

UC Davis

IDAV Publications

Title

High-Quality Rendering of Smooth Isosurfaces

Permalink

<https://escholarship.org/uc/item/44q319s8>

Journal

Journal of Visualization and Computer Animation, 10

Authors

LaMar, Eric C.
Hamann, Bernd
Joy, Ken

Publication Date

1999

Peer reviewed

- [25] SCHRÖDER, P., AND STOLL, G. Data parallel volume rendering as line drawing. pp. 25–32.
- [26] SHEKHAR, R., FAYAD, E., YAGEL, R., AND CORNHILL, F. Octree-based decimation of marching cubes surfaces. In *Proceedings of Visualization'96* (Sept. 1996), pp. 335–342.
- [27] TOFFOLI, T., AND QUICK, J. Three-dimensional rotations by three shears. *Graphical Models and Image Processing* 59 (1997), 89–96.
- [28] WEBBER, R. E. Ray tracing voxel data via biquadratic local surface interpolation. *The Visual Computer* 6, 1 (Feb. 1990), 8–15.
- [29] WILHELMS, J., AND VAN GELDER, A. A coherent projection approach for direct volume rendering. *Computer Graphics* 25, 4 (July 1991), 275–284.
- [30] WYVILL, B., MCPHEETERS, C., AND WYVILL, G. Animating soft objects. *The Visual Computer* 2, 4 (1986), 235–242.
- [31] YAGEL, R., AND KAUFMAN, A. Template-based volume viewing. *Computer Graphics Forum* 11, 3 (Sept. 1992), 153–167.

- [13] JAIN, A. K. *Fundamentals of Digital Image Processing*. Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [14] LACROUTE, P. Real-time volume rendering on shared memory multiprocessors using the shear-warp factorization. In *Proceedings of the 1995 Parallel Rendering Symposium* (New York, 30–31 Oct. 1995), S. N. Spencer, Ed., ACM Press, pp. 15–22.
- [15] LACROUTE, P. Analysis of a parallel volume rendering system based on the shear-warp factorization. *IEEE Transactions on Visualization and Computer Graphics* 2, 3 (Sept. 1996), 218–231.
- [16] LACROUTE, P., AND LEVOY, M. Fast volume rendering using a shear-warp factorization of the viewing transformation. In *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)* (July 1994), A. Glassner, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH, ACM Press, pp. 451–458. ISBN 0-89791-667-0.
- [17] LAUR, D., AND HANRAHAN, P. Hierarchical splatting: A progressive refinement algorithm for volume rendering. *Computer Graphics (SIGGRAPH '91 Proceedings)* 25, 4 (July 1991), 285–288.
- [18] LEVOY, M. Display of surfaces from volume data. *IEEE Computer Graphics and Applications* 8, 3 (Feb. 1987), 29–37.
- [19] LEVOY, M. Efficient ray tracing of volume data. *ACM Transactions on Graphics* 9, 3 (July 1990), 245–261.
- [20] LEVOY, M. Volume rendering by adaptive refinement. *The Visual Computer* 6, 1 (Feb. 1990), 2–7.
- [21] LITTLEWOOD, D. *A University Algebra*. William Heinemann, Ltd., London, England, 1950.
- [22] LORENSEN, W. E., AND CLINE, H. E. Marching cubes: A high resolution 3D surface construction algorithm. In *SIGGRAPH '87 Conference Proceedings (Anaheim, CA, July 27–31, 1987)* (July 1987), M. C. Stone, Ed., Computer Graphics, Volume 21, Number 4, pp. 163–170.
- [23] NIELSON, G. M., FOLEY, T. A., HAMANN, B., AND LANE, D. A. Visualizing and modeling scattered multivariate data. *IEEE Computer Graphics and Applications* 11, 3 (May 1991), 47–55.
- [24] NIELSON, G. M., AND HAMANN, B. The asymptotic decider: Removing the ambiguity in marching cubes. In *Proceedings of Visualization '91* (1991), pp. 83–91.

References

- [1] BAJAJ, C. L., PASCUCCI, V., AND SCHIKORE, D. R. Fast isocontouring for improved interactivity. In *1996 Volume Visualization Symposium* (Oct. 1996), IEEE, pp. 39–46.
- [2] CATMULL, E., AND ROM, R. A class of local interpolating splines. In *Computer Aided Geometric Design*, R. Barnhill and R. Riesenfeld, Eds. Academic Press, 1974, pp. 317–326.
- [3] CLEARY, J. G., AND WYVILL, G. Analysis of an algorithm for fast ray tracing using uniform space subdivision. *The Visual Computer* 4, 2 (July 1988), 65–83.
- [4] CLINE, H. E., LORENSEN, W. E., LUDKE, S., CRAWFORD, C. R., AND TEETER, B. C. Two algorithms for the reconstruction of surfaces from tomograms. *Medical Physics* 15, 3 (June 1988), 320–327.
- [5] DEROSE, T. D., AND BARSKY, B. A. Geometric continuity and shape parameters for Catmull-Rom splines. In *Graphics Interface '84 Proceedings* (1984), S. MacKay, Ed., pp. 57–64.
- [6] FARIN, G. *Curves and Surfaces for Computer Aided Geometric Design*. fourth edition, Academic Press, Boston, 1997.
- [7] FRANKE, R., AND NIELSON, G. Smooth interpolation of large sets of scattered data. *International Journal for Numerical Methods in Engineering* 15, 11 (1980), 1691–1704.
- [8] FRANKE, R., AND NIELSON, G. M. Scattered data interpolation and applications: A tutorial and survey. In *Focus on Scientific Visualization*, H. Hagen, H. Mueller, and G. Nielson, Eds. Springer-Verlag, New York, 1995, pp. 131–159.
- [9] GLASSNER, A. S. *Uniform Sampling and Reconstruction*. Morgan Kaufmann, San Francisco, CA, 1995, pp. 331–368.
- [10] HAMANN, B., TROTTS, I. J., AND FARIN, G. E. On approximating contours of the piecewise trilinear interpolant using triangular rational-quadratic Bézier patches. *IEEE Transactions on Visualization and Computer Graphics* 3, 3 (July 1997), 215–227.
- [11] HARDY, R. L. Multiquadric equations of topography and other irregular surfaces. *Journal of Geophysical Research* 76 (1971), 1906–1915.
- [12] HARDY, R. L. Theory and applications of the multiquadric-biharmonic method: 20 years of discovery 1968–1988. *Computers and Mathematics with Applications* 19 (1990), 163–208.


the rotation/resampling phase. We shall develop better methods to identify grid cells containing isosurfaces, more efficient ones than our current linear-search scheme. We intend to apply our technique to other volume visualization techniques, *e.g.*, Levoy's method, and possibly to the rendering of vector fields. We also will investigate the elimination of the rotation steps in this algorithm (see Cleary and Wyvill [3]).

If a data set has the same value over large regions, methods to smoothly approximate these regions by high-degree schemes do not improve smoothness: they require many unnecessary coefficients to be stored. Adaptive methods, such as an octree-based approach, could possibly be used to determine regions that hardly vary, and we could use low-degree schemes in such regions. Additional work must be done to improve the ray tracing phase of the algorithm: There is a high degree of correlation among ray-isosurface intersections for adjacent pixels, which we do not consider at this point. We plan to investigate this in the future.


9. Acknowledgments

This work was supported by the National Science Foundation under contract ACI 9624034 (CAREER Award), the Office of Naval Research under contract N00014-97-1-0222, the Army Research Office under contract ARO 36598-MA-RIP, the NASA Ames Research Center through an NRA award under contract NAG2-1216, the Lawrence Livermore National Laboratory through an ASCI ASAP Level-2 under contract (W-7405-ENG-48 and B335358, B347878), and the North Atlantic Treaty Organization (NATO) under contract CRG.971628 awarded to the University of California, Davis. We also acknowledge the support of Silicon Graphics, Inc.

The equine metacarpal dataset was provided by Dr. C.M. Les, of the JD Wheat Veterinary Orthopedic Research Lab at the University of California, Davis. We would like to thank all members of the Visualization Group at the Center for Image Processing and Integrated Computing (CIPIC) at the University of California, Davis, for their help.




(a)




(b)

FIGURE 11: Equine metacarpal data set: (a) using tricubic sampling and trilinear interpolation in the ray tracing step (Note the bands running diagonally through the image); (b) using tricubic sampling and tricubic interpolation in the ray tracing step. The banding has effectively disappeared.



(a)



(b)

FIGURE 12: A second view of the equine metacarpal data set: (a) using tricubic sampling and trilinear interpolation; (b) using tricubic sampling and tricubic interpolation. Again, the banding virtually disappears.




FIGURE 10: The skull data set rendered with tricubic sampling in the rotation step and tricubic interpolation in the ray tracing step. The artifacts are nearly gone.

effect of tricubic sampling in the rotation step and trilinear interpolation in the ray tracing step. Here, most of the artifacts are muted, but still visible. Figure 9 illustrates the effect of Hardy sampling in the rotation step and tricubic interpolation in the ray tracing step. In this image the artifacts due to the algorithm are muted, but the data artifacts look blurred. Figure 10 illustrates the effect of tricubic sampling in the rotation step and tricubic interpolation in the ray tracing step. The artifacts have nearly disappeared.

In animation, these vertical bands appear to move and pulse across the isosurface. In an animation generated with tricubic sampling and tricubic interpolation in the ray tracing step the images appear smooth and very few artifacts are visible.

8. Conclusions

We have discussed a new method for rendering smooth isosurfaces of trivariate scalar fields using C^1 -continuous spline functions with an efficient ray-tracing scheme. We have shown how a rotation-resampling approach can be used to reduce the complexity of ray-isosurface intersection calculations.

We plan to develop error metrics to characterize and quantify the error introduced during




FIGURE 8: The skull data set rendered with Hardy sampling in the rotation step and tricubic interpolation in the ray tracing step. Note that the artifacts are reduced, but they are slightly blurred.




FIGURE 9: The skull data set rendered with tricubic sampling in the rotation step and linear interpolation in the ray tracing step. One can see a substantial reduction in the artifacts generated by the algorithm.

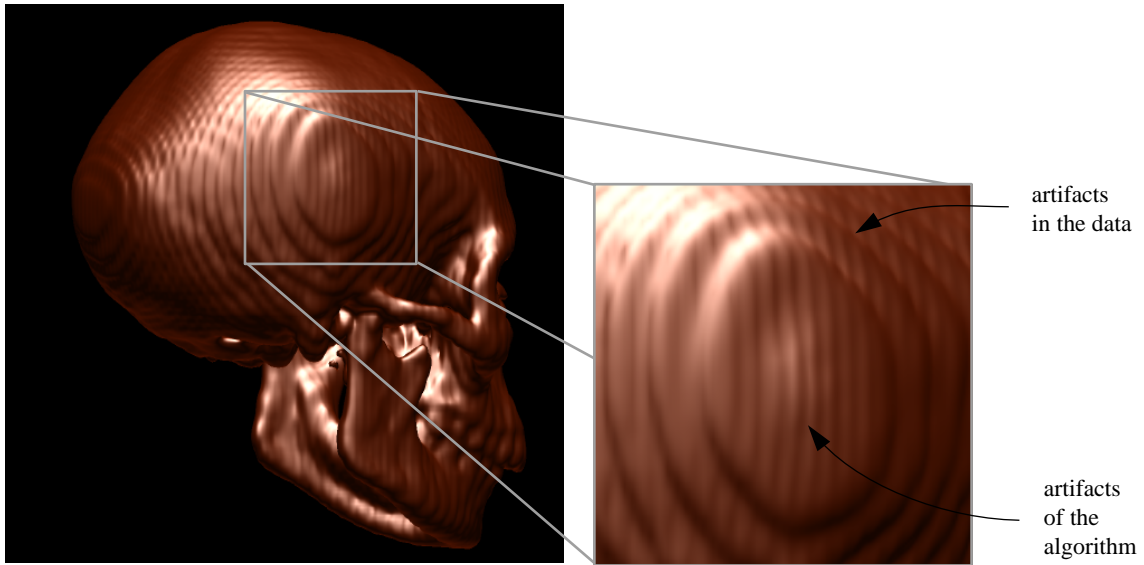


FIGURE 6: The artifacts generated from the algorithm and from the data in the skull data set. The artifacts inherent to the data are the ring patterns, while the artifacts from the algorithm are the vertical patterns in the image.




FIGURE 7: The skull data set rendered with trilinear sampling in the rotation step and tricubic interpolation in the ray tracing step. Notice that most of the artifacts remain in the image.

each cell and determine whether the isosurface exists within the cell. If the interval bounds the isovalue, trace the ray through the grid cell and determine the possible intersections with the isosurface. If the ray intersects the isosurface, then we compute the intersection closest to the image plane and shade the surface using the gradient of the isosurface. Ray tracing terminates when all pixels have been processed.

We have implemented our algorithm on several Silicon Graphics workstations, including systems with multiple processors. Our implementation takes advantage of multiple processors by executing all compute-intensive activities in parallel. We have found a near-linear speedup between the implementation on an SGI O2 (single-processor system) and an Origin2000 (16-processor system).

The mapping of source grid vertices to target grid vertices is done in parallel. The mapping of vertices is structured such that the size of the needed memory stride is kept at a minimum. Each processor runs to completion, and waits until the other processors are done. We have experienced that the difference in completion time, considering all processors, is usually less than ten percent.

The ray-isosurface intersection calculations are parallelizable as well. Since rays may pass through different numbers of grid cells and different types of computations may be performed, we use a more sophisticated load-balancing scheme: We use a job queue that consists of scan line indices. A processor reads a scan line index from the job queue and casts rays for all the pixels in the scan line. The time needed to acquire and update the job queue is very small compared to line scanning.

7. Results

Figure 6 shows the data set and artifacts that brought this research problem to our attention. This data set is the result of a computerized-axial-tomography (CAT) scan of a skull. The size of the data set is $64 \times 64 \times 68$, and its small resolution produces artifacts in the renderings. These artifacts are the “circular” patterns shown in Figure 6. However, when we rendered the skull using conventional techniques, we obtained artifacts that were not inherent to the data. These are also visible in Figure 6 as the vertical “stripe” pattern. When animated, the artifacts inherent to the data set stay fixed, but the stripe artifacts ripple and move.

Figures 7-10 illustrate our methods as applied to this data set. Figure 7 illustrates the effect of trilinear sampling in the rotation step and tricubic interpolation in the ray tracing step. We note that most of the artifacts due to the algorithm remain in the image. Figure 8 illustrates the

- If $D < 0$, then the equation has three different real roots. The roots are given by

$$s_1 = k \cos \frac{\alpha}{3}, s_2 = k \cos \frac{2\pi + \alpha}{3}, \text{ and } s_3 = k \cos \frac{4\pi + \alpha}{3},$$

where

$$k = \sqrt{-\frac{4p}{3}} \text{ and } \cos \alpha = \frac{-4q}{k^3}.$$

- If $D = 0$, then the equation has three real roots, of which two are equal. The roots are given by

$$s_1 = \sqrt[3]{\frac{q}{2}} \text{ and } s_2 = -2\sqrt[3]{\frac{q}{2}},$$

where s_1 is the double root.

To accelerate the ray-isosurface intersection computations, we store an interval with each grid cell in T , which represents a lower and upper bound of the function values within the cell. The Catmull-Rom interpolants are polynomial splines, and we find the lower and upper bound by determining minimal and maximal spline coefficients of the corresponding Bernstein-Bézier representation (see Section 3). If, for a given grid cell, the desired isosurface value is outside the interval associated with this cell, then the cell cannot contain any part of the isosurface.

We also store intervals for each “column of grid cells”.² Thus, a ray can intersect an isosurface in a grid cell of a particular column only if the column-specific interval contains the particular isovalue. The calculation of lower and upper bounds turns out to be beneficial for “sparse data sets,” *i.e.*, data sets containing large homogeneous areas, or thin features.

6. Implementation

The algorithm requires an initial grid S and an isovalue \bar{f} , and it outputs an image of the isosurface $f = \bar{f}$. To reduce sampling artifacts, we use a resolution for the grid T that is at least the twice the resolution of the S grid (Nyquist limit) in each direction [9, 13]. We estimate function values in T from those in S by using one of the sampling methods discussed in Section 4. We compute intervals for each grid cell in T , and intervals for each column of grid cells in T .

For each pixel in the image plane, we cast a ray $\mathbf{r}(t)$ into the grid T . If the column-specific interval of the column of cells intersected by $\mathbf{r}(t)$ does not contain the specific isovalue, we skip to the next pixel. Otherwise, for each grid cell intersected by the ray, examine the intervals for

²A column of grid cells in T consists of all the cells of T intersected by an axis-aligned ray.

If \mathbf{p} is a vertex of the grid T , we consider the 64 data sites given by the vertices of the “stencil” in S of the cell containing \mathbf{p} (see Figure 5). Assuming that the spacing of the S grid is uniform, the distances $\|\mathbf{p}_j - \mathbf{p}_i\|$ are the same throughout the grid and one needs to invert the 64×64 matrix only once. Hardy’s method is considerably slower than the tricubic spline approximation due to the computation of the square roots. The visual results of the two approximation schemes are similar.

5. Ray-Isosurface Intersection Tests

To intersect the parametric ray $\mathbf{r}(t) = (x_0, y_0, z_0 + tz_d)$ with the isosurface

$$f(x, y, z) = \sum_{k=0}^3 \sum_{j=0}^3 \sum_{i=0}^3 c_{i,j,k} x^i y^j z^k = 0 \quad (4)$$

we substitute the coordinate functions for x , y , and z into equation (4) leading to the (normalized) cubic equation

$$t^3 + at^2 + bt + c = 0. \quad (5)$$

To find the roots of this equation, we use Cardan’s solution [21], which first transforms equation (5) by substituting $s = t - \frac{a}{3}$ leading to the polynomial

$$s^3 + ps + q, \quad (6)$$

where

$$p = -\frac{a^3}{3} + b \quad \text{and} \quad q = 2\left(\frac{a}{3}\right)^3 - \frac{ab}{3} + c.$$

Equation (6) has three roots. The characteristics of these roots are determined by the discriminant of the cubic equation

$$D = 4p^3 + 27q^2,$$

which leads to three cases:

- If $D > 0$, then the cubic equation has one real root and two complex roots. The single real root is given by

$$s_1 = A + B,$$

where

$$A = \sqrt[3]{\frac{q}{2} + \sqrt{\frac{D}{27}}} \quad \text{and} \quad B = \sqrt[3]{\frac{q}{2} - \sqrt{\frac{D}{27}}}.$$

the cell in S containing \mathbf{p} are

$$\begin{aligned} u &= x_p - \lfloor x_p \rfloor, \\ v &= y_p - \lfloor y_p \rfloor, \text{ and} \\ w &= z_p - \lfloor z_p \rfloor. \end{aligned} \tag{3}$$

Thus, the function value $f(\mathbf{p})$ is

$$f(\mathbf{p}) = \sum_{k=0}^1 \sum_{j=0}^1 \sum_{i=0}^1 f_{\lfloor x_p \rfloor + i, \lfloor y_p \rfloor + j, \lfloor z_p \rfloor + k} L_i(u) L_j(v) L_k(w),$$

where the Lagrange polynomials are $L_0(t) = 1 - t$ and $L_1(t) = t$.

We can extend the trilinear spline approximation scheme to tricubic spline approximation by using the Catmull-Rom spline. In this case the function value for a vertex \mathbf{p} of T is approximated by considering the 64 vertices of the “stencil cell” in S associated with the cell containing \mathbf{p} . If u , v and w denote the local parameter values of \mathbf{p} (equation (3)), then the function value $f(\mathbf{p})$ is

$$f(\mathbf{p}) = \sum_{k=0}^3 \sum_{j=0}^3 \sum_{i=0}^3 f_{\lfloor x_p \rfloor + i, \lfloor y_p \rfloor + j, \lfloor z_p \rfloor + k} C_i(u) C_j(v) C_k(w),$$

where $C_i(t)$ denotes the cubic Catmull-Rom blending functions.

Hardy’s multiquadric method [8, 11, 12, 23] is a standard scattered data approximation scheme. It is used in cases where the data sites are randomly distributed and (typically) no connectivity information is known for the data sites. Hardy’s approximant smoothly interpolates the values at the data sites and we use it for estimating function values for the vertices of the T grid.

Given a vertex \mathbf{p} in the T grid and n vertices $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_n$ of S , with associated function values f_0, f_1, \dots, f_n , Hardy’s multi-quadric interpolant is defined by the linear system

$$f_j = \sum_{i=0}^n \alpha_i \sqrt{(\|\mathbf{p}_j - \mathbf{p}_i\|)^2 + R^2}, \quad j = 0, \dots, n,$$

and the estimate at \mathbf{p} is

$$f(\mathbf{p}) = \sum_{i=0}^n \alpha_i \sqrt{(\|\mathbf{p} - \mathbf{p}_i\|)^2 + R^2},$$

where R^2 is some positive constant. The computation of the coefficients α_i requires the inversion of an $(n + 1) \times (n + 1)$ matrix.




FIGURE 5: A two-dimensional illustration of resampling in three dimensions: Given a vertex \mathbf{p} of the grid T , trilinear approximation uses the eight vertices of the three-dimensional grid cell in S containing \mathbf{p} to approximate a function value at \mathbf{p} ; tricubic spline approximation and Hardy's method use the 64 vertices of S in the stencil of the cell containing \mathbf{p} to approximate a function value.

The rotation of S into T is effectively performed by resampling the scalar field approximation at the vertices of T , *i.e.*, we must approximate new function values for the vertices of each cell in T from the original values at the vertices in S . These resampling problems arise in many visualization applications and they can be resolved in a number of ways [7, 15, 16, 23, 31]. We compare three different approximation techniques for the estimation of function values at the vertices of T : (1) trilinear approximation ; (2) tricubic spline approximation using the Catmull-Rom spline; and (3) Hardy's multiquadric method [11, 12].

Trilinear interpolation is the standard technique used to estimate function values for rectilinear data sets. In this method the function value for a vertex \mathbf{p} of T is approximated by trilinearly interpolating the eight values at the vertices of the cell in S that contains \mathbf{p} (see Figure 5). Denoting the coordinates of \mathbf{p} in S by x_p , y_p , and z_p , \mathbf{p} 's local parameter values with respect to




FIGURE 3: Casting rays through the grid S : ray $\mathbf{r}(t)$ passes through a pixel center in the image plane I and intersects the arbitrarily oriented grid cells of S .




FIGURE 4: Definition of the grid T : One family of grid planes of T is parallel to the image space plane I . Ray $\mathbf{r}(t)$ is axis-aligned with T . After resampling at the vertices of T , ray tracing can be performed with axis-aligned rays.

with the isosurface $f = \bar{f}$. To intersect $\mathbf{r}(t)$ with the isosurface, we substitute the x , y and z values of equation (2) into equation (1) and obtain, in general, a polynomial of degree nine in t . We reduce this problem to solving a third-degree equation by using rays that are aligned with the z axis, *e.g.*, $x_d = y_d = 0$.

Our algorithm is an image-space algorithm, *i.e.*, we cast rays through each pixel in the image plane I and find intersections with an isosurface. A ray, perpendicular to the image plane I , is cast through the source grid S (see Figure 3). We replace the grid S by a second rectilinear grid T , a grid with one family of grid planes parallel to the image plane. Rotating grid S into grid T allows us to perform ray-isosurface computations with axis-aligned rays (see Figure 4).

Bézier “control coefficients” are given by

$$\begin{bmatrix} b_{k-1} \\ b_k \\ b_{k+1} \\ b_{k+2} \end{bmatrix} = M_B^{-1} M_C \begin{bmatrix} f_{k-1} \\ f_k \\ f_{k+1} \\ f_{k+2} \end{bmatrix}.$$

We define I as the smallest interval containing all four values b_{k-1} , b_k , b_{k+1} , and b_{k+2} . The convex-hull property ensures that the curve $f(x)$ is contained within the interval I .

The univariate Catmull-Rom spline can be extended to the trivariate case by constructing a tensor-product spline. In the trivariate case, an individual tricubic segment is given by

$$f(x, y, z) = \sum_{k=0}^3 \sum_{j=0}^3 \sum_{i=0}^3 f_{\lfloor x \rfloor + i - 1, \lfloor y \rfloor + j - 1, \lfloor z \rfloor + k - 1} C_i(x - \lfloor x \rfloor) C_j(y - \lfloor y \rfloor) C_k(z - \lfloor z \rfloor).$$

We use this spline scheme to represent a piecewise tricubic approximation of a scalar field. We must consider a stencil of 64 data values to define the Catmull-Rom spline for a particular grid cell. These 64 data values include the scalar values at the corners of a grid cell C , and the scalar values at the corners of the 26 neighboring cells sharing at least one vertex with C . The 64 values are the set

$$\{f_{i,j,k} : i \in \{i_0 - 1, i_0 + 2\}, j \in \{j_0 - 1, j_0 + 2\}, k \in \{k_0 - 1, k_0 + 2\}\},$$

where f_{i_0, j_0, k_0} is the value at the vertex of C with minimum i, j, k indices. Again, it is possible to determine the range of function values for a particular tricubic polynomial by converting it to Bernstein-Bézier form and computing the minimal and maximal coefficient.

4. Rotation and Resampling

We use ray tracing as the mechanism for finding and rendering an isosurface of a trivariate scalar function. Suppose we are to render the isosurface defined by the expression

$$f(x, y, z) = \sum_{k=0}^3 \sum_{j=0}^3 \sum_{i=0}^3 c_{i,j,k} x^i y^j z^k = \bar{f} \quad (1)$$

for a particular voxel. We must intersect the ray

$$\mathbf{r}(t) = \mathbf{p}_0 + t\vec{\mathbf{d}} = \begin{pmatrix} x_0 + tx_d \\ y_0 + ty_d \\ z_0 + tz_d \end{pmatrix} \quad (2)$$




FIGURE 2: A Catmull-Rom spline: The curve interpolates the control points at integer values of the parameter; the slope of the curve at the integer knots is given by the central difference of the two values on either side of each knot.

where

$$\begin{bmatrix} C_0(x) & C_1(x) & C_2(x) & C_3(x) \end{bmatrix} = \begin{bmatrix} 1 & x & x^2 & x^3 \end{bmatrix} M_C,$$

and the functions $C_j(x)$ denote the Catmull-Rom basis functions.

We can bound the values of a Catmull-Rom spline for a particular knot interval by converting it to Bernstein-Bézier form. This conversion is given by

$$f(x) = \begin{bmatrix} B_0(x) & B_1(x) & B_2(x) & B_3(x) \end{bmatrix} M_B^{-1} M_C \begin{bmatrix} f_{k-1} \\ f_k \\ f_{k+1} \\ f_{k+2} \end{bmatrix},$$

where

$$M_B = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -3 & 3 & 0 & 0 \\ 3 & 6 & 3 & 0 \\ -1 & 3 & -3 & 1 \end{bmatrix}.$$

The matrix M_B converts from the cubic power basis to the Bernstein basis $B_j(x)$ [6]. The convex-hull property of Bézier curves ensures that each segment of the spline lies between the minimal and maximal values of its four defining Bernstein-Bézier control points. The Bernstein-

Given a set of scalar values f_0, f_1, \dots, f_n , the univariate cubic Catmull-Rom spline is a piecewise cubic polynomial function

$$f(x) = \sum_{i=0}^n f_i(x)F_i(x)$$

where each $f_i(x)$ is a function of the values f_i , and the $F_i(x)$ denotes a set of blending functions [2]. The functions $f_i(x)$ depend on the f_i values but are allowed to vary with x . Catmull and Rom discovered that certain choices of the functions $f_i(x)$ allowed the curve to interpolate the control points f_i .

For our application, it is sufficient to assume that the knots of the spline curve are the integers, and, by utilizing the basis functions $F_i(x) = (x - [x])^i$, we can obtain an interpolating spline segment (see the discussing in [2] by setting $f_j(x) = 0$ for $j < [x] - 1$ and for $j > [x] + 2$, and

$$\begin{bmatrix} f_{k-1}(x) \\ f_k(x) \\ f_{k+1}(x) \\ f_{k+2}(x) \end{bmatrix} = M_C \begin{bmatrix} f_{k-1} \\ f_k \\ f_{k+1} \\ f_{k+2} \end{bmatrix}$$

where $k = [x]^1$, and

$$M_C = \frac{1}{2} \begin{bmatrix} 0 & 2 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ 2 & -5 & 4 & -1 \\ -1 & 3 & -3 & 1 \end{bmatrix}.$$

Thus one can write $f(x)$ as

$$f(x) = \begin{bmatrix} 1 & x - k & (x - k)^2 & (x - k)^3 \end{bmatrix} M_C \begin{bmatrix} f_{k-1} \\ f_k \\ f_{k+1} \\ f_{k+2} \end{bmatrix},$$

where $k = [x]$. For $x = k$ the spline interpolates the point f_k , and the derivative at $x = k$ is $\frac{1}{2}(f_{k+1} - f_{k-1})$. Figure 2 illustrates a Catmull-Rom spline scheme.

A segment of the Catmull-Rom spline can also be written as

$$f(x) = \sum_{i=0}^3 f_{[x]+i-1} C_i(x - [x]),$$

¹ $[x]$ denotes the floor of x , the greatest integer less than x .

size regions in a three-dimensional data set where the (scalar) function values are in a certain range, one can assign higher weights to data in these regions by increasing their opacity values. Furthermore, one can use the gradient to emphasize boundaries between different “tissue types”: High gradients imply such boundaries and can therefore be used to increase the opacity of boundary regions. Levoy describes an efficiency improvement of his algorithm in [19, 20].

Several volume visualization algorithms that utilize and render the data directly are based on transforming a rectilinear data set so that it coincides with an “image-space cube” [14, 15, 25]. These methods factor the viewing matrix into a three-dimensional shear and a two-dimensional “warp,” thus projecting the given data set into image space. The shear-warp algorithm is a re-sampling scheme which is advantageous for algorithms utilizing SIMD architectures [27].

Yagel and Kaufman [31] present an algorithm for volume rendering that is based on exploiting coherency between rays in parallel projections. Rays are cast from a base plane, which simplifies the paths through the volume in a serial image-order algorithm, insuring uniform sampling of the volume.

Webber [28] describes an algorithm that casts rays through the grid directly. This algorithm constructs a biquadratic isosurface over each grid cell, and intersects the ray with this surface. This biquadratic surface is derived by examining the $3 \times 3 \times 3$ volume of voxels around the target voxel and using one of the 3×3 sides of this voxel array as a base for a biquadratic function. This approximating function is used for ray-isosurface intersection tests. This method allows the derivation of a precise intersection and a precise normal over a voxel.

Our approach combines ray-casting methods with a tricubic spline interpolation scheme for high-quality image generation and animation. To eliminate the root-finding for ninth-degree polynomials, we rotate and resample the original rectilinear data set at the vertices of a grid with one family of grid planes parallel to the image plane. We cast rays that are normal to the image plane and thus perform ray-intersection calculations for third-degree polynomials. These methods are all independent calculations and can be efficiently implemented on a parallel processing system.

3. Tricubic Interpolation

To solve the general ray-isosurface intersection problem, we must compute the intersection of a (parametric) ray and an isosurface of a tricubic approximant. We utilize cubic tensor-product Catmull-Rom splines [2, 6] to represent the scalar field for all grid cells. We briefly review this spline scheme for the univariate case.

of its coordinate “directions” is parallel to the ray direction, and by resampling the field at the vertices of the “ray-aligned” grid. Thus, each intersection calculation is reduced to root-finding for a third-degree polynomial.

In Section 2, we review research related to this problem. In Section 3, we review the tricubic spline that we use to approximate a three-dimensional field. In Section 4, we discuss the rotation and resampling steps that allow us to efficiently render smooth isosurfaces by using axis-aligned rays. The ray-isosurface intersection method is discussed in Section 5. In Section 6, we discuss ray-isosurface intersection and implementation details of the algorithm. We present results of our approach in Section 7 and conclude with an outlook on future work in Section 8.

2. Related Work

Two basic approaches are used in scientific visualization for the identification, extraction, and rendering of isosurfaces: those that generate an intermediate representation of the isosurfaces for rendering purposes and those that render the data directly.

The marching-cubes algorithm [30, 22, 24] generates an intermediate triangle mesh from a set of trilinear interpolants. Considering each trilinear interpolant independently, the algorithm determines a set of triangles that approximates the isosurface within each voxel. The resulting triangle mesh is displayed using conventional rendering hardware. The number of cases to be considered for generating the mesh within a voxel can be simplified by symmetry and rotation to 14 cases, and the algorithm uses a look-up table to quickly generate the triangular mesh approximating an isosurface. The marching-cubes method has been adapted to tetrahedral meshes by Shekhar *et al.* [26], where only three cases must be considered.

Hamann *et al.* describe a method for constructing a “smoother” isosurface by fitting rational-quadratic surface patches to the triangle mesh resulting from the marching-cubes algorithm [10]. Since the contour of the bilinear interpolant on a face of a grid cell is a hyperbolic arc, rational quadratic curves can represent this curve exactly. These boundary curves, together with boundary curves in the interior of a grid cell, can be used to generate the control nets of triangular rational-quadratic patches that approximate the isosurface within the cell. The use of this type of surface patch yields a “smoother” isosurface, but tangent plane discontinuities still exist at the boundaries of each grid cell.

Levoy [18] introduced the concept of casting rays directly through volumetric data to obtain transparent images. His algorithm uses trilinear spline interpolation along with the spline gradient to assign normal vectors points along the rays for shading purposes. In order to empha-

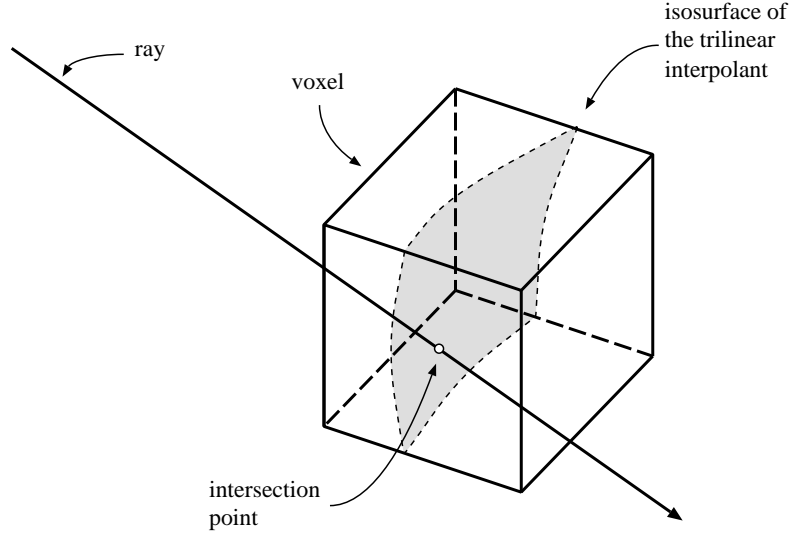


FIGURE 1: Ray-isosurface intersection in a voxel: The shaded area shows the isosurface of the trilinear approximant. The bounding curves of the isosurface on the faces of the voxel are hyperbolic arcs.

on an isosurface we insert a parametric ray equation

$$\mathbf{r}(t) = \mathbf{p}_0 + t\vec{\mathbf{d}} = \begin{pmatrix} x_0 + tx_d \\ y_0 + ty_d \\ z_0 + tz_d \end{pmatrix}$$

into the individual tricubic polynomial segments characterizing field over each voxel, *i.e.*,

$$f(x, y, z) = \sum_{k=0}^3 \sum_{j=0}^3 \sum_{i=0}^3 c_{i,j,k} x^i y^j z^k.$$

Thus, we have to solve the equation

$$\sum_{k=0}^3 \sum_{j=0}^3 \sum_{i=0}^3 (x_0 + tx_d)^i (y_0 + ty_d)^j (z_0 + tz_d)^k = \bar{f}$$

to determine points on the isosurface $f = \bar{f}$. This is an equation of degree nine in t , and so the intersection of a ray with the isosurface of a tricubic polynomial requires the solution of a ninth-degree polynomial – a difficult task.

However, if the direction vector $\vec{\mathbf{d}}$ of the ray \mathbf{r} is parallel to one of the axes of the coordinate system in which $f(x, y, z)$ is defined, then two of the values of x_d , y_d , and z_d must be zero, and the intersection problem leads to a (solvable) third-degree equation. Thus, we reduce the general ray-isosurface intersection problem by rotating the given rectilinear grid such that one

1. Introduction

With ever increasing speed, data sets are being generated that represent three-dimensional, volumetric information. They arise as a result of complex computational simulations or of empirical data-collection procedures, and present a challenge to current visualization techniques. As these are large data sets, the approximation techniques used to render the data produce artifacts that are visible in the resulting images, and especially in the animations. These artifacts frequently present misleading information about the data and may reduce the reliability of the visualization. It is crucial to clearly display all information contained in a data set, and not to simplify the data for the sake of rendering ease.

Scientific data sets are typically scalar- or vector-valued and defined on a grid in three-dimensional space. In many cases, the grid is rectilinear and the data represents either point values associated with the nodes (or vertices) or constant values associated with the voxels (or grid cells). Three methods have been developed to visualize these data sets: the construction of isosurfaces, *i.e.*, surfaces defining the region in space for which a particular field variable is constant [18, 22, 24]; volume visualization methods, where the algorithm casts rays through the volume, averaging color and opacity values along the ray [18]; and splatting, which is based on projecting the voxel data onto the image plane and considering relative depth to simulate transparency [4, 17, 29].

We describe a method based on a C^1 -continuous interpolation scheme that leads to normal-continuous isosurfaces. We construct a cubic tensor product spline interpolating the function values at the grid points ensuring an overall C^1 -continuous representation of the field. The algorithm extends the de facto standard of piecewise-trilinear interpolation of the data in a cell to piecewise-tricubic interpolation.

Piecewise trilinear interpolation (Figure 1) of rectilinear data interpolates the eight data values at the corners of a cell and yields an overall C^0 -continuous approximation of a field. The resulting isosurfaces have normal discontinuities on the shared boundary faces of each pair of neighboring voxels. The resulting normal discontinuities are clearly visible when rendering isosurfaces. Images obtained by ray casting also suffer from the lack of smoothness in the trilinear spline interpolation. The resulting artifacts – “ring patterns” – are particularly disturbing in animations of a data set.

Piecewise tricubic interpolation utilizes the 64 vertices at the corners of a $3 \times 3 \times 3$ array of cells, and calculates a tricubic approximant to the values within the cell. To determine a point

HIGH-QUALITY RENDERING OF SMOOTH ISOSURFACES

Eric LaMar¹,
Bernd Hamann²,
Kenneth I. Joy³

Center for Image Processing and Integrated Computing (CIPIIC)
Department of Computer Science
University of California
Davis, CA 95616-8562

Abstract

Animation and visualization of rectilinear data require interpolation schemes for smooth image generation. Piecewise trilinear interpolation, the de facto standard for interpolating rectilinear data, usually leads to significant visual artifacts in the resulting imagery. These artifacts reduce the confidence in the resulting visualization, and may even lead to false interpretations of the data. This paper is concerned with the generation of smooth isosurface image sequences, obtained by casting rays through the image plane and computing their intersections with an isosurface. We describe a novel solution to this problem: We replace trilinear interpolation by tricubic interpolation, smoothing out the artifacts in the images; and we simplify the ray-isosurface intersection calculations by rotating and resampling the original rectilinear data in a second rectilinear grid – a grid with one family of grid planes parallel to the image plane. Our solution significantly reduces artifacts in individual images and leads to smooth animations.

Keywords: animation; approximation; trilinear splines; volume visualization; ray casting; isosurfaces; resampling; Catmull-Rom splines; Hardy's multiquadric scheme.

¹Department of Computer Science, University of California, Davis, CA 95616-8562, USA; e-mail: lamar@cs.ucdavis.edu

²Co-Director of CIPIIC, Department of Computer Science, University of California, Davis, CA 95616-8562, USA; e-mail: hamann@cs.ucdavis.edu

³Corresponding Author; Department of Computer Science, University of California, Davis, CA 95616-8562, USA; e-mail: joy@cs.ucdavis.edu