering time, future research might focus on better methods of fairly fast non-interactive rendering of complex implicit surfaces.

2. PREVIOUS WORK

As mentioned to earlier, our work was primarily inspired by recent work in non-photorealistic rendering (NPR). NPR in general describes any method of drawing a simple abstraction of a complex model in order to highlight important details, or to add extra information or feeling [13]. There are at least two big benefits to rendering this way. The first is the perceptual gain–being able to get more information or feeling from an image or series of images. The other is related to efficiency–we can sometimes save computational expense by computing and rendering only a select portion of a model which is felt to be the most important. Since implicit surfaces are expensive to render at all, we focused on using NPR to make the algorithm fast. Future work in the other area might include shading the whole surface in an expressive way, such as with a pen-and-ink style, and providing a visual way to identify cusps, singularities, or other interesting parts of the model.

Our work was directly inspired by Markosian et al. [7] who produce real-time silhouette-renderings of a polygonal model. After seeing that work, we tried to adapt the idea of rendering only outlines to the problem of implicit surface rendering.

But that work was not the first to produce line drawings of surfaces. An early line drawing application was developed by Scott Roth [8]. It draws just object outlines, but does so with a ray-tracing algorithm. Similarly, it was developed to give a speed improvement over regular ray-tracing, which it obtains because it is not required to evaluate lighting equations. Unlike the other work, it was designed to operate on CSG models constructed from a few implicitly defined primitive shapes (cubes, spheres, cones, and cylinders).

Dobkin et al. [4] have developed an algorithm that will trace implicit sets of functions from $\mathcal{R}^n$ to $\mathcal{R}^k$; if we take our implicit function f from $\mathcal{R}^3$ to $\mathcal{R}$, its gradient ∇f, and the view direction $\mathbf{v}$, we can create a new function

$$g : \mathcal{R}^3 \to \mathcal{R}^2 : \mathbf{x} \mapsto (f(\mathbf{x}), \nabla f(\mathbf{x}) \cdot \mathbf{v});$$

The inverse image of $(0, 0)$ under g is exactly the silhouette set that we compute here. The primary differences in our approaches are that (a) we do occlusion testing and some shading of the surface, which is not possible with the results of the Dobkin et al. algorithm, and (b) we work with the function f, and assume that both it and its first two derivatives are available, whereas Dobkin et al. work with a fixed-resolution piecewise-linear approximation. In the language of that paper, we taken an "infinitesimal" approach and they take a "local" approach. Nonetheless, their ideas could easily be adapted for the "tracing" portion of this work; we could then use our methods for the occlusion testing and the shading portions.

As a side note, there is a rendering style called contour line drawing which bears a superficial resemblance to ours. This method slices planes through the implicit surface and draws the curves formed by the intersection [2]. The big difference between that method and ours is that it uses curves across the whole surface, in order to show the surface shape, whereas ours draws far fewer curves and does so just to indicate the surface's outline.

3. THE ALGORITHM

3.1. Overview

We need algorithms to find the silhouette edges and information for shading, to test the edges for occlusion, and to draw the edges and shading information. However, if while examining the model the user pauses for a moment, the algorithm need not waste time recomputing silhouettes.

Figure 3 shows the process schematically. The pseudocode is as follows:

```
For each frame
    If the camera moved or model changed
        Find new silh. edges and shading info
        Test edge sections for occlusion
        Draw the edges and shading strokes
    Else
        Look for missed edges and shading info
        Draw all edges and shading strokes
```

3.2. Notation and assumptions

We assume that the surface model, S, that we are rendering is the zero-set of a twice continuously differentiable function f on R^3. To make the explanation simpler, suppose that the solid bounded by the surface is the region in which $f < 0$, which causes gradients to point away from the surface rather than inward.

In order to make the algorithm work with a wide variety of model representations, it treats the implicit function defining the model as a "black box" from which it only needs to be able to get, at any point, the function value, the gradient, and the Hessian (matrix of second derivatives) of the function.

We further assume that the virtual camera is orthographic with view direction $\mathbf{v}$, that all points on the film plane are outside the object (i.e., $f > 0$ on the film plane), and that the surface S lies entirely within some known region of space (a sphere of radius 20 about the origin in our particular implementation).

We denote points and vectors by boldface letters; the point $\mathbf{x}$ has coordinates x_1, x_2 and x_3.
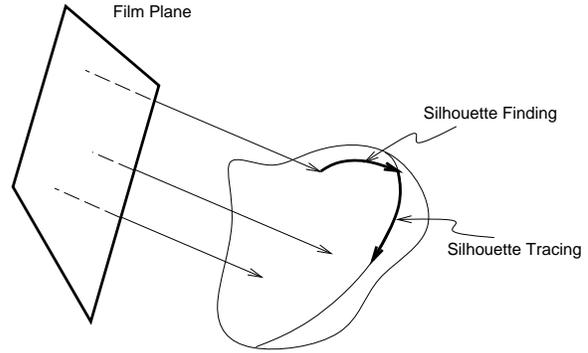

Figure 3: Rays from the film plane are traced along the view direction until they hit the surface. Then we trace in the direction of the view-plane projection of the gradient to try to find a silhouette. Once one is found, we trace along it.

We also assume, to make ray-surface intersection easier, that there is a constant $K > 0$ such that at every point $\mathbf{x}$, the gradient of f,

$$\nabla f(\mathbf{x}) = \left(\frac{\partial f}{\partial x_1}(\mathbf{x}), \frac{\partial f}{\partial x_2}(\mathbf{x}), \frac{\partial f}{\partial x_3}(\mathbf{x})\right)$$

has magnitude bounded by K. Finally, we denote the *Hessian* of f, the matrix of second partial derivatives, by Hf, so that

$$Hf(\mathbf{x}) = \left[\begin{array}{ccc} \frac{\partial^2 f}{\partial x_1 \partial x_1}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_1 \partial x_2}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_1 \partial x_3}(\mathbf{x}) \\ \frac{\partial^2 f}{\partial x_2 \partial x_1}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_2 \partial x_2}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_2 \partial x_3}(\mathbf{x}) \\ \frac{\partial^2 f}{\partial x_3 \partial x_1}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_3 \partial x_2}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_3 \partial x_3}(\mathbf{x}) \end{array} \right].$$

3.3. Silhouette Finding

Silhouette finding is a three-step process:

1. Locate a point on the surface through ray-surface intersection

2. Trace along the surface to a point on a silhouette

3. Trace out the silhouette

3.3.1. Ray-surface intersection

We apply a modification of Kalra and Barr's [2] implicitization algorithm to do ray-surface intersection: the idea is that for functions with bounded gradients, we can search for ray-surface intersections by stepping along a ray and be guaranteed to miss no intersection: if, while searching for an intersection along the ray $\mathbf{p} + t\mathbf{v}$, we are at location $\mathbf{x}$, then we can take a step of size $f(\mathbf{x})/K$ to location $\mathbf{x}' = \mathbf{x} + (f(\mathbf{x})/K)\mathbf{v}$ and be confident that $f(\mathbf{x}') \geq 0$. We search along rays from the eye until the function value is nearly zero, and call the resulting point a surface point. If the search proceeds far enough, our assumption that the surface
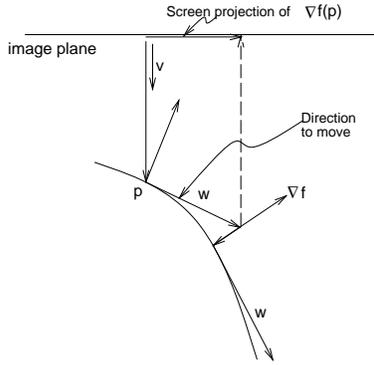
Figure 4: When we reach a point **p** of the surface, we find a tangent vector whose screen projection is in the same direction as that of the gradient (that tangent vector's called w here) and move in that direction to find a silhouette.
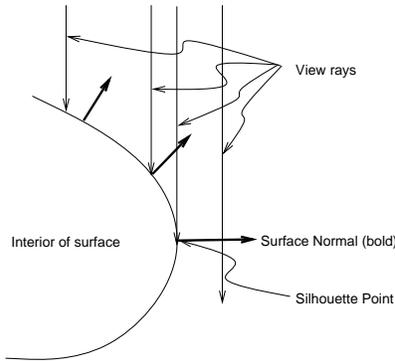

Figure 5: When a view ray hits the interior of the surface the surface normal (shown in bold) points back towards the eye. When the view ray grazes the silhouette of the surface, the view ray and the surface normal are orthogonal.

lies within a bounded region of space lets us terminate the search. In this case, the silhouette-finding process does not continue with silhouette-point finding, but rather with ray-surface intersection using a different ray.

3.3.2. Silhouette point finding

When we find a ray-surface intersection, we try to use it to locate a silhouette by walking along the surface in the direction of the screen projection of the gradient at each point (see Figure 4).

We take the ray-surface intersection **p**, compute the function gradient at **p**, and use this to find a tangent vector in the plane spanned by **v** and ∇f; i.e., we let:[1]

$$F(\mathbf{p}) = \frac{\nabla f(\mathbf{p}) \times (\mathbf{v} \times \nabla f(\mathbf{p}))}{\|\nabla f(\mathbf{p})\|^2}$$

---

[1] The division by the square of the gradient is designed to make the vector field independent of the scale of f: if we replace f with αf, the vector field is unchanged.
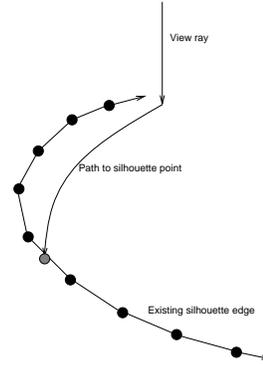

Figure 6: A new silhouette point, shown in grey, should be tested to see if it lies on a silhouette edge that was already traced out, to avoid recomputing the position of a silhouette edge.

The vector field F is tangent to the surface and lies in the plane spanned by **v** and $\nabla f(\mathbf{x})$. We trace along this vector field until the dot product of **v** and ∇f changes sign, which indicates that we have passed a silhouette. A silhouette point is a point of the surface where the tangent plane contains the view direction; it may be obscured by some other part of the surface, but we call it a silhouette point nonetheless. In the mathematical literature, it's sometimes called a "fold point." See Figure 5.

So far, we have just described the vector field to be integrated, not the method of integration. See Section 4 for a description of the integration method.

Often, after a few silhouettes have been traced out, the newly found point will lie on one of the silhouettes already traced out. So before proceeding, we test to see if this is just such a point. See Figure 6.

Silhouettes are represented as 3D polylines. We do a proximity test between the found point and all the silhouette points which may be near it. (These points are stored in a hash table based on location in three-space.) All the points in the polyline are approximately ϵ units apart, so if the found point is within ϵ units of any found point, it is discarded, and the silhouette finding process starts over with a new ray-intersection.

3.3.3. Silhouette tracing

A curve $h : \mathcal{R} \to \mathcal{R}^3$ lies on the silhouette of the surface S viewed along **v** if

- $h(t) \in S$ for every t, and

- The tangent plane to the level surface at $h(t)$ contains **v** for every t.

These two conditions can be rephrased as

$$
\begin{aligned}
f(h(t)) &= 0 \\
\mathbf{v}^t \nabla f(h(t)) &= 0.
\end{aligned}
$$

Rather than try to solve for $h(t)$ analytically, we instead use this implicit description to determine the tangent vector of h, from which we can determine h by numerical integration.

Differentiating each equation with respect to t, applying the chain rule, and using $\mathbf{w}$ to denote $h'(t)$, we get

$$\begin{aligned}
\nabla f(h(t)) \cdot h'(t) &= 0 \\
\mathbf{v}^t H f(h(t)) h'(t) &= 0,
\end{aligned}$$

i.e.,

$$\begin{aligned}
\nabla f(h(t)) \mathbf{w} &= 0 \\
\mathbf{v}^t H f(h(t)) \mathbf{w} &= 0.
\end{aligned}$$

Thus the tangent vector to a silhouette curve must be orthogonal both to the gradient at its basepoint, and to the product of the Hessian at the basepoint with the view direction. This makes it proportional to the cross product of these:

$$\mathbf{w} \propto \nabla f(h(t)) \times \mathbf{v}^t H f(h(t)).$$

We can therefore trace along a silhouette by computing

$$\mathbf{w} = \nabla f(\mathbf{p}) \times \mathbf{v}^t H f(\mathbf{p})$$

and finding an integral curve for this vector field. Of course, at locations where $\mathbf{w} = 0$ the tracing process stagnates. This happens, for example, at cusps like the one shown in Figure 7. Our silhouette-tracing algorithm starts from a silhouette point, found previously, and traces out the silhouette by taking a series of steps of size ϵ until the tracing process stagnates or returns to the starting point; if it stagnates, we return to the starting point and trace in the other direction. Section 4 discusses the details of the tracing scheme.

In order to avoid stepping past the starting point, after each step we find the distance between the new point and the starting point. If the distance is less than ϵ, tracing stops and the 3D polyline is closed.

3.4. Occlusion testing

We test the vertices of the silhouette for occlusion by first checking every nth (4th in our implementation) vertex for occlusion, and then, for those between which occlusion status changes, testing the intervening vertices as well. To check occlusion of a single vertex of a silhouette polyline, we start at the vertex, move back towards the film plane from it, and do a ray-surface intersection test back into the scene. If the ray intersects the surface at a place much closer to the film plane than our vertex, we declare the vertex invisible; otherwise it's visible.

Note that because of numerical issues, the ray may not intersect the surface exactly at the silhouette vertex (which may, indeed, not actually lie exactly on

Figure 7: A cusp occurs at the end of a silhouette. A slight "hook" is conventionally drawn at the cusp to indicate its shape.

the surface), so the "much closer" test is important. Unfortunately, if the true silhouette is just barely obscured by some nearer piece of surface, and the silhouette vertex still happens to seem to be visible, we draw still the silhouette.

3.5. Rendering

3.5.1. Silhouette edges

The silhouette edges could be simply drawn as polylines, as would be done in a basic implementation. But, to convey extra information about the surface's shape near the silhouette, we alter the drawing style based on the local curvature. At a point $\mathbf{p}$ of the silhouette, the Hessian can be used to determine the curvature of the surface in the plane defined by the view and the normal to the silhouette. To be more precise, if we consider the plane through $\mathbf{p}$ spanned by the gradient and the view direction, its intersection with S is a curve. The gradient to S is normal to this curve, and the rate of change of this normal in the view direction is proportional to the curvature of the curve. But the rate of change of the normal as we move the basepoint in some direction $\mathbf{u}$ through the point $\mathbf{p}$ is simply $\mathbf{u}^t H f(\mathbf{p})$; to compute its component in the view direction, we take the inner product with the view vector. We can therefore compute the curvature in the view direction as $\mathbf{v}^t H f(\mathbf{p}) \mathbf{v}$. Unfortunately, this computation depends on the scaled f; we normalize it by dividing by the magnitude of the gradient of f. Hence we compute

$$\kappa(\mathbf{p}) = \frac{\mathbf{v}^t H f(\mathbf{p}) \mathbf{v}}{\|\nabla f(\mathbf{p})\|}$$

for points on the silhouette, and use it to help us draw shading near the silhouettes to indicate curvature.

So we draw not just the silhouette but several parallel copies of it, with the inter-copy spacing proportional to $1/\kappa(\mathbf{p})$. Thus tightly-curved sections get closely-spaced curves, and areas of shallow curvature get widely-spaced ones. See Figure 8.
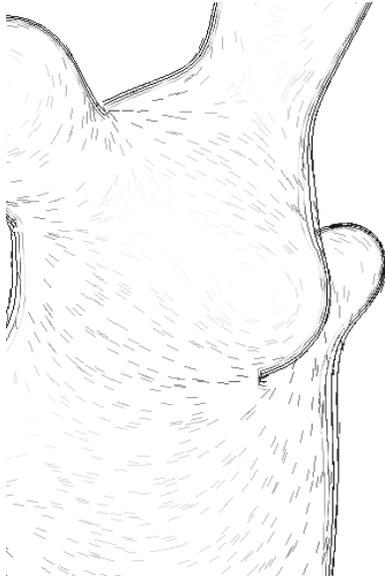
Figure 8: The silhouette drawing style varies with the surface curvature. The tip of the bunny's left arm is drawn with tightly spaced lines, whereas the lines defining its torso are widely spaced.

We could also experiment with stroke styles as did Markosian et al. [7].

3.5.2. Interior shading

Interior shading strokes are drawn using a very simple lighting model – we assume that the light in the scene is arriving from behind the virtual camera and that the surface is diffuse, so that the illumination is proportional to the dot-product of the view direction and the (unit) surface normal. At interior points, when a ray strikes the surface, we immediately compute the gradient at the intersection point so that we can start searching for a silhouette. We use this computed gradient to determine two additional things:

- the direction $\mathbf{v} \times \nabla f$ that is tangent to the curve of constant illumination (isophote), and

- the lightness s of the surface ($\frac{\mathbf{v} \cdot \nabla f}{\|\nabla f\|}$) at the intersection point.

We then pick a color $|s|bg + (1 - |s|)db$, where bg is the background color (a neutral gray) and db is a dark blue, and draw a short stroke in this color, tangent to the isophote.

These "free" shading lines accumulate as rays are shot at the surface in search of silhouette edges, and help convey the interior shape of the surface. See Figure 9.

Figure 9: Several images of the same model, with progressively more strokes filled in. Note the light strokes on places perpendicular to the view direction, such as the bunny's nose, and dark strokes near silhouette edges.

4. EFFICIENCY CONSIDERATIONS

4.1. Approximate tracing

In both silhouette finding and silhouette tracing, we need to "walk along" the implicit surface, guided by a vector field. In each case the general algorithm we use is Euler integration: $\mathbf{p}$ is replaced by $\mathbf{p} + \epsilon F(p)$, where F is the vector field and ϵ is some small number. This approach is notoriously unstable; using it to walk along the circumference of a circle (i.e., along the vector field $F(x, y) = (-y, x)$) leads to the sort of spiral shown in Figure 10(a). But if in addition to knowing that we want to be guided by a vector field, we have some other constraint, we can use this to help stabilize the process. For example, in silhouette finding, we know that we not only want to move along the surface in the direction determined by the gradient and view direction but also want to remain on the surface. By adding a "penalty" term to the vector field – a term that's zero on the surface, but drives us back to the surface when we're off it – we can ensure that the integral curves don't wander too far. In the case of the tangent vector field to the circle, we can use the vector field $G(x, y) = K * (x^2 + y^2 - 1) * (-x, -y)$ as a "corrector" field; when we add this to F, the integral curves, even with Euler integration, now lie close to the circle rather than following a diverging spiral (see Figure 10(b)). This idea is closely related to the constraint-satisfaction method in Barzel and Barr's "Dynamic Constraints" work [1]. The constant K determines the degree of penalty for falling off the circle: if K is small, the curve will not stay close; if K is made too large, however, the curve can oscillate across the circle.

In the case of silhouette finding, we know that we want to remain on the surface as we search for a silhouette. Our first implementation took small steps and then, at the end of each step, did a ray-surface intersection to "fall back" onto the surface. Our revised version instead uses the stabilization method: instead of using the vector field

$$F(\mathbf{x}) = \frac{\nabla f(\mathbf{x}) \times (\mathbf{v} \times \nabla f(\mathbf{x}))}{\|\nabla f(\mathbf{x})\|^2}$$

defined on the surface, we define (on all of $\mathcal{R}^3$) the vector field

$$F(\mathbf{x}) = \frac{\nabla f(\mathbf{x}) \times (\mathbf{v} \times \nabla f(\mathbf{x})) - f(\mathbf{x})\nabla f(\mathbf{x})}{\|\nabla f(\mathbf{x})\|^2}.$$

The additional term $- \frac{f(\mathbf{x})\nabla f(\mathbf{x})}{\|\nabla f(\mathbf{x})\|^2}$ – is a field that points towards the surface at all points of space. Hence when Euler integration takes our curve off the surface, the additional term tries to coax it back onto the surface. Just as in the case of the spiraling circle, the correction is imperfect: the "stabilized" path still does not lie exactly on the surface. But it does not diverge from it either, and the expensive ray-surface
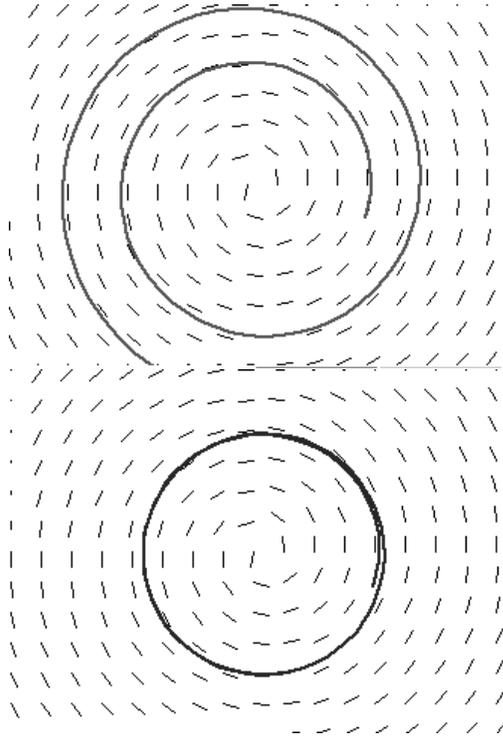

Figure 10: (a) When the tangent field to a family of circles is integrated with Euler integration, the result is a growing spiral. (b) When we add a penalty term for distance from the starting circle, the result is a non-diverging circle (albeit slightly displaced from the starting point's circle).

intersection is eliminated. By the way, this is just a specialized type of predictor-corrector integrator; the novelty is in its application to finding silhouettes for isosurfaces.

For the case of silhouette *tracking*, our initial corrector takes the predicted location and does a ray-surface intersection (moving in the negative gradient direction) to fall back to the surface, and then a silhouette-finding operation to fall back to the silhouette. In the current implementation, these two steps use the same algorithms that initially are used to find the surface and then a silhouette point. With this implementation, on reaching a cusp the corrector no longer works correctly, which is why tracing stops at cusps.

We have since implemented a corrector like the one described for silhouette finding. The pictures here, however, use the original method, since we have not thoroughly tested the new corrector. For this new corrector, we have two additional constraints: we want to find integral curves of the vector field

$$G(\mathbf{p}) = \nabla f(\mathbf{p}) \times \mathbf{v}^t H f(\mathbf{p})$$

on the surface, but Euler integration will wander off the surface and off the silhouette curve. Again, we can add a correction of the form $\frac{-f(\mathbf{x})\nabla f(\mathbf{x})}{\|\nabla f(\mathbf{x})\|^2}$ to keep the curve on the surface. We can also add a correction to keep the curve running along the silhouette: just

as adding $-f\nabla f$ tends to drive f to zero, we can add $-g\nabla g$, where $g(\mathbf{p}) = \mathbf{v} \cdot \nabla f(\mathbf{p})$ to drive g to zero, i.e., to drive us onto a silhouette. This expression simplifies to

$$-(\mathbf{v} \cdot \nabla f(\mathbf{p}))\mathbf{v}^t H f(\mathbf{p}).$$

Fortunately, ∇f and $H f$ are already computed in getting the basic vector field to walk along.

In summary, we find an integral curve of

$$\frac{1}{\|\nabla f(\mathbf{p})\|^2}(\nabla f(\mathbf{p}) \quad \times \quad (\mathbf{v} \times \nabla f(\mathbf{p})) - f(\mathbf{p})\nabla f(\mathbf{p})$$
$$-K * (\mathbf{v} \cdot \nabla f(\mathbf{p}))\mathbf{v}^t H f(\mathbf{p}))$$

and it will not only follow the silhouette, but if it (because of Euler steps) wanders from the silhouette, will be driven back towards it.

In our initial tests, setting K to 1 has led to some oscillatory behavior; $K = 0.5$ seems to work well, however.

For this new tracer, the behavior at cusps is more complex than before: the silhouette curve goes from being visible to being invisible by briefly heading directly away from the viewer. In general, this seems to have worked fine, and silhouette tracing now ends when the tracer returns to its starting point. But what if, near a cusp, the tracer overshoots? Figure 11 shows that the silhouette-tracing field, just beyond the cusp, has a clockwise-spiral projection onto the surface; when the corrector field is added, we will find that any overshoot into the region beyond the cusp will get swept back clockwise up to the silhouette edge, or will stall out on the line immediately below the cusp, where the (uncorrected) field is orthogonal to the surface.

As an alternative to Euler integration, we could use Runge-Kutta integration. Runge-Kutta integration requires more computation for each step, but diverges much more slowly from the ideal curve than Euler integration, allowing bigger steps to be taken. But both methods still diverge, so regardless of which we use, we would still want to take advantage of the special correction information available to us, which lets us pull the curve directly back toward the surface or curve on which it should lie. Future work might include testing to see if, and how much, the speed gained from Runge-Kutta's bigger steps offsets the cost of extra computation.

4.2. Choosing good rays to shoot

Our algorithm begins by shooting rays from the film plane along the view direction (orthogonal to the film plane in our implementation) into the scene, hoping to find silhouettes. The silhouette-finding algorithm can easily get stuck in "valleys" in the surface, so some rays produce nothing of interest. On the other hand, a ray that falls near a silhouette will rapidly
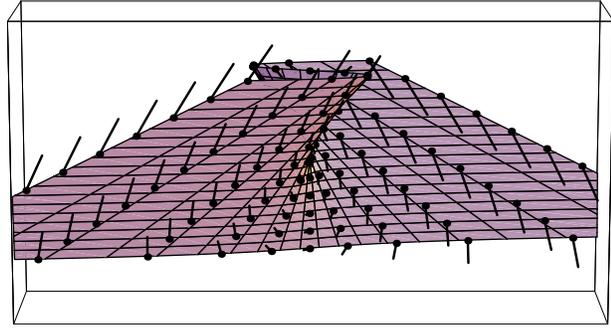

Figure 11: A prototypical cusp in the graph of $y = x^3 + xz$, as seen from $(5, 0, -1)$ looking at $(0, 0, 0)$; the (un-corrected) silhouette-tracing field is indicated by the dot-and-line icons: the dot is the basepoint of the vector, the line shows the direction. Along the silhouette edges, the field is evidently tangent to the silhouette. In the lower half of the figure, beyond the silhouette, the field has a sort of sprial form, so that tracing it from points below the cusp should lead to a clockwise circular path back to a location above the cusp (although such curves may well stagnate if they hit the line directly below the cusp, where the field is normal to the surface).

lead to productive results. Because we are trying to render at interactive speeds, we have some confidence that inter-frame differences in the image are small, so silhouettes in a frame are likely to be near their locations in the previous frame. Thus former silhouette points are good candidates for ray-starting-points in the current frame. If we displace these points slightly "inward" along the surface normal, then surface (or camera) translations are less likely to cause them to miss the surface when they're re-shot. We therefore, in choosing rays to shoot in each frame, preferentially select starting points that lie on silhouettes from previous frames; we also use some randomly chosen rays, in hopes of finding new silhouettes that may appear far from any previous silhouette.

5. TIMING

Initial timing tests for an execution that involves frequent camera motion suggest that the bulk (60%) of the algorithm's time is spent determining silhouette curves. Of this, half (30%) is spent shooting rays, many of which miss the surface (although this depends on the screen-area occupied by the surface), a quarter (15%) on silhouette-finding, and the remaining quarter on silhouette tracing.

Another 30% is spent on occlusion testing, virtually all of it in ray-surface intersection computations.

The remaining 10% is spent drawing the shapes, doing object-creation in Java3D, and handling thread synchronization and other tasks unrelated to the algorithm.

By contrast, during a model-creation session, about 80% of the time was spent determining silhouette curves (70% ray-shooting, 10% silhouette-finding, 20% silhouette-tracing), and about 20% doing occlusion testing. A small amount was spent creating a drawing shapes in Java3D.

6. LIMITATIONS AND FUTURE WORK

The algorithm described here has some serious limitations. We require that f, ∇f and Hf all be available at all points of the model that we render. For sampled data, these might be provided by performing some tricubic interpolation of the samples, although we have not implemented this. Further, our ray-surface intersection requires the bound on the gradient magnitude, although it could be replaced with some other method if no such bound is available.

The occlusion testing uses only samples of the silhouettes, and hence is prone to small errors. If we knew that we had computed all silhouettes, and projected them to 2D, we could apply the methods used by Markosian et al. [7] instead. It may well turn out that this is more efficient, because it would drastically reduce the number of ray-surface intersection tests we need to perform. Furthermore, it would allow us to do 2D region-fill operations to make the surface interior a different color from the background, which would presumably help in indicating the object's shape.

There are two additional cues to the shape that could probably be shown by using a perspective camera: motion parallax and the fact that objects diminish in size with increasing distance from the viewer. The first seems most effective in a system like ours which views models at interactive rates, although the second may be worth considering as well. The current implementation uses only an orthographic camera; replacing it with a perspective camera is a small change, but the vector $\mathbf{v}$, which is constant for an orthographic camera, becomes dependent on the viewed point for a perspective camera, which would add some modest computation.

Our system cannot render texture maps on the surfaces, and indeed, since we sample as few points on the surface as we can, we see no way to include this.

We would also like to push the limits of NPR further. For example, the rendering near cusps, where silhouette edges disappear, has a disappointing (to us) appearance, with the "shading curves" fanning out. Hand-drawn cusps like the one in Figure 7 present a far more attractive appearance, and we'd like somehow to copy this. It would be nice to be able to find and detect other interesting features, such as sharp edges and singularities. In addition, we might experiment with a slower version of the algorithm which would draw shading across the whole surface, perhaps in a pen-and-ink style.

The tradeoff between step size and speed is only partly successful: if we increase the step size too much, either we spend excessive time in the explicit correctors (re-intersect surface, re-find silhouette) or the implicit correctors can fail because the assumption that the point is not far from the surface, so that gradient forces can bring it back on, fails.

As mentioned earlier, it would be good to try other methods of integrating the silhouette curve, such as with Runge-Kutta integration, to see if we can make a gain in efficiency.

Our use of Java3D is unsatisfactory: it seems foolish to create curves in 3-space so that a 3D renderer can redraw them for us in 2D. But with the optimizations in Java3D, it appears (at least on our Sun workstations) to be faster to do this than to draw directly in 2D.

7. FINAL NOTES

The Java classes implementing this work will be made available through the homepage of the authors, at `www.cs.brown.edu/people/jfh/is/is.html`, which also contains instructions for using the application. The application does, however, use Java3D, requiring that users download this library from Sun and install it on their local machines.

8. ACKNOWLEDGMENTS

9. REFERENCES

[1] Ronen Barzel and Alan H. Barr. A modeling system based on dynamic constraints. In John Dill, editor, *Computer Graphics (SIGGRAPH '88 Proceedings)*, volume 22, pages 179–188, August 1988.

[2] Jules Bloomenthal, editor. *Introduction to Implicit Surfaces*. Morgan Kauffman Publishers, Inc., 1997.

[3] Cassidy J. Curtis, Sean E. Anderson, Joshua E. Seims, Kurt W. Fleischer, and David H. Salesin. Computer-generated watercolor. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 421–430. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[4] David P. Dobkin, Silvio V. F. Levy, William P. Thurston, and Allan R. Wilks. Contour tracing by piecewise linear approximation. *ACM Transactions on Graphics*, 9(4):389–423, 1990.

[5] E. Bruce Goldstein. *Sensation and Perception.* Brooks/Cole Publishing Company, 1996.

[6] Victoria Interrante. Perceiving and representing shape and depth. SIGGRAPH 97 Course Notes for Principles of Visual Perception and Its Applications in Computer Graphics, August 1997.

[7] Lee Markosian, Michael A. Kowalski, Samuel J. Trychin, Lubomir D. Bourdev, Daniel Goldstein, and John F. Hughes. Real-time nonphotorealistic rendering. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 415–420. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[8] Scott D. Roth. Ray casting for modeling solids. *Computer Graphics and Image Processing*, 18(2):109–144, 1982.

[9] Michael P. Salisbury, Sean E. Anderson, Ronen Barzel, and David H. Salesin. Interactive pen–and–ink illustration. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 101–108. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.

[10] Michael P. Salisbury, Michael T. Wong, John F. Hughes, and David H. Salesin. Orientable textures for image-based pen-and-ink illustration. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 401–406. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[11] Mike Salisbury, Corin Anderson, Dani Lischinski, and David H. Salesin. Scale-dependent reproduction of pen-and-ink illustrations. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 461–468. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.

[12] Barton T. Stander and John C. Hart. Guaranteeing the topology of an implicit surface polyganization for interactive modeling. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 279–286. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[13] Georges Winkenbach and David H. Salesin. Computer–generated pen–and–ink illustration. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 91–100. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.

[14] Georges Winkenbach and David H. Salesin. Rendering parametric surfaces in pen and ink. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 469–476. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.