Speed-up of Volume Rendering by Adaptive Distance Sampling

H. Chen, J. Hesser, R. Männer

Lehrstuhl für Informatik V, University of Mannheim

This paper presents an algorithm to speed up volume rendering by adaptive sampling distance coding. Our algorithm extends current approaches by not only considering empty and homogenous regions, but also by reducing the sampling frequency where the object opacity varies slowly. This is numerically realized by a replacement of the rectangle rule for integrating the voxel contributions along a ray by a numerical integration of piecewise linear opacity changes.

The algorithm consists of two stages: sampling distance coding as preprocessing and ray casting. Theoretical investigations show that the overall complexity of volume rendering can be reduced to $O(N^2)$ for any choice of the opacity function. Implementations of this algorithm show that the speed-up for several typical examples is in the range of 1-6.3. No speedup occurs only in the cases where algorithmic approaches based on space-leaping and early-ray-termination are efficient, i.e. where the object has hard surfaces and the remaining volume is empty. The highest speedup is observed for semi-transparent objects. In such regions the average sampling distance for voxels is extended from 0.5 to 5.2 units of the grid size.

To determine the sampling interval for the semi-transparent voxels we differentiate the results from two different quadrature methods. Provided that the difference between the two quadratures is less than a given value, the maximal integration interval will be the distance from the current sample point to the next point. While one quadrature method estimates the integral with the rectangle rule using evenly spaced sample points along the sample interval (in voxel distance), the second method estimates the integral with the trapezoid rule using only the two end points of the sampling interval.

This precalculation, programmed in a straightforward way, needs 15 hours computation time on a PIII 550. By using local properties of the distance coding criterion we could reduce this time to 36 seconds for a volume of 256×256×123 voxels.

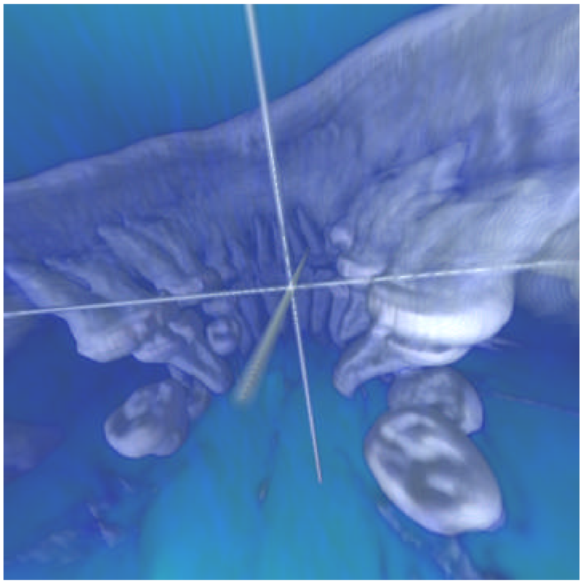

Fig 1.

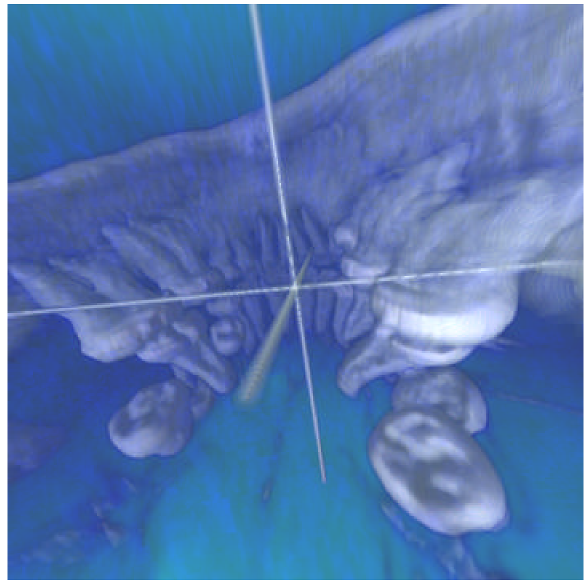

Fig 2.

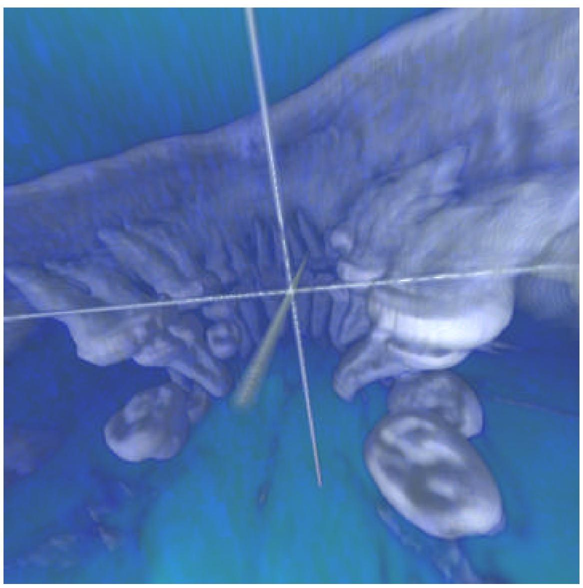

Fig 3.

Fig 1 is a reference image accelerated by space leaping and early ray termination; Fig 2 additionally accelerated by adaptive sampling with exhaustive precalculation; Fig 3 is the same as Fig 2 but using optimized precalculation. The lines in the volume are one voxel wide, the bones are surrounded by semi-transparent soft tissue. CT artifacts from the dataset are visible.