Illumination in Direct Volume Rendering

Frida Hernell[1,2] (frida@hernell.se), Patric Ljung[3] and Anders Ynnerman[1,2]
1.Visual Information Technology and Applications
2. Center for Medical Image Science and Visualization
3.Siemens Corporate Research
Linköpings universitet

Introduction

Interaction with patient specific Computed Tomography (CT) data through direct volume rendering (DVR) has become an invaluable tool in medical diagnosis. To further promote the use of DVR in medicine, new rendering methods that more efficiently and intuitively lead the user to areas of interest while maintaining the integrity of the visual data representations are needed. One area of development that shows promise in supporting this is the inclusion of more sophisticated lighting models that provide visual cues to enhance the perception of shapes and depth. However, general methods for the solution of the full light transport equations are computationally very demanding and many applications, including medical, must run at interactive frame rates. We present two efficient illumination methods with inclusion of self shadowing for improved perception.

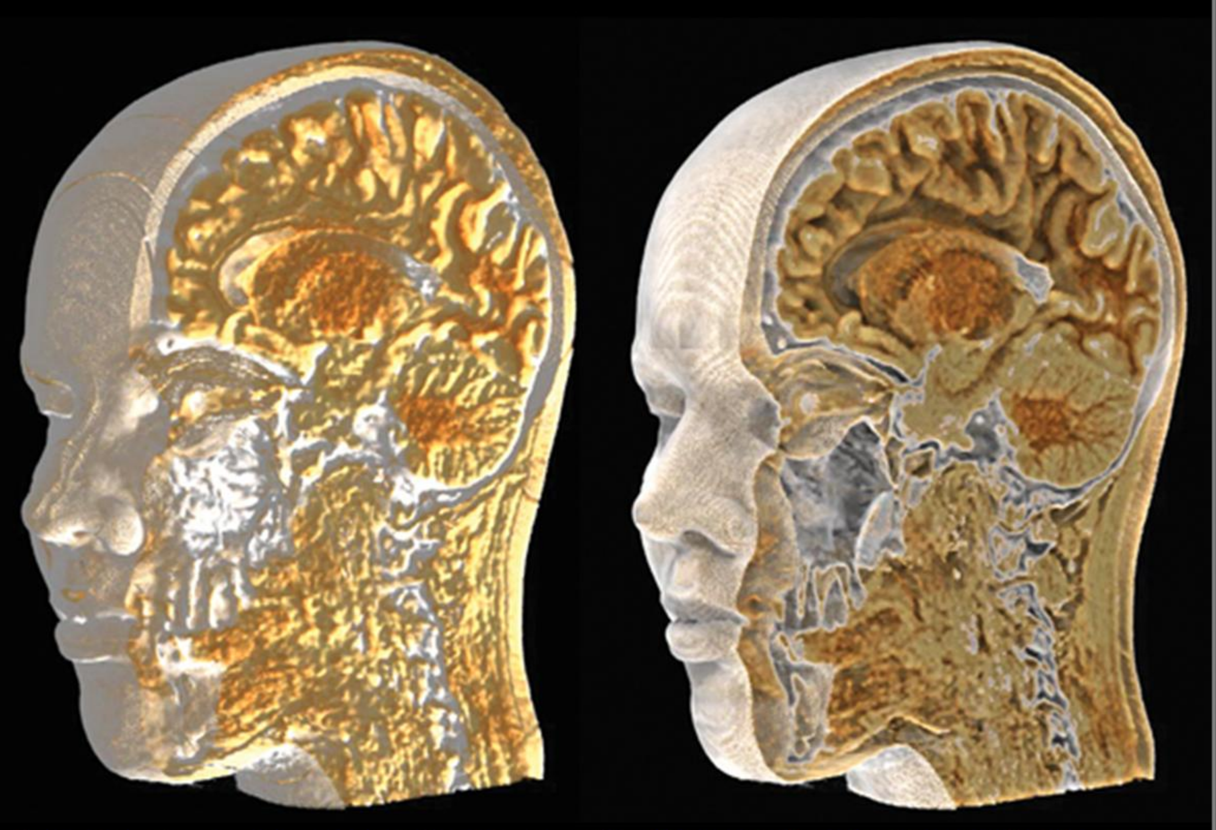

A magnetic resonance (MR) volume with different illumination methods. The undefined gradients at the clipped boundary results in misleading illumination when using diffuse shading (left). Shapes are more easily perceived in the LAO illuminated volume(right).

Local Ambient Occlusion (LAO)

In our first method we present an efficient approach for ambient occlusion, which approximates light arriving to each voxel from all directions, just like illumination on a cloudy day. Integrating ambient occlusion for each voxel is highly computational demanding, as illustrated in fig. 1(left). Faster calculations are possible by restricting the region of incident light to a local neighborhood (fig. 1(right)). Each voxel is thereby only occluded by structures in the vicinity. The intensity of incident light is derived by integrating the optical depth along rays cast from the voxel towards a local sphere in a number of directions. Incremental refinements of the incident light directions are employed to reach high frame rates during illumination updates.

LAO Concept

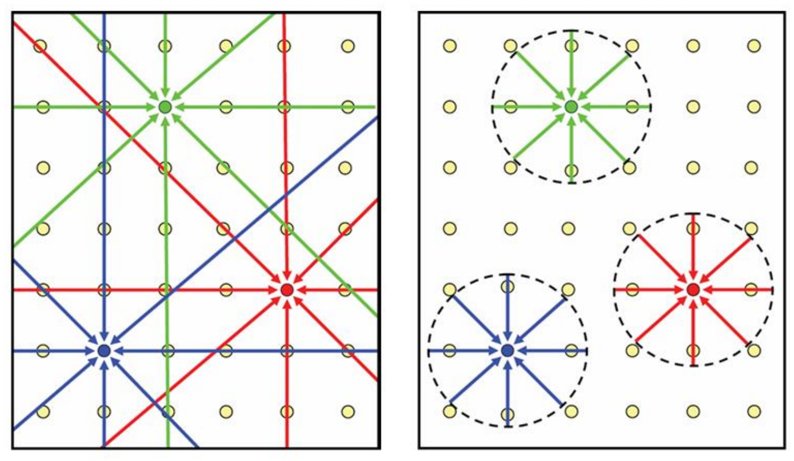

Fig. 1. Incident light is integrated along rays to compute ambient occlusion. This image illustrates computations for three points in the volume, red, green and blue. For a global approach (left) the rays must traverse the whole volume. In the local approach (LAO), the integration of incident light is only computed for rays within a sphere around each voxel (right).

(a) Diffuse Illumination (b) LAO with emission (c) Diffuse Illumination × LAO (with emission)

Example images showing the enhanced information from the emissive materials. The bullet and fragments are clearly visible in the abdomen. The effect of the LAO in revealing the bone structure is also very clear. A combination of LAO and diffuse shading intensifies fine details, for example the structure of the pelvis.

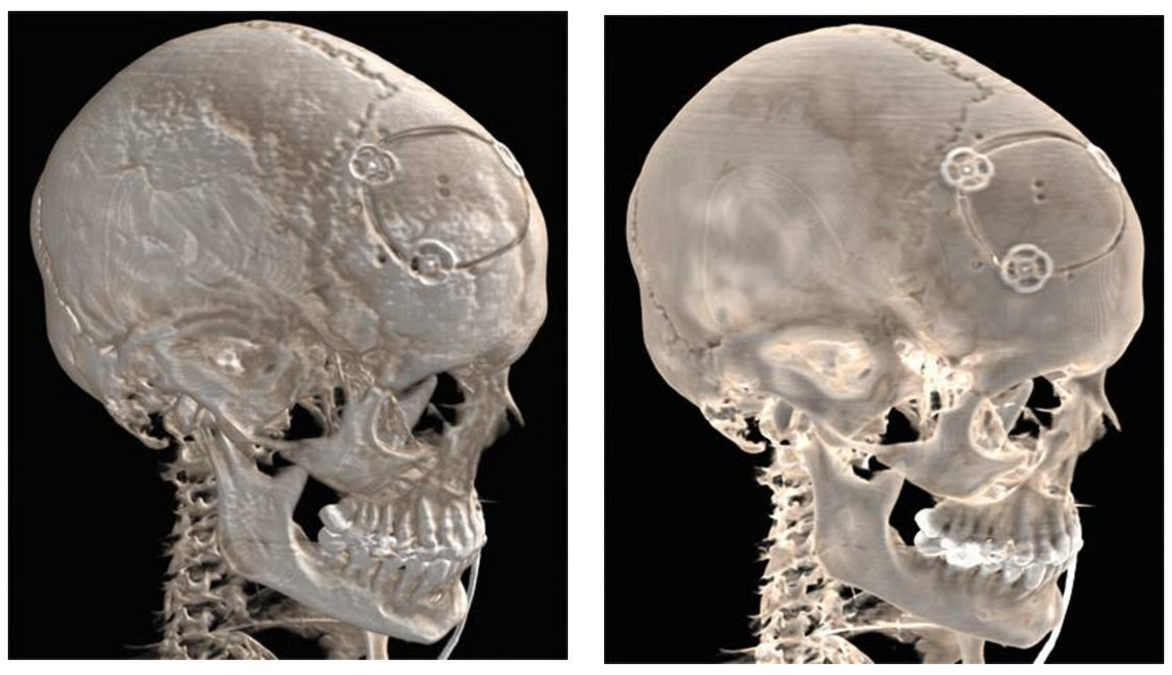

The density of the skull appears better with LAO (right) compared to diffuse shading (left). The LAO method is not sensitive to gradient data which can be an advantage for noisy data sets.

Local Piecewise Integration

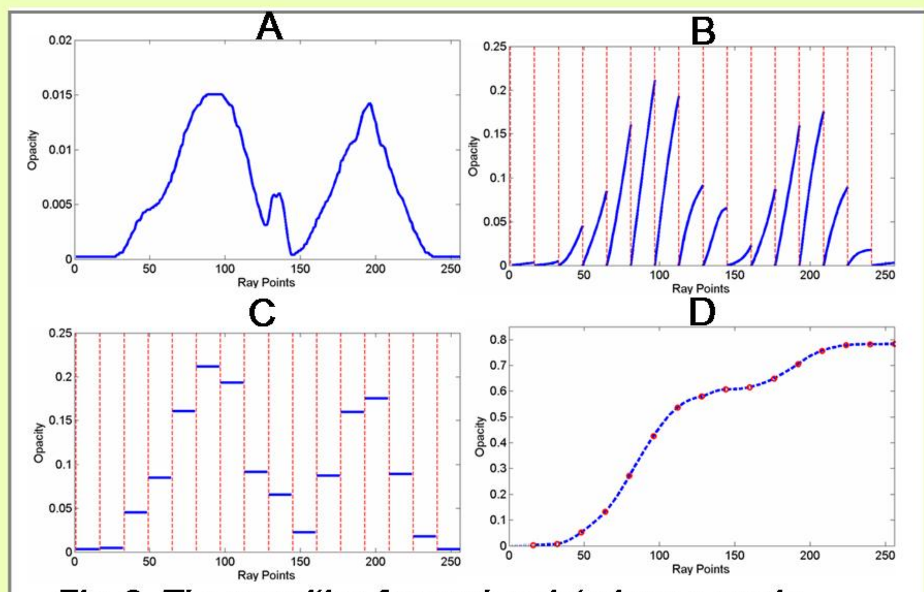

Fig. 2. The opacities for each point along a ray is illustrated in (a). Opacities of local segments are integrated in (b) and the resulting opacities for each segment are shown in (c). Computing opacity along the whole ray, in (d), from the first point to the light source, is performed by sampling (red circles) the opacities in (c), which allows for a sparse sampling density. Note that the local piecewise integration must be performed and stored for each point along the ray to allow the global sampling to start at an arbitrary point.

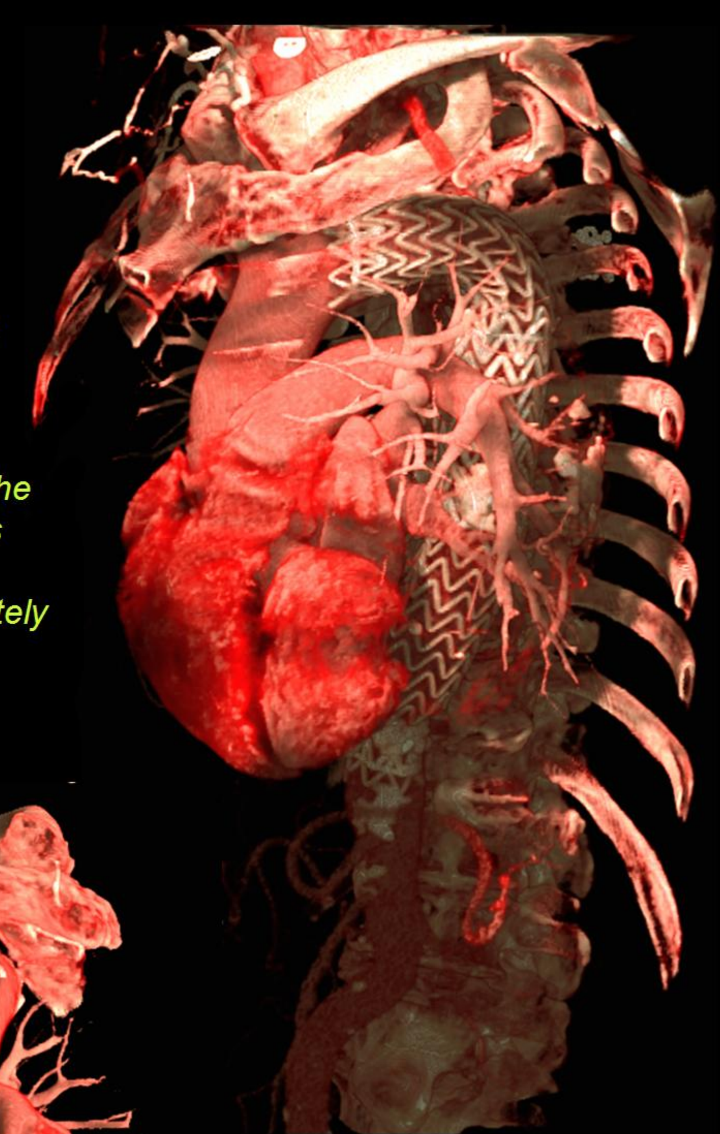

Two examples, with varying position of the light source, using the presented illumination technique. The frame rate is 2.5 fps when the light source is moving, 4 fps during in-scattering refinements and approximately 15 fps when the light computations are updated.

Global Light Propagation using Piecewise Integration

In our second method we present an efficient approximation of volumetric light propagation from a directional light source with first order scattering effects. The method progresses in several steps to simulate the light transport in the volume. Opacities are composited into a global shadow volume, in the fist step, for a given light source and transfer function settings. The calculation is based on piecewise pre-integration techniques (fig. 2) on the GPU and configurable sparse representations of intermediate results. The method then proceeds to include first order scattering of the global light arriving at voxels by integration of scattered light in a local spherical neighborhood around each voxel (similar to LAO). The resulting application enables updates of light position and transfer functions while maintaining interactive frame rates and still simulating a nearly physically correct light model. To speed up the interaction during an illumination update the final intensity is progressively refined by adding computations for one additional in-scattering direction, for each frame.

Implementation

Both illumination methods are implemented on the graphics processing unit (GPU) for fast computations. Additional improvement in performance is reached by employing a flat blocking multiresolution data management, which allows for empty space skipping in the illumination calculations and reduces the storage space of intermediate results.

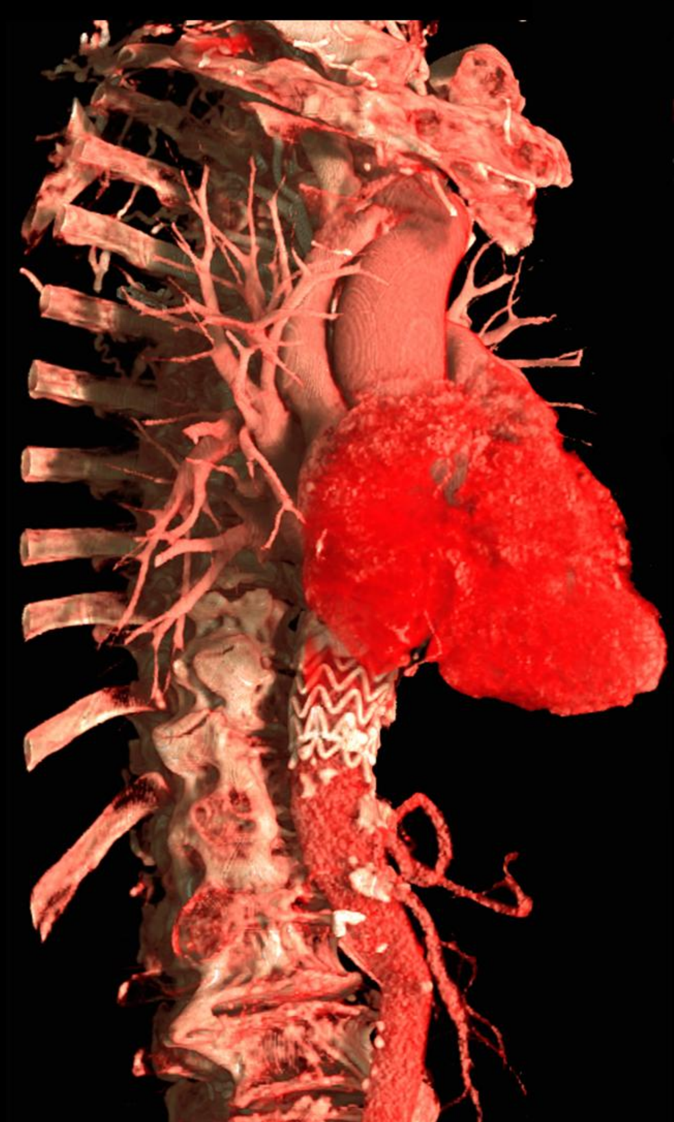

The medical volume used in the images, above and to the left, has a resolution of 512^3 voxels. A stent is inserted into the aorta.

A CT scan of a lion (512x512x1100 voxels) rendered using the proposed method for interactive volume rendering. An initial global light distribution is evaluated in 300 ms and refined incrementally in 185 ms for single scattering.

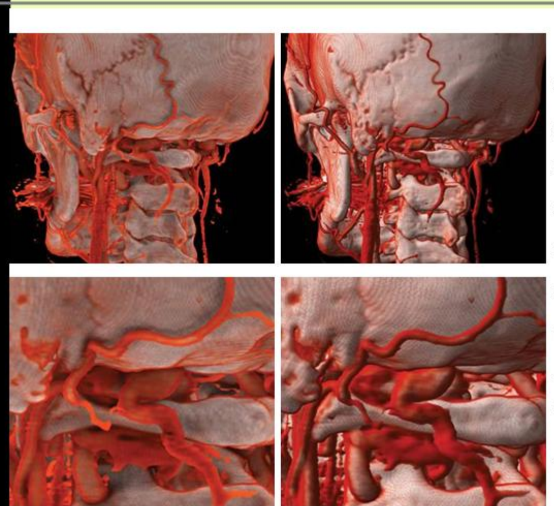

Improved perception of depth is obtained using LAO (left column) compared to diffuse illumination (right column) since shadows give visual clues of spatial relations.

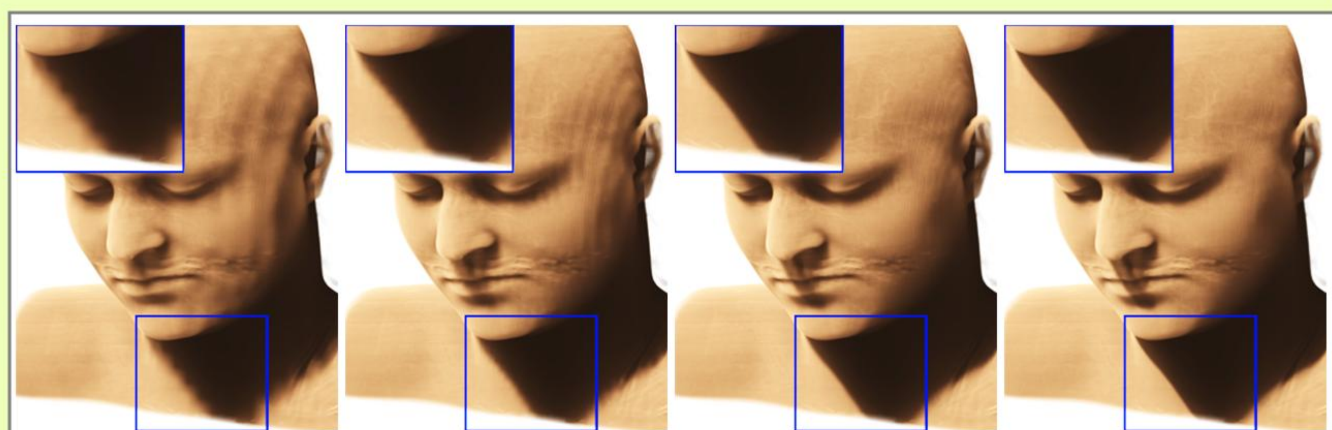

Computations for global light propagation using reduced sizes of the global shadow volume. From left to right: 32^3, 64^3, 128^3 and 256^3 for a 512^3 volume.

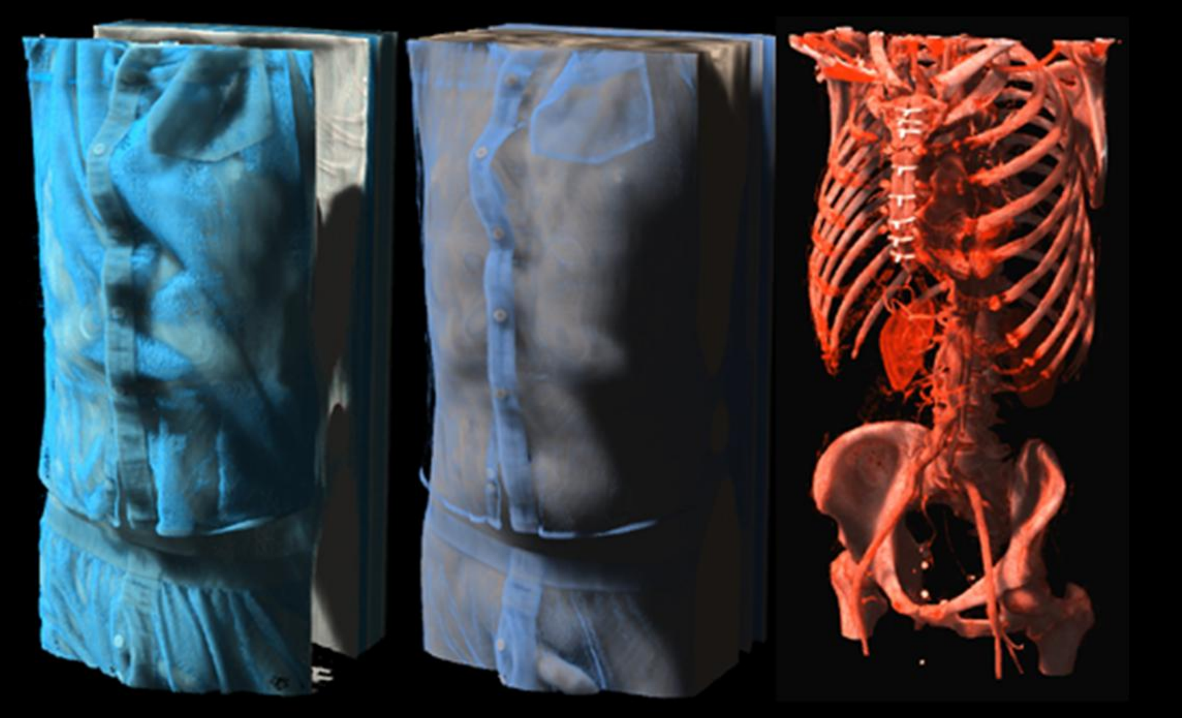

Illumination computed for a CT volume using local piecewise integration and different TF settings.