A MORPHOLOGICAL APPROACH TO THE VOXELIZATION OF SOLIDS

J. Andreas Bærentzen†, Miloš Šrámek‡, and Niels Jørgen Christensen†

†Graphical Communication, Technical University of Denmark, DK2800 Lyngby, Denmark, {jab|njc}@gk.dtu.dk

‡Austrian Academy of Sciences, Sonnenfelsgasse 19/2, A-1010 Vienna, Austria, Milos.Sramek@oeaw.ac.at

ABSTRACT

In this paper we present a new, morphological criterion for determining whether a geometric solid is suitable for voxelization at a given resolution. The criterion embodies two conditions, namely that the curvature of the solid must be bounded and the critical points of the distance field must be at a certain distance from the boundary of the solid. For solids that fulfill this criterion, we present an analytic and an empirical bound for the trilinear reconstruction error. Additionally, we give a theoretical argument as to why the distance field approach to voxelization is more sound than the prefiltering technique. The essence of the argument is that while sampling and interpolation must always introduce some error, the latter method (but not the former) introduces an error in the surface position independently of the sampling.

Keywords: Voxelization, Morphology, Geometric modeling, Curvature, Hesse normalform

1 INTRODUCTION

Volume graphics is the broad term used to describe a set of techniques in 3D computer graphics that employ discrete representations of 3D objects rather than continuous implicit or parametric representations. Volume graphics has important applications in certain areas of computer graphics such as the modeling of amorphous objects (clouds, smoke &c.) and the interactive modeling of certain types of solids. The latter application is usually known as volume sculpting [Galye91, Wang95, Bæren98]. Volume sculpting is not yet a very widespread technique, but we believe that it may soon become more popular, since the volumetric representation allows for very intuitive tools, and is more amenable to modeling objects with organic and complex shapes than boundary representations.

However, one of the impediments to a widespread use of volume graphics is that some of the fundamental operations still need theoretical work. The aim of this paper is to improve the underpinnings of one of these operations, namely the voxelization of solids.

Hitherto, two main paradigms for voxelization have been proposed:

- The prefiltering approach [Wang93], where a geometric solid is numerically convolved with a bandlimiting filter in the continuous domain, before sampling.
- The distance field approach [Breen98, Gibso98, Šráme99, Šráme98], where the idea is to sample the distance to the solid. This approach can also be modified so that a function of the distance is sampled rather than the distance itself [Šráme99].

After some preliminary definitions in section 2, we discuss the *prefiltering* and *distance field* techniques for voxelization in section 3 and argue why the latter is preferable. In section 4 we present a set of conditions for whether a solid is suitable for voxelization at a given resolution. In section 5 a criterion for whether a geometric solid is suitable for voxelization is presented. In section 6 we present two error bounds for the reconstruction of voxelized solids that fulfill the criterion. The first error bound is analytic the second is based on empirical data.

We only investigate the reconstruction error for the trilinear interpolation function, because trilinear reconstruction is fast, usually adequate for *volume*

graphics (even if it is not always adequate for visualization tasks) and usually the chosen interpolation for hardware implementations such as in the cube architecture which has recently been implemented in the VolumePro system [Pfist99].

Lastly, we draw conclusions and discuss future work in the sections 7 and 8.

2 DEFINITIONS

Solid By a solid X we understand a closed subset of R^3. We define the interior of X as the set itself ($\mathbf{p} \in X$), the exterior as the complement ($\mathbf{p} \notin X$), and the boundary ($\mathbf{p} \in \partial X \subset X$) as the subset where any neighborhood contains non members. Dangling boundaries are not allowed, i.e. there must be a path from a boundary point to a non–boundary point which does not touch other boundary points. The boundary of a solid is, of course, a *surface* in R^3 and the words surface and boundary will be used interchangeably.

Inside–outside function The inside–outside function returns 0 for points outside the object and 1 for points inside

$$B_X(\mathbf{p}) = \begin{cases} 1 & \mathbf{p} \in X \\ 0 & \mathbf{p} \notin X \end{cases} \tag{1}$$

Distance field By a distance field, we understand a scalar field associated with a solid X. The value of the field is given by a function $F_X : R^3 \rightarrow R$ that maps a point in space to the distance from that point to the closest point on ∂X.

$$F_X(\mathbf{p}) = \begin{cases} -\min_{\forall q \in \partial X}(|\mathbf{p}-\mathbf{q}|) & \mathbf{p} \in X \\ 0 & \mathbf{p} \in \partial X \\ \min_{\forall q \in \partial X}(|\mathbf{p}-\mathbf{q}|) & \mathbf{p} \notin X \end{cases} \tag{2}$$

As is apparent from (2) we use the convention that F_X is positive outside and negative inside the solid, so it is really an oriented distance function which is also called a *Hesse normalform* [Hartm99]. The normalform has several properties that we will need later. For instance, $|\nabla F| = 1$ and the principal curvatures can be inferred from the *Hessian* (i.e. the matrix of second order derivatives) of the normalform. It should be noted that the distance field (normalform) is only known explicitly for planes and spheres, and the normalform calculations must, in general, be done numerically.

Maximum curvature In this paper, we will not use Gaußian or average curvature, so *curvature* means normal curvature [Carmo76]. By the maximum curvature at a point $\mathbf{p}$, we mean the numerically greatest principal curvature [Carmo76] at $\mathbf{p}$ of the isosurface that contains $\mathbf{p}$. By the maximum curvature of a distance field in some region $S \subset R^3$ we understand the maximum of the maximum curvatures of all $\mathbf{p} \in S$

Voxel A *voxel* is usually defined either as a small rectangular box or a point sample of a 3D function. In this paper only the latter definition is used. The voxels are arranged in a rectangular, isotropic 3D lattice, and a neighboring voxel is one of the six voxels that are closest along one of the six principal directions of the lattice.

Voxel unit The voxel unit *vu* is the distance between two neighboring voxels. All distances are in voxel units.

2.1 An example

A typical example of a solid is the sphere. The interior of a sphere with centre $\mathbf{p_0}$ and radius r is given by

$$X = \{\mathbf{p} : |\mathbf{p} - \mathbf{p_0}| \leq r\} \tag{3}$$

The boundary is given by

$$\partial X = \{\mathbf{p} : |\mathbf{p} - \mathbf{p_0}| = r\} \tag{4}$$

and the distance function is

$$F_X(\mathbf{p}) = |\mathbf{p} - \mathbf{p_0}| - r \tag{5}$$

3 VOXELIZATION TECHNIQUES

The first work on non–binary volume sampling of geometric primitives (solids or polygons) was done by Wang and Kaufman in [Wang93]. Their method, known as *prefiltering* was to convolve the inside–outside function of a geometric primitive with a Bartlett filter[1] before sampling in order to band-limit the function. It is only necessary to know the value of the convolution at voxel positions, hence a numerical solution is feasible, and the method was successful in producing voxelized objects with few visible aliasing artifacts.

Recently, another and simpler technique has been employed for solid voxelization by e.g. Šrámek [Šráme98, Šráme99], Gibson [Gibso98], and Breen [Breen98]. The idea is to simply sample the distance

[1] Also known as the hypercone filter. The filter has its maximum in the centre of the support and the value decreases linearly with the distance to the centre to 0 at the edge of the support

function (2) or a function that is proportional to the distance function. It is possible to sample and interpolate the distance function just like the convolved inside–outside function, but this approach has the advantage that it is simpler (in fact the prefiltering method uses the distance field), and has experimentally been shown to yield superior results [Šráme98].

Various reconstruction filters may be applied to the voxel raster to reconstruct the value at arbitrary locations, and even the trilinear filter yields quite good results. Šrámek shows experimentally that the surface reconstruction error for a sphere decreases as the radius increases and reports an average error of less than 0.05 *vu* [Šráme99] for the reconstruction of a sphere of a radius of 4 *vu*. Both Gibson and Šrámek conjecture that the error is curvature dependent, and Gibson also notes that certain special cases must be taken into account. These special cases are when critical points in the distance field come so close to the surface that they are within the support of reconstruction or gradient reconstruction filters. This can either be due to sharp edges or, in the case of an object with a smooth surface, due to two surfaces (or surface components) that are close to each other.

3.1 Prefiltering

The enticing thing about the prefiltering approach is that the operation of bandlimiting by convolving with a smoothing filter is a well known operation that is frequently used in computer graphics. However, in volume graphics, the method has a drawback. Surfaces of solids are almost always defined as isosurfaces in the scalar field used to represent the solid, and the value of the convolution at a point on ∂X is only constant for all $\mathbf{p} \in \partial X$ if the curvature of the surface of X is constant. Therefore, there is, in general, no isovalue σ for which

$$\{\mathbf{p} : (B_X \star Ba)(\mathbf{p}) = \sigma\} = \partial X \tag{6}$$

where $\star$ denotes convolution and Ba is the Bartlett filter. This problem does not exist in the distance field approach since by definition of the distance field

$$\{\mathbf{p} : F_X(\mathbf{p}) = 0\} = \partial X \tag{7}$$

The problem is illustrated in figure 1 where we observe that only a planar surface divides a spherical support in two identical halves when the centre of the support is exactly on the surface of the solid. The greater the curvature at the boundary point, the greater the difference between the part of the support that intersects the solid and the part that does not, and as the filter is non–zero within the support, the result of the convolution will also differ. Note that the error is not a byproduct of sampling and interpolation,

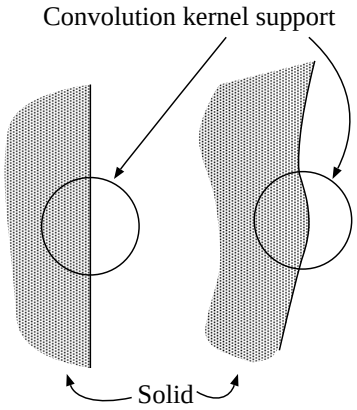

Figure 1: Intersection of solid and filter support

but an intrinsic problem with the method which suggests that prefiltering may not be the best paradigm for voxelization.

3.2 Distance field sampling

It has been mentioned that high curvature is known to reduce the quality in voxelization according to the distance field approach, and furthermore certain special cases should be avoided. These special cases have in common that they occur whenever the medial surface of the solid or the complement of the solid comes too close to the surface of the solid. A medial surface (See also appendix) is the locus of points that are equidistant from at least two points on the boundary of the solid. These points are the critical points of the distance field where the gradient is not defined, and since the gradient is used in shading (to estimate the surface normal), the gradient filter should not use samples that are distributed on both sides of a medial surface. The medial surface may be close to the

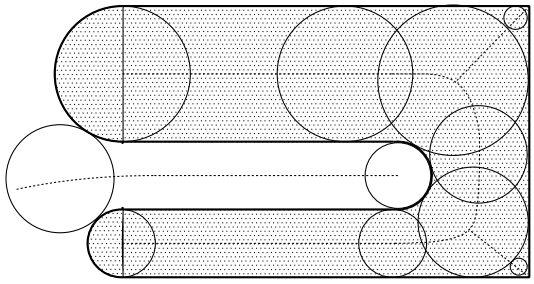

Figure 2: Medial axis of a solid and parts of the medial axis of the complement.

boundary either due to a sharp edge (where the medial surface touches the surface of the solid) or a very thin structure. A 2D example of medial surfaces is shown in figure 2 where the dashed lines indicate the

medial axes. (The 2D counterpart of medial surfaces are medial axes).

4 CONDITIONS FOR VOXELIZATION SUITABILITY

The observations in the previous section can be presented more concisely as two conditions for whether an object is suitable for voxelization

Condition 1 The curvature should be low relative to the resolution. This reduces reconstruction error.

Condition 2 The reconstruction and gradient reconstruction filters should not use samples that are distributed on both sides of a medial surface of X or X^c.

The second condition can be restated as: No point on the medial surface should be closer to ∂X than $\sqrt{6}$. Because, if $\mathbf{p}$ is a point on the surface of X then the gradient value is calculated at the eight nearest voxels and trilinearly interpolated at $\mathbf{p}$. The values at a total 28 voxels are used. (The voxel configuration is shown in figure 3). $\mathbf{p}$ must by construction be within the

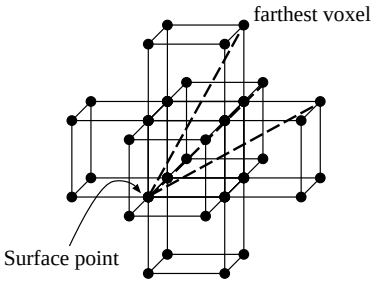

Figure 3: Voxels used in gradient computation

cube whose corners are the eight nearest voxels, and it is possible to ascertain by visual inspection that the greatest distance from a point within that cube to any voxel in the configuration is $\sqrt{6} = \sqrt{2^2 + 1^2 + 1^2}$.

The aim of the next section is to define a single criterion that comprises both of the above conditions.

5 MORPHOLOGICAL CRITERION

What is needed is some sort of measure that takes both curvature and overall feature size into account. Fortunately, such measures exist in mathematical morphology (see appendix for definitions).

Let S_r denote a sphere of radius r and X a solid which is S_r–open and S_r–closed, then X has the following two additional properties

Property 1 Given a point $\mathbf{p}$ for which $F_X(\mathbf{p}) = \sigma$ where $-r < \sigma < r$ the following holds for κ, the maximum curvature at $\mathbf{p}$

$$\kappa \leq \frac{1}{r - |\sigma|} \tag{8}$$

Property 2 The medial surface is nowhere closer to the boundary ∂X than r.

Property 2 follows directly from the definition of the medial axis (see appendix).

Proof of property 1

Without loss of generality, we assume that $\mathbf{p}$ is a point in the interior of X. Let $\mathbf{p}_0$ be the point on ∂X closest to $\mathbf{p}$. It is required that X is S_r–open and S_r–closed. This means that S_r can be translated so that it touches $\mathbf{p}_0$ from either side. The exterior instance, S_{r_1}, of S_r does not touch interior points of X and the opposite holds for the interior instance S_{r_2}. The configuration of $\mathbf{p}$ and the translated instances of S_r is shown in figure 4. It is clear that the translated instances of

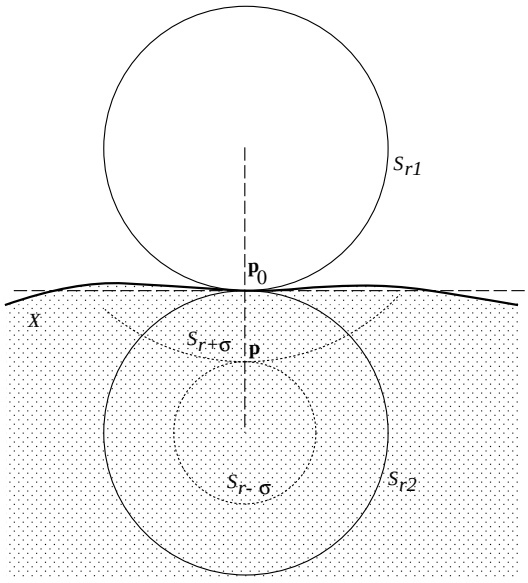

Figure 4: Translated instances of S_r touch $\mathbf{p}$ from either side.

S_r must share tangent plane with each other and with ∂X.

Now, let $S_{r+\sigma}$ be a sphere of radius $r + \sigma$ which has the same centre as S_{r_1}, and let $S_{r-\sigma}$ be a sphere of

radius $r - \sigma$ with the same centre as S_{r_2}. Any new point $\mathbf{p}_1$ near $\mathbf{p}$ on the same isosurface $F_X(\sigma)$ must lie on or between the two spheres $S_{r+\sigma}$ and $S_{r-\sigma}$, because assuming otherwise leads to a contradiction:

Assume that $\mathbf{p}_1$ is inside $S_{r-\sigma}$. Since the distance from $\mathbf{p}_1$ to the surface is σ the surface intersects S_{r_2} which violates the S_r–openness.

Assume that $\mathbf{p}_1$ is inside $S_r + \sigma$. By the S_r–closedness property there must be a point on S_{r_1} which has shorter distance to $\mathbf{p}_1$ than σ violating that $F_X(\mathbf{p}_1) = \sigma$.

If all points on the σ isosurface lie between $S_{r+\sigma}$ and $S_{r-\sigma}$ then the smallest osculating sphere of a curve on the σ–isosurface at the point $\mathbf{p}$ is $S_{r-\sigma}$. Hence, the greatest normal curvature at $\mathbf{p}$ is indeed $\frac{1}{r-|\sigma|}$.

5.1 Putting the criterion together

There is an obvious correspondence between the two properties of this section and the two conditions from the previous section. In fact, if r is chosen large enough, property 1 ensures that condition 1 is fulfilled. Likewise, if $r > \sqrt{6}$ it follows from property 2 that condition 2 must be fulfilled. More concisely:

Voxelization suitability criterion

A geometric solid X is suitable for voxelization at a given resolution, if X is S_r–open and S_r–closed where $r > \sqrt{6}$ is chosen so that the reconstruction error is sufficiently low for the application.

Note that by choosing a radius r we also choose resolution, since r is in voxel units.

6 ERROR BOUNDS

The above criterion is, of course, only really interesting if we can say something about the error so that it is possible to determine whether the error for a given r is "sufficiently low". In this section, we will develop a (somewhat loose) analytic error bound for the reconstruction error and afterwards a tighter empirically based error bound. These error–bounds can then be used to determine what r to plug into the criterion we found in the preceding section.

First, we need a theorem about linear interpolation: Let $f(x)$ be a function which is continuous on $[a, b]$ and twice differentiable on (a, b), and let there be given a linear interpolation function

$$h(x) = \frac{f(a)(b-x) + f(b)(x-a)}{b-a} \tag{9}$$

which interpolates between the value of f at a and b.

$$f(x) - h(x) = \frac{(x-a)(x-b)}{2} f''(c) \tag{10}$$

where $c \in (a, b)$. Using (10), it is easy to show that given a bound on the second order derivative $|f''(x)| \leq M$ we also have a bound on the interpolation error

$$|f(x) - h(x)| \leq \frac{(b-a)^2}{8} M \tag{11}$$

The proofs of the above may be found in [Young88].

6.1 Analytic error bound

Using (11) we will now derive an error bound for trilinear interpolation in a voxelized distance field[2]. Given a distance field $F : R^3 \rightarrow R$ and a line segment between two neighboring voxels $\mathbf{a}$ and $\mathbf{b}$, we know that the value of the field along the line from $\mathbf{a}$ to $\mathbf{b}$ is

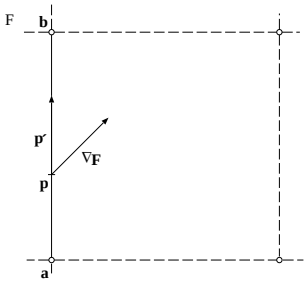

Figure 5: Line segment from $\mathbf{a}$ to $\mathbf{b}$ in distance field F

$$f(s) = F(\mathbf{p}(s)) \tag{12}$$

where $\mathbf{p}$ is a parameterized line

$$\mathbf{p}(s) = s(\mathbf{b} - \mathbf{a}) + \mathbf{a} \tag{13}$$

and $|\mathbf{p}'(s)| = 1$ since $\mathbf{a}$ and $\mathbf{b}$ are neighboring voxels.

To find the derivative of the function f, we apply the chain rule to the right hand side of (12) yielding

$$f'(s) = \nabla F \cdot \frac{dp}{ds} = \begin{pmatrix} F_x(\mathbf{p}(s)) \\ F_y(\mathbf{p}(s)) \\ F_z(\mathbf{p}(s)) \end{pmatrix} \begin{pmatrix} b_x - a_x \\ b_y - a_y \\ b_z - a_z \end{pmatrix} \tag{14}$$

[2] recall from section 2 that a distance field is just a scalar field, where the scalar value is the distance to the surface of the solid that is represented by the field

The dot product yields a three term expression for f', and to get f'' all we need to do is to apply the chain rule to each of these three terms. The result is a nine term sum, where each term is the product of one of the second order partial derivatives of F and the corresponding two components of $\mathbf{p}'$. This nine term sum can be written in matrix notation in the following way

$$f''(s) = \mathbf{p}'(s) H(\mathbf{p}(s)) \mathbf{p}'(s) \tag{15}$$

where H is the Hessian of F, i.e. the matrix of the second order partial derivatives

$$H = \begin{pmatrix} F_{xx} & F_{xy} & F_{xz} \\ F_{yx} & F_{yy} & F_{yz} \\ F_{zx} & F_{zy} & F_{zz} \end{pmatrix} \tag{16}$$

To find a bound for $|f''|$ all we need to do is find the numerical maximum of the right hand side of (15).

This turns out to be simple, because F fulfills the requirements of a *Hesse normalform* [Hartm99], and it is known from the theory about such, that the Hessian of the normalform (i.e. the Hessian of F) has three eigenvalues $\lambda_0 = 0$, $\lambda_1 = \kappa_{min}$, $\lambda_2 = \kappa_{max}$ corresponding, respectively, to the direction of the gradient ($\mathbf{n}$) and the directions of minimum and maximum curvature ($\mathbf{t}_{min}$ and $\mathbf{t}_{max}$. Since any vector $\mathbf{v} \in R^3, |v| = 1$ can be expressed as a linear combination of these three eigenvectors,

$$\mathbf{v} = a\,\mathbf{n} + b\,\mathbf{t}_{min} + c\,\mathbf{t}_{max} \tag{17}$$

where $|\mathbf{v}| = \sqrt{a^2 + b^2 + c^2} = 1$, we know that

$$\begin{aligned} |\mathbf{v}^T H \mathbf{v}| &= |a\lambda_0 + b\lambda_1 + c\lambda_2| \\ &\leq |\lambda_2| = |\kappa_{max}| \end{aligned} \tag{18}$$

Consequently,

$$|f''(s)| \leq |\kappa| \tag{19}$$

where κ is the maximum curvature at all points on the line segment between $\mathbf{a}$ and $\mathbf{b}$ (See section 2 for a definition of maximum curvature). Using (19) and (11) we obtain

$$\text{lin\_err} = \frac{1}{8}|\kappa| \tag{20}$$

Of course, our real interest is in the trilinear interpolation function. A trilinear interpolation may be perceived as a linear interpolation of two values that are pairwise linearly interpolated between four values which are interpolated between the eight original voxels. These seven linear interpolations are shown in figure 6. To do a worst case analysis of the cumulative error, let us begin with the value IA0. IA0 is linearly interpolated between the voxels V0 and V1 and the maximum interpolation error is known to be

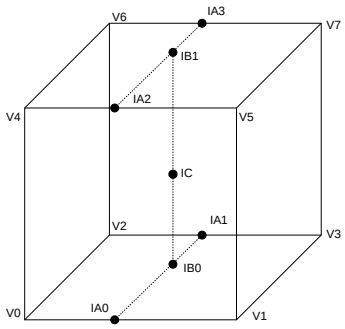

Figure 6: The seven linear interpolations that constitute trilinear interpolation

lin_err. IA1 has the same maximum error. IB0 is interpolated between IA0 and IA1. If we knew the exact values at IA0 and IA1, it would follow that the maximum error at IB0 would be just lin_err. However, we must take into account that we are interpolating between interpolated values. Fortunately, we know that (for linear interpolation) the difference between interpolation between exact values and interpolation between imprecise values can not be greater than the greatest of the two errors associated with the imprecise values. In the present case, the interpolation is between IA0 and IA1 both of which differ from the exact values by at most lin_err. Therefore, to obtain a bound for the total error at IB0, we must add lin_err to the linear interpolation error bound at IB0 yielding a total error bound of 2 lin_err. By a similar argument, we may conclude that the total error bound at IC which is interpolated between IB0 and IB1 is 3 lin_err, hence

$$\text{trilin\_err} = \frac{3}{8}|\kappa| \tag{21}$$

where $|\kappa|$ is the maximum curvature within the cell.

The final important question is to find the maximum curvature within the cell. According to property 1, we can find the maximum curvature by finding the greatest distance from any point in the cell to the surface of the solid and plugging that distance into (8). We are only interested in cells which intersect the surface, so the greatest possible distance from the surface of any point in the cell is $\sqrt{3}$, and the final expression for the reconstruction error as a function of the radius r of our structuring element S_r becomes

$$\text{err}(r) = \frac{3}{8(r - \sqrt{3})} \tag{22}$$

where, according to the suitability criterion, $r > \sqrt{6}$. It is obvious, unfortunately, that the bound is somewhat loose, since we have to make worst case assumptions at every step, but it is difficult to make a tighter bound without making assumptions about the

shape of the solid or the configuration of the solid and the trilinear cell. A plot of err(r) can be seen in figure 7

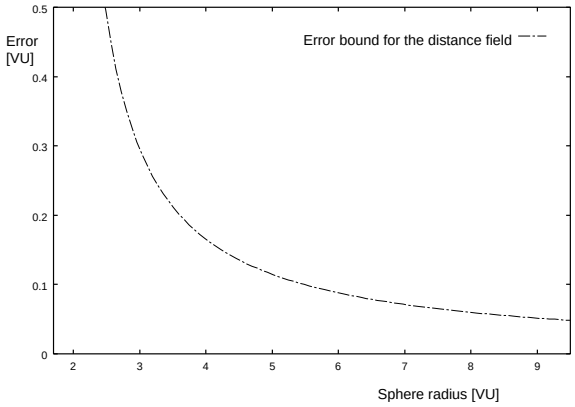

Figure 7: The error function.

6.2 Empirical error bound

The analytic error-bound shows us that the reconstruction error decreases as the curvature decreases. For a general solid X which fulfills the suitability criterion for some sphere S_r, we know that the curvature of isosurfaces in X is less than or the same as that of isosurfaces corresponding to the same isovalues in S_r. Therefore, we would assume that the worst case reconstruction error of X is not worse than the worst case reconstruction error of S_r. In light

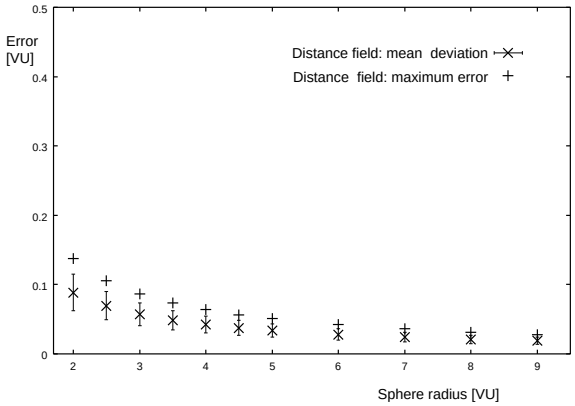

Figure 8: Maximum and mean reconstruction error for a sphere as a function of radius. Standard deviation for mean error is also shown.

of this hypothesis, we propose a much tighter bound which is based on our empirical results. We voxelized spheres of radii ranging from 2 *vu* to 9 *vu* and sent rays from the centre of the spheres towards their periphery. Where the rays hit the level 0 isosurface, we measured the error as the distance to the true sphere surface. Figure 8 shows the maximum and mean error. Note that this error measure is slightly different from that of the previous section. The analytic error bound bounds the greatest difference between the value of the true and interpolated distance functions, while the empirical error shown in figure 8 is the geometric shortest distance from the point on the interpolated isosurface to the true sphere.

We notice that at a sphere radius of $r = 3 > \sqrt{6}$ the error has fallen below 0.1 voxel unit, and for many applications this error should be acceptable.

7 CONCLUSIONS

The prefiltering approach to voxelization has been shown, experimentally, to yield less precise volumetric models than the approach based on distance field sampling [Šráme98]. In this paper, we have given a theoretical argument as to why the prefiltering approach is problematic, namely that the method does not, in general, produce an isosurface which corresponds, precisely, to the original solid.

It is known that the reconstruction error when (trilinearly) reconstructing distance field sampled volumetric data is due to curvature [Šráme98, Gibso98]. In addition, certain special cases due to critical points in the vicinity of the solid boundary must be taken into account [Gibso98]. We have shown that by formulating a suitability criterion in terms of the morphological properties *openness* and *closedness*, it is possible to take into account the quality loss due to curvature as well as the problems that are due to these special cases. Furthermore, we have provided error bounds for the reconstruction error of solids that fulfill the suitability criterion. While the analytic error bound is loose, we believe that the empirical error bound should be a practical tool for choosing voxelization resolution.

8 FUTURE WORK

For simple geometric solids whose shape and curvature are known, it is not difficult to verify whether they fulfill the criterion. For more complex, perhaps composite, solids, it is frequently obvious that they do not fulfill the criterion (e.g. if we know the object has a sharp edge), but we want to voxelize them anyway. Therefore, a general method for finding out whether a given (implicit) solid fulfills the criterion would probably be less useful than a method for filtering complex solids so that they fulfill the criterion. This filtering can, obviously, be performed by applying the digital versions of the morphological open and close

filters to the solid. These filters should be applied before, or maybe as a part of the sampling process. However, some difficulties are strewn along the way since a naïve implementation would either introduce gross imprecisions or be very computationally demanding, and furthermore the sequence of operations is significant since, in general, $O[C[X]] \neq C[O[X]]$.

Lastly, a purely analytic error bound is the theoretically most pleasing, and a tightening of the analytic error bound is, indeed, a part of our plans for future work.

REFERENCES

[Breen98] David E. Breen, Sean Mauch, and Ross T. Whitaker. 3d scan conversion of csg models into distance volumes. In Stephen Spencer, editor, *Proceedings of IEEE Symposium on Volume Visualization*, October 1998.

[Bæren98] Andreas Bærentzen. Octree–based volume sculpting. In Craig M. Wittenbrink and Amitabh Varshney, editors, *LBHT Proceedings of IEEE Visualization '98*, October 1998.

[Carmo76] Manfredo P. Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.

[Galye91] Tinsley A. Galyean and John F. Hughes. Sculpting: An interactive volumetric modeling technique. *ACM Computer Graphics*, 25(4), July 1991.

[Gibso98] Sarah F.F. Gibson. Using distance maps for accurate surface representation in sampled volumes. In Stephen Spencer, editor, *Proceedings of IEEE Symposium on Volume Visualization*, October 1998.

[Hartm99] Erich Hartmann. On the curvature of curves and surfaces defined by normal forms. *Computer Aided Geometric Design*, 16(5):355–376, 1999.

[Pfist99] Hanspeter Pfister, Jan Hardenbergh, Jim Knittel, Hugh Lauer, and Larry Seiler. The volumepro real–time ray–casting system. In *SIGGRAPH 1999 Conference Proceedings*, pages 251–260, 1999.

[Šráme98] Miloš Šrámek and Arie Kaufman. Object voxelization by filtering. In Stephen Spencer, editor, *Proceedings of IEEE Symposium on Volume Visualization*, October 1998.

[Šráme99] Miloš Šrámek and Arie Kaufman. Alias–free voxelization of geometric objects. *IEEE Transactions on Visualization and Computer Graphics*, 5(3), July/September 1999.

[Wang93] Sidney Wang and Arie E. Kaufman. Volume sampled voxelization of geometric primitives. In Gregory M. Nielson and Dan Bergeron, editors, *Proceedings, Visualization 93*. IEEE, 1993.

[Wang95] Sidney Wang and Arie E. Kaufman. Volume sculpting. In *1995 Symposium on Interactive Graphics*. ACM SIGGRAPH, 1995.

[Young88] David M. Young and Robert Todd Gregory. *A Survey of Numerical Mathematics*. Dover, 1988.

APPENDIX A: MORPHOLOGY

Open and close

The open operation of a set X with respect to a structuring element b is

$$O[X, b] = \bigcup \{b_{\mathbf{x}} : b_{\mathbf{x}} \subset X\} \tag{23}$$

The close operation of X with respect to b may be expressed in terms of the open operation

$$C[X, b] = O[X^c, b]^c \tag{24}$$

where $\mathbf{x}$ is a vector and $b_{\mathbf{x}}$ is the structuring element b translated according to $\mathbf{x}$. Intuitively, the open operation corresponds to moving the structuring element b around inside the set X. The result of the operation is the subset of X where b fits. The little protrusions where b does not fit are cut off. Similarly, the close operation fills out the cavities where b does not fit.

One of the important properties shared by both open and close is *idempotence*:

$$O[O[X, b], b] = O[X, b] \tag{25}$$
$$C[C[X, b], b] = C[X, b] \tag{26}$$

If we have already applied open or close to an object, further applications of the operator do not change the result. A set X which is not changed by an open operation with a structuring element b is called b–open. A set which is not changed by a close operation with a structuring element b is called b–closed.

The medial surface

Let $\mathbf{p} \in X$. $F_X(\mathbf{p})$ is the distance to ∂X. If there is more than one point $\mathbf{p}_i \in \partial X$ so that $|\mathbf{p}_i - \mathbf{p}| = F_X(\mathbf{p})$ we say that $\mathbf{p}$ is in the medial surface. More intuitively: Let X be a solid. If $\mathbf{p}$ is the centre of a sphere S_1, and there is no sphere of greater radius S_2 which properly includes S_1 whilst itself being included in X, then $\mathbf{p}$ belongs to the medial surface.