

“VolumeViz”: Interactive Visualization of Cross-Sections in Scientific and Medical Datasets

Jacob Stern and Robin Martens
Advisor: John Hughes

May 1, 2014

Three-dimensional datasets produced by procedures such as MRI and CT scans are important in many fields of research, medicine, and education. Often our primary means of interfacing with these datasets is through two-dimensional cross sections such as those often associated with MRI results. “VolumeViz” is an application that attempts to combine three-dimensional and two-dimensional visualization and manipulation of these datasets in real-time via an easy to use interface.

The rendering pipeline in VolumeViz is hardware-accelerated via OpenGL and CUDA. Voxel datasets are stored in GPU memory as a 3D texture prior to rendering. To create the 3D visualization, start and end points for ray marching are precomputed in an OpenGL fragment shader. The resultant textures are then bound to a CUDA context and the volume is rendered by ray marching from the front to the back of the bounding cube in a CUDA kernel, sampling from the 3D texture until a desired opacity is reached. Voxels in the texture are shaded using a transfer function, coded as a 1D texture in the CUDA context, and optionally a simple Phong shading model. To compute gradients for Phong shading, a slab-based ray casting approach is used wherein the volume is divided into image-space “slabs” which are mapped to individual CUDA blocks. This allows localized caching of voxels sampled from

the 3D texture when computing gradients.¹ The 2D cross-section visualization is computed concurrently in a separate CUDA kernel.

Our application provides a three-pane interface for specification of cross sections in the visualization. The first pane allows the user to specify one of three possible canonical cross sections in the volume: sagittal, coronal, and horizontal. The second pane allows for specification of a wider number of cross sections parameterized by origin and rotation of the slicing plane. The third pane allows for arbitrary cross sections through a novel image-space UI wherein the user may click and drag in the image to specify a cross section plane.




Figure 1: "Engine" dataset with Phong shading.



Figure 2: MRI cross section.

¹ See [Mensmann et al. 2010](#).