

Volume Rendering

Murad Aliyev
801 Leroy Pl., POBox 2644,
Socorro, NM 87801
maliyev@nmt.edu

ABSTRACT

This paper shows topical research on volume rendering and volume rendering techniques. Different areas of applications are presented and related problems are discussed.

General Terms

Measurement, Performance, Experimentation.

Keywords

Volume rendering, direct volume rendering.

1. INTRODUCTION

Visualization of three-dimensional scalar fields, commonly called volume rendering, is an area that have been developed since 1980s. The nature of data is that the interior of visualized object can be rendered as well as the surface of it.

Volume rendering is a technique used to display a 2D projection of a 3D discretely sampled data set. A typical 3D data set is a group of 2D slice images acquired by a scanner. Usually these are acquired in a regular pattern (e.g., one slice every millimeter) and usually have a regular number of image pixels in a regular pattern. This is an example of a regular volumetric grid, with each volume element, or voxel represented by a single value that is obtained by sampling the immediate area surrounding the voxel.

Direct volume rendering is a computationally intensive operation that has become a valued and often preferred visualization tool. Direct volume rendering is the common name that describes the viewing of volume data as a semi transparent cloudy material. Its advantages are that much or all of the volume may be visible to the observer at one time; there is no need to introduce intermediate geometry that doesn't really exist in the data. We assume the input data is a scalar field sampled at the vertices of a 3D rectilinear lattice – a situation often encountered in medical and simulation data.

The challenges of volumetric visualization arise from the size of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Visualization, 2007, Socorro, NM.

Copyright 2007 NMT Computer Science Department ...\$5.00.

volume data as well as the complexity of the data. One of the key goals of volume graphics is the interactive exploration of volumetric data. Exploration of data includes two important aspects — the change of the viewpoint as well as the change of the visible parts of the data (manipulated through the transfer function.)

Much research has focused on the quick rendering of the data in order to enable (real-time) interaction and hence to improve the 3D understanding of the data. The interactive manipulation of the viewpoint allows us to recover a 3D understanding from a series of 2D screen projections making use of the *kinetic depth effect*. Another important aspect of volume data exploration is the ability to manipulate color and opacity of voxels to adjust their visibility and appearance.

2. VOLUME RENDERING TECHNIQUES

Volume data is in its simplest form a quadruple (x,y,z,w) where x, y, z is a scalar describing a position in three-dimensional space and w is the actual data value. This quadruple is often called a voxel. A volume, or dataset as it is sometimes called, is hence a collection of volume data generated from some sort of measurement apparatus or simulation.

A voxel can have an arbitrary number of data values associated with it. In this case, the dataset is said to be multivariate when more than one data value is used. Multivariate data is common with data flow simulations or when the result of a number of data sources is combined.

Data structure is categorized as:

- Cartesian: The data elements are structured, and arranged in a cubic grid
- Regular: As Cartesian, but the grid can be rectangular as well.
- Rectilinear: As Regular, with the acceptance of non-uniformly sized grids
- Structured / Curvilinear: Non-linear grids.
- Block structured: Several structured systems are allowed.
- Unstructured: Unstructured data, data points can be anywhere within the volume.
- Hybrid: A combination of unstructured and structured data.

The most important distinction between structured and unstructured data is that they often require different visualization methods. For example, shear-warp is a common technique used

for structured data, while projected tetrahedras is used for unstructured data.

Direct Volume Rendering is a volume rendering method that uses the raw data directly as it is acquired, without any intermediate states, like surface extraction, so we are continuously working on original scans. A direct volume renderer requires every sample value to be mapped to opacity and a color. This is done with a “transfer function” which can be a simple ramp, a piecewise linear function or an arbitrary table. Once converted to an RGBA (for red, green, blue, alpha) value, the composed RGBA result is projected on correspondent pixel of the frame buffer. The way this is done depends on the rendering technique.

A combination of these techniques is possible. For instance, a shear warp implementation could use texturing hardware to draw the aligned slices in the off-screen buffer. There are several DVR methods, quite similar in general, all involving some way of compositing the voxels.

2.1 Volume Ray Casting

The simplest way to project the image is to cast rays through the volume using ray casting. In this technique, a ray is generated for each desired image pixel. Using a simple camera model, the ray starts at the center of the projection of the camera (usually the eye point) and passes through the image pixel on the imaginary image plane floating in between the camera and the volume to be rendered.

The ray is clipped by the boundaries of the volume to save time then the ray is sampled at regular intervals throughout the volume. The data is interpolated at each sample point, the transfer function applied to form an RGBA sample, the sample is composited onto the accumulated RGBA of the ray, and the process repeated until the ray exits the volume. The RGBA color is converted to an RGB color and deposited in the corresponding image pixel. The process is repeated for every pixel on the screen to form the completed image.

In its basic form, the volume ray casting algorithm comprises four steps:

1. **Ray casting.** For each pixel of the final image, a ray of sight is shot (“cast”) through the volume. At this stage it is useful to consider the volume being touched and enclosed within a bounding primitive, a simple geometric object — usually a cuboid — that is used to intersect the ray of sight and the volume.
2. **Sampling.** Along the part of the ray of sight that lies within the volume, equidistant sampling points or samples are selected. As in general the volume is not aligned with the ray of sight, sampling points usually will be located in between voxels. Because of that, it is necessary to trilinearly interpolate the values of the samples from its surrounding voxels.
3. **Shading.** For each sampling point, the gradient is computed. These represent the orientation of local surfaces within the volume. The samples are then shaded, i.e. colored and lighted,

according to their surface orientation and the source of light in the scene.

4. **Compositing.** After all sampling points have been shaded; they are composited along the ray of sight, resulting in the final colour value for the pixel that is currently being processed. The composition is derived directly from the rendering equation and is similar to blending acetate sheets on an overhead projector. It works back-to-front, i. e. computation starts with the sample farthest from the viewer and ends with the one nearest to him. This work flow direction ensures that masked parts of the volume do not affect the resulting pixel.




Figure 1. Four basic steps of Volume Ray Casting. 1 - Ray casting, 2 - Sampling, 3 - Shading, 4 - Compositing. http://upload.wikimedia.org/wikipedia/commons/e/ec/Volume_ray_casting.png

2.2 Splatting

This is a technique which trades quality for speed. Here, every volume element is splatted (like snow balls) on to the viewing surface from in back to front order. These splats are rendered as disks whose properties (color and transparency) vary diametrically in normal (Gaussian) manner. Flat disks and those with other kinds of property distribution are also used depending on the application.

2.3 Shear Warp

The shear-warp technique was introduced in 1994 and is a mixture of several prior methods and it offers excellent speed without sacrificing the picture quality. It is still today regarded as the fastest software rendering method and it is the basis for many hardware rendering methods.

The shear-warp method tries to combine the picture quality of image-based techniques such as raycasting with the efficiency of object-order techniques. The main disadvantage of raycasting is that a lot of computational time is used for tracing the viewing rays through the volume and the interpolation of voxels. Shear-warp is based on the idea that if raycasting is always done perpendicularly to the volume, then the data can potentially be traversed much more efficiently by exploiting how data in rectilinear grids are stored. Likewise can interpolation be made easier compared to when the rays are cast into the volume at an

angle. This is the core idea of shear-warp rendering. The technique assumes that the volume data is arranged in rectilinear grids, and consequently, it will not work on unstructured data.

2.4 Texture mapping

Many 3D graphics systems use texture mapping to apply images, or textures, to geometric objects. Commodity PC graphics cards are fast at texturing and can efficiently render slices of a 3D volume, with realtime interaction capabilities.

These slices can either be aligned with the volume and rendered at an angle to the viewer, or aligned with the viewing plane and sampled from unaligned slices through the volume. Graphics hardware support for 3D textures is needed for the second technique.

Volume aligned texturing produces images of reasonable quality, though there is often a noticeable transition when the volume is rotated. View aligned texturing creates images of similar high quality to those of ray casting, and indeed the sampling pattern is identical.

2.5 Hardware-Accelerated Volume Rendering

A recently exploited technique to accelerate rendering is the use of modern graphics cards to accelerate traditional volume rendering algorithms such as ray-casting. Starting with the programmable pixel shaders that appeared around the year 2000, people recognized the power of parallel operations on multiple pixels and began to perform general purpose computations on the graphics chip. The pixel shaders, once called "register-combiners" were able to read and write randomly from texture memory, perform some basic mathematical and logical calculations. These SIMD processors, now called GPUs, were harnessed to perform general calculations such as ray tracing polygons and signal processing. With OpenGL version 2.0, the pixel shaders now are able to function as MIMD processors (now able to independently branch) with as many as 48 parallel processing elements utilizing up to 1GB of texture memory and high bit depth numerical formats. With such power, virtually any algorithm such as volume ray casting or CT reconstruction can be performed with tremendous acceleration.



3. APPLICATION AREA

Volume rendering technique is widely used in different areas. These include:

3.1 Medicine

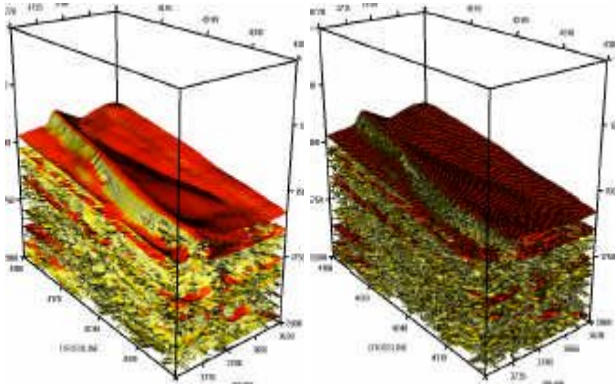
With the evolution of image acquisition techniques, both in resolution and tissue distinctiveness, the capacity and fidelity of image diagnosis were extended. Once the medical image is obtained it has to be analyzed by a radiologist. Medical images are used to visualize the structure or function of parts of the human body and to detect when the characteristics of some disease are present.

The physician could use the computer to manipulate the images, e. g. change the colors of an image, to allow a better visualization. Among Image Processing techniques, digital filtering (e.g. to remove noise), registration and segmentation are the most used in medical systems. Registration refers to the alignment process of several images obtained either from a single scanner or from different types of scanners. Segmentation consists in dividing an image in parts according to its characteristics, and it is used to detect objects or regions of interest.



3.2 Seismic data analysis

In modern seismic interpretation, these techniques have been employed as basic tools to view structural and stratigraphic features from 3D seismic data. Direct volume rendering is particularly interesting for previewing seismic datasets. Using DVR the user can interactively manipulate the transfer function to control the iso-surface displayed.



4. OPEN QUESTIONS

The biggest open question is increase in image quality. This factor is constrained to development of better hardware that could handle big amounts of visualized data. Even at the moment, it is required to have very good computers to achieve high-quality pictures.

The other question is optimization of known algorithms for volume rendering to obtain faster results. Research in this area is being done at the moment.

5. CONCLUSIONS

Volume rendering is a big research field that is has seen a lot of activity over the past decade and will continue to do so in the future with the introduction of better hardware accelerators. The original techniques which all had a number of drawbacks initially have received many updates that have made them faster with better image quality and more versatile. Hardware acceleration of volume rendering is still in its infancy and it is very likely that it will see a similar development as the hardware matures.

While the software rendering techniques have enjoyed very good picture quality, they have also lacked an important quality, and that is interactive framerates. The time to render a single image is sometimes exceeding three minutes, which means that finding a good transfer function configuration is a very time consuming procedure. Getting interactive framerates combined with a good picture quality is a very precious characteristic since the data can be analyzed much faster that way.

The applications of volume rendering are many and varied. The driving force is primarily the medical field which can use the rendered images to help diagnose illnesses, or to practice performing virtual surgeries. The special-purpose hardware acceleration cards that are available to these purposes are expensive and are not widely available. Hardware acceleration with general purpose hardware brings volume rendering with the same performance to anyone with a PC which means that both price and availability is far better.

6. REFERENCES

- [1] Isabel Harb Manssour, Carla Maria Dal Sasso Freitas “COLLABORATIVE VISUALIZATION IN MEDICINE”, The 8th Internation Conference in Central Europe on Computer Graphics, Plzen, Czech Republic, February 7-11, 2000
- [2] Pedro Mário Silva, Marcos Machado, Marcelo Gattass, “3D Seismic Volume Rendering”, 2003, 11Tecgraf/PUC-Rio, Petrobras
- [3] Steven P. Callahan, Milan Ikits, Joaõ L.D. Comba, and Cla´udio T. Silva, “Hardware-Assisted Visibility Sorting for Unstructured Volume Rendering”, IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS, VOL. 11, NO. 3, MAY/JUNE 2005
- [4] www.wikipedia.org
- [5] <http://home.student.uu.se/c/cebl4603/>
- [6] <http://www.cs.umu.se/education/examina/Rapporter/MarcusJonsson.pdf>