MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Calculating the Distance Map for Binary Sampled Data

Sarah F. Frisken Gibson

TR99-26 December 1999

Abstract

High quality rendering and physics-based modeling in volume graphics have been limited because intensity-based volumetric data do not represent surfaces well. High spatial frequencies due to abrupt intensity changes at object surfaces result in jagged or terraced surfaces in rendered images. Use of a distance-to-closest-surface function to encode object surfaces allows accurate reconstruction of objet surfaces for volumetric data. However, constructing the distance map for distance-based rendering requires a prior model of the object surface. Here we present a number of methods that can be used to estimate the distance map from a binary segmented volume, where no prior knowledge of object surfaces exists.

Calculating the Distance Map for Binary Sampled Data

Sarah F. F. Gibson

MERL – A Mitsubishi Electric Research Laboratory
201 Broadway, Cambridge, MA, 02139

Abstract

High quality rendering and physics-based modeling in volume graphics have been limited because intensity-based volumetric data do not represent surfaces well. High spatial frequencies due to abrupt intensity changes at object surfaces result in jagged or terraced surfaces in rendered images. Use of a distance-to-closest-surface function to encode object surfaces allows accurate reconstruction of objet surfaces for volumetric data. However, constructing the distance map for distance-based rendering requires a prior model of the object surface. Here we present a number of methods that can be used to estimate the distance map from a binary segmented volume, where no prior knowledge of object surfaces exists.

Keywords: Volume Rendering, Volume Graphics, Medical Applications, Surgical Simulation, Physics-based Graphics, Haptic Feedback

1 Introduction

In [2], it was shown that surfaces can be encoded into volumetric data by storing the distance-to-closest-surface value at each sample point in the data. The distance map has some attractive properties. First, the zero-value of the distance map locates surfaces while the gradient of the distance map yields surface normals. Second, when the sampling rate of the volumetric data is adequate (i.e. when it is large relative to the surface curvature), a low-cost 6-point central difference gradient estimator applied to the distance map can accurately reconstruct surface normals near the surface. Third, although folds and object edges or corners introduce singularities and nonlinearities into the distance field that cause shading artifacts with a central difference gradient estimator, the presence of these irregularities in the distance map can be easily detected during rendering so that higher order filters or more sophisticated gradient estimation methods can be locally applied. Fourth, in addition to applications in volume rendering, the distance map approach can be used to reconstruct surface normals and penetration distances for applications in haptics and physics-based modeling.

While the distance map enables high quality rendering and shading of surfaces in volumetric data, it requires a prior model of the surface. However, in binary-sampled data, exact knowledge of the underlying object surface is missing and distance maps must be approximated from the binary data rather than from the true surface. In this working paper, we present the results of experiments to test various existing methods for estimating distance maps from binary segmented data. The results of these tests show that existing methods do not provide artifact-free distance maps. Five different distance metrics are compared for a 2D circle and the artifacts that result from errors in the distance map are illustrated for a volume rendered sphere. This report uses a simple sphere so that artifacts can be easily compared. In practice, we have found that artifacts are particularly severe in data that is not isometric – where the data is sampled at different rates along different major axes. This is often the case in MRI imaging, where the in-plane sampling rate is frequently much higher that the spacing between planes.

Because these existing methods do not generate satisfactory distance maps from binary data, we have developed a new method for calculating distances from the surface of a binary object. This method, Constrained Elastic Surface Nets, is detailed in [1].

2 Calculating Distances from Binary Data

Binary volumetric data no longer contains an accurate representation of object surfaces. This section analyzes several techniques for estimating surfaces and their corresponding distance maps from binary databy comparing errors in these distance maps for a circular object. Volume rendered images for distance maps of a binary sphere that were created with some of these 5 methods illustrate the artifacts produced by errors in the distance approximations.

Perhaps the most common source of binary volumetric data is segmented medical data, where different tissues or structures are each assigned a unique classification or type. This paper does not advocate the use of a distance map with grey-scale data because the gradient of the grey-scale data often provides a good estimate of surface normals[3]. However, Figure 1, a 2D slice from a 3D MRI image of a human knee, illustrates a case where the grey-scale gradient does not give a good surface normal estimate. Figure 1b) was calculated by applying a central difference operator to Figure 1a) at hand-segmented edge points along the femur, one of the bones in the knee. Because the bone surface is generally smooth and of uniform density, we expect surface normals to have a relatively constant magnitude and a slowly varying direction. However, because the gradient depends on the thickness and image intensity of materials that are adjacent to the bone surface, and because there are a variety of materials adjacent to the femur, the estimated normals vary dramatically around the edge of the femur. Even when the magnitude of the gradient vector is normalized as in Figure 1c), unexpected and sudden changes in the direction of the image gradient would introduce severe artifacts into an image shaded using grey-scale gradients. For this reason, it is sometimes necessary to calculate surface normals from a binary segmentation of objects rather than from the original measured grey-scale data.

In binary-sampled data, exact knowledge of the underlying object surface is missing. Distance maps must be approximated from the binary data rather than from the true surface. Two basic strategies are: 1) estimate distances directly from the binary values; and 2) estimate the surface based on local binary values and calculate distances from this surface. In this section, we describe methods from both of these categories, illustrate errors in these methods by comparing estimated 2D distance maps to the true distance function of a circle, and illustrate the effect of the errors on surface normal calculation in images of volume rendered spheres.

Five different methods were used to calculate the distance from the edge of a binary 2D circle. Methods calculated directly from the binary values include the chessboard distance, the city-block distance, and the Euclidean distance. These methods use the distance metrics illustrated in Figure 2 and are detailed for 2D dis-

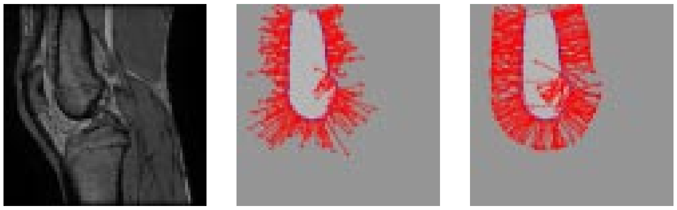

a) b) c)

Figure 1: *a) 2D Magnetic Resonance Image (MRI) cross section through a human knee. b) image gradient vectors calculated using central differences on the grey-scale data at edge points of the segmented femur. c) image gradient vectors with a normalized magnitude. The direction and magnitude of image gradients vary much more than we would expect the surface normals of the knee bone to vary, in some cases pointing inward when we expect an outward facing normal. Hence, applying a gradient operator to the grey-scale data does not always provide a good estimate of surface normals. (Data and segmented image courtesy of R. Kikinis, Brigham and Women's Hospital, Boston MA).*

4	3	2	3	4
3	2	1	2	3
2	1	0	1	2
3	2	1	2	3
4	3	2	3	4

2	2	2	2	2
2	1	1	1	2
2	1	0	1	2
2	1	1	1	2
2	2	2	2	2

$\sqrt{8}$	$\sqrt{5}$	2	$\sqrt{5}$	$\sqrt{8}$
$\sqrt{5}$	$\sqrt{2}$	1	$\sqrt{2}$	$\sqrt{5}$
2	1	0	1	2
$\sqrt{5}$	$\sqrt{2}$	1	$\sqrt{2}$	$\sqrt{5}$
$\sqrt{8}$	$\sqrt{5}$	2	$\sqrt{5}$	$\sqrt{8}$

Figure 2: Distance metrics used by the (left to right): city-block, chessboard and Euclidean distance methods. The figures illustrate the assigned distance of neighboring pixels from the central pixel for the three methods.

tances in [5][1]. Extending these algorithms to 3D is straightforward. The city-block distance and Euclidean distance methods have been implemented in 3D.

The remaining two methods approximate the surface from local binary values and then estimate distances to the approximate surface. The first method is analogous to the Marching Cubes algorithm [4]. It assumes that the surface lies in cubes bounded by 8 elements with different binary values. For each cube containing a surface, the surface is assumed to cross cube edges mid-way between elements with different binary values. In 2D, the constructed edge becomes a sequence of straight lines with endpoints midway between elements with different binary values. Once the edge is defined, distances from the edge are calculated using a modification of the Euclidean distance filter. In the second method, the constructed surface connects points at the center of each cube containing a surface. The distance map is computed from the mesh of center points using Euclidean distances. This central-point method has been implemented in 3D.

Figure 3 shows the true 2D distance map for a circle and the approximate distances calculated using the chessboard distance metric and the central-point distance method. Only small differences are visible from these intensity images. Table 1 reports the maximum and minimum errors of the estimated distance compared to the true distance for various methods. For some of the methods, various smoothing filters were applied to the distance maps and the resultant errors for these estimates are also presented. The smoothing filters utilized were a 3x3 Gaussian filter, and 5x5 Gaussian and box averaging filters. Figure 4 are images of the difference between the true distance map and the calculated distance maps for various distance estimation methods. In the images, a mid-level grey value represents zero error while black or white represent large negative or positive errors respectively. The magnitude and variance of the error along the edge of the circle in Figure 4 shows that for a circle, the smoothed central-point distance method provides the best estimate of the true distance.

[1] In [5], true Euclidean distances are approximated to allow for an integer-based algorithm in the interest of speed. In our implementation, floating point distances were used.

Dist. Measure	*Filter*	*Max Error*	*Min Error*
city block	none	42.39	-8.67
chessboard	none	22.00	-29.61
euclidean	none	5.15	-2.08
euclidean	3x3 Gaussian	5.15	-1.92
euclidean	5x5 Gaussian	5.15	-1.83
euclidean	5x5 box	5.15	-1.73
marching cubes	none	1.02	-1.02
central point	none	0.71	-0.69
central point	5x5 box	2.49	-0.69

Table 1: *Errors (in pixels) between the true distance map and estimated distance maps for a binary circle of radius 50 pixels using various distance measures. As indicated, in some cases the approximate distance fields were smoothed with a Gaussian or a box averaging filter.*

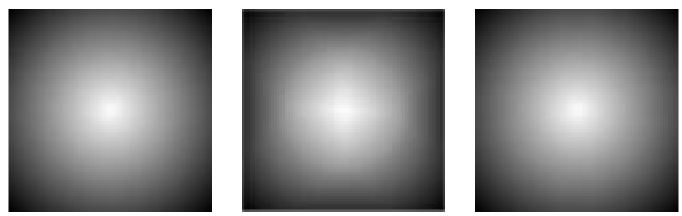

a) b) c)

Figure 3: *Distances from the edge of a circle of radius 50 centered in a 100x100 image. A grey value of 255 (white) represents a positive distance, a grey value of 0 (black) represents a negative distance and a grey value of 127 represents a distance of zero. a) is the true distance map for a 2D circle. b) is the distance map approximated from a binary circle using the chessboard distance and c) is the distance map approximated from the central-point distance method.*

Figure 5 presents images rendered from various distance maps estimated from a binary representation of a sphere of radius 30 voxels. The distance map for the images were calculated using the city-block distance, the Euclidean distance, the central-point distance, and the central-point distance smoothed with a 5x5x5 box averaging filter. The image generated from the smoothed central-point distance estimator has a relatively smooth surface that represents the original data reasonably well.

3 Conclusions

The ability to represent surfaces accurately in volume graphics opens up many research directions in volume rendering, haptics, and physics based modeling for applications such as surgical simulation which require a volumetric object representation. Using a distance map representation allows accurate reconstruction of object surfaces. However creating a distance map requires an estimate of the object surface.

For binary objects, the accuracy of the distance map is limited

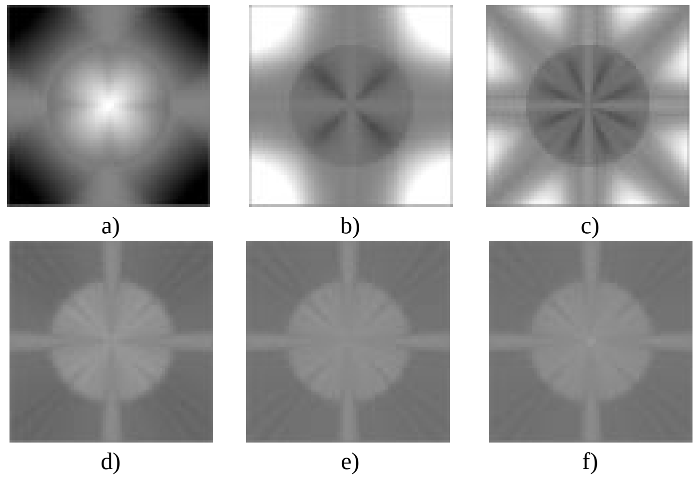

Figure 4: *Errors in distance maps estimated from a binary circle using: a) chessboard distance; b) cityblock distance; c) Euclidean distance; d) marching-cubes edge distance; e) central-point distance; and f)central point distance smoothed with a 5x5 box averaging filter. The error images were calculated by subtracting the true distance from the estimated distance and linearly scaling and centering the error so that zero error corresponds to a grey value of 127, and positive errors are brighter and negative errors are darker than 127. The city block and chessboard errors were significantly larger and were scaled here by a factor of 12.75. The other errors were scaled here by a factor of 51.*

by the best available estimate of the original surface. Most existing techniques and all of the techniques analyzed above use local filtering methods to estimate surfaces from binary data. However, especially when surfaces are at shallow angles to the sampling grid, or when the data sampling is not isometric, local filters give unsatisfactory surface estimates. In [1], a new method is presented that uses a more global approaches to surface estimation by fitting an elastic net over the binary surface that relaxes to smooth out the shape of the surface but is constrained to adhere to the original binary segmentation. This method provides much better results than the distance metrics that have been described in this working paper.

References

[1] S. Gibson. Constrained elastic Surface Nets: generating smooth surfaces from binary segmented data. In *submitted to Medical Image Computation and Computer Assisted Surgery, 1998*, pages proceedings, MICCAI'98, 1998.

[2] S. Gibson. Using distance maps for accurate surface representation in sampled volumes. In *submitted to IEEE Visualizaiton, 1998*, pages proceedings, Visualization'98, 1998.

[3] K. Hohne, M. Bomans, A. Pommert, M. Riemer, C. Schiers, U. Tiede, and G. Wiebecke. 3d visualization of tomographic volume data using the generalized voxel model. *The Visual Computer*, 6(1):28–36, February 1990.

[4] W. Lorensen and H. Cline. Marching cubes: a high resolution 3d surface construction algorithm. In *Computer Graphics Proceedings, Annual Conference Series*, Proceedings of SIGGRAPH 87, pages 163–169. ACM SIGGRAPH, 1989.

[5] J. Russ. *The Image Processing Handbook*. CRC Press, Boca Raton, FL, 1992.

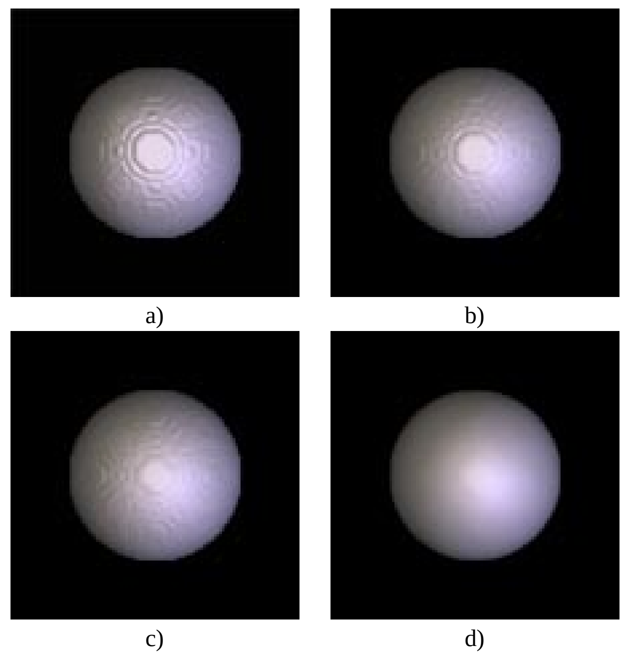

Figure 5: *Volume rendered images shaded using central differences applied to distance maps estimated from a binary sampled sphere of radius 30 voxels. The distance estimation methods were: top left, city block distance; top right, Euclidean distance; bottom left, central point distance; and bottom right, the central point distance map smoothed with a 5x5x5 box averaging filter. The last method produces acceptable shading for a 30 voxel radius binary sphere.*