

An Immersive Virtual Environment for DT-MRI Volume Visualization Applications: a Case Study

S. Zhang, Ç. Demiralp, D. F. Keefe
M. DaSilva, D. H. Laidlaw, B. D. Greenberg*
Brown University

P.J. Basser
C. Pierpaoli†
National Institutes of Health

E.A. Chiocca
T.S. Deisboeck‡
Massachusetts General Hospital

Abstract

We describe a virtual reality environment for visualizing tensor-valued volumetric datasets acquired with diffusion tensor magnetic resonance imaging (DT-MRI). We have prototyped a virtual environment that displays geometric representations of the volumetric second-order diffusion tensor data and are developing interaction and visualization techniques for two application areas: studying changes in white-matter structures after gamma-knife capsulotomy and pre-operative planning for brain tumor surgery. Our feedback shows that compared to desktop displays, our system helps the user better interpret the large and complex geometric models, and facilitates communication among a group of users.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual Reality; I.4.10 [Computer Graphics]: Image Processing and Computer Vision—Image RepresentationVolumetric; J.3 [Computer Applications]: Life and Medical Sciences—Biology and Genetics

Keywords: Scientific Visualization, DT-MRI, Diffusion, Medical Imaging, Virtual Reality.

1 Introduction

Diffusion tensor magnetic resonance imaging (DT-MRI) is a new imaging modality with the potential to measure fiber-tract trajectories in fibrous soft tissues such as nerves, muscles, ligaments, and tendons [1]. However, the datasets produced are volumetric, with six values at each spatial location, and hence a significant challenge to visualize and understand. The potential of visualization tools to reveal the internal structures from these multi-valued volume images is likely to be great but is, as yet, not well exploited.

A number of methods have been developed that display the data as 3D models [2, 3, 4]. Some methods, including ours, track the fiber-tract trajectories in DT-MRI's [1, 2, 5]. Our method generates streamtubes to follow these trajectories and streamsurfaces to represent planar structures; we also display anisotropy values on both streamtubes and streamsurfaces [5]. To provide anatomical context, we generate ventricle surfaces and skull surfaces.

The resulting set of geometric models is large and complex. With traditional monocular single-screen displays, it is difficult to interpret these complex geometric models. The situation is exacerbated when the complexity of the models is increased to meet user needs for more detail (see color Figs. 1 and 2).

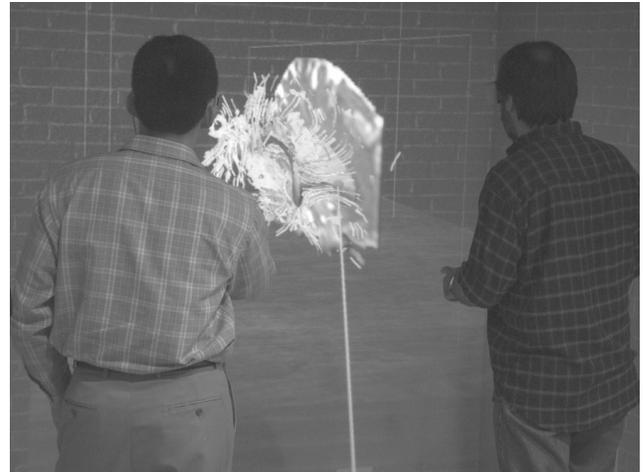


Figure 1: Two users stand in a virtual room within the Cave and manipulate a virtual brain dataset. The environment gives them a better sense of the structures in the 3D data than static images or a desktop display. See color Fig. 3.

An immersive virtual environment such as the Cave [6] has the advantages of a head-tracked stereo display, a large display surface, and interactivity [7] (see Fig. 1 and color Fig. 3 to 5). In such environments, it has been demonstrated that users are able to improve performance involving similar complex models by 200% compared to working with a static image [8]. While virtual environments have already been applied to related applications such as reconstructive surgery and other medical applications [9, 10], they have not yet been used to visualize DT-MRI datasets.

This paper describes our efforts to use the Cave for DT-MRI visualization. Our targeted applications include studying changes in white matter structures before and after gamma-knife capsulotomy and studying anatomy during pre-operative planning for brain tumor surgery.

2 Virtual Reality Setup

Our Cave is an $8 \times 8 \times 8$ foot cube with rear-projected front and side walls and a front projected floor, shown in Fig. 2. A user wears a pair of LCD shutter glasses that alternate images to each eye; the graphics system calculates stereo imagery in synchrony with the glasses. The glasses also have an attached tracker that relays their position and orientation to the computer. With that information, the graphics system can display the virtual environment so that it appears stationary to the user moving through it.

The goal of our setup is to let the user feel comfortable and oriented in the Cave and focus his attention on the visualization. The virtual environment consists of a room the size of the Cave itself and a table under the visualization of the DT-MRI data. Since the

* {sz,cad,dfk,mds,dhl}@cs.brown.edu, bdg@butler.org

† {pj)basser,carlo}@helix.nih.gov

‡ {chiocca,deisboeck}@helix.mgh.harvard.edu

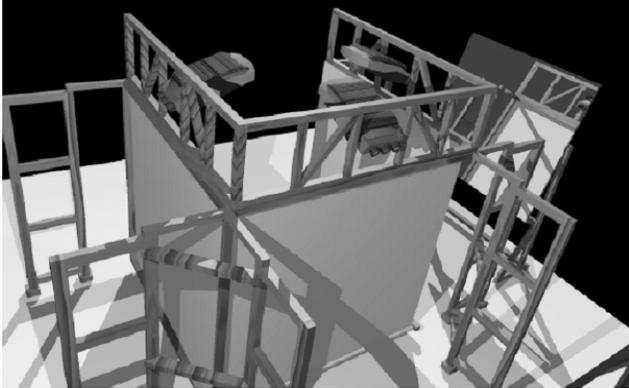


Figure 2: A model of the Cave viewed from behind the back wall. The user stands inside the three rear-projected walls.

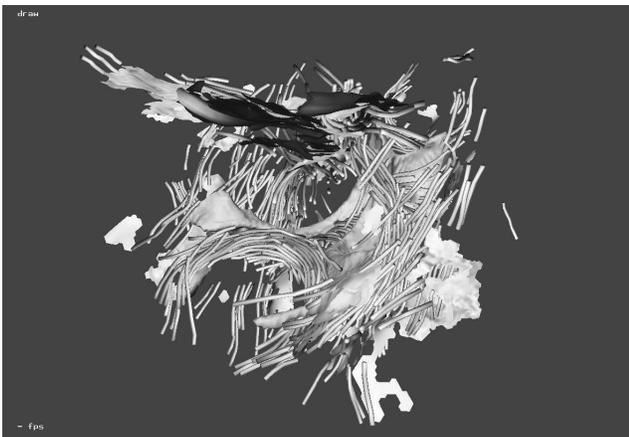


Figure 3: A set of geometric models representing the DT-MRI dataset. Tubes represent linear structures, surface fragments represent planar structures, and an isosurface shows the shape of the ventricles. See color Fig. 1.

user can see both the virtual world and real world objects through the shutter glasses, we texture map the virtual walls and floor in the same positions as the real walls and floor. This method decreases visual conflicts. We display geometric models of streamtubes, streamsurfaces, ventricles, and other isosurfaces inside the data boundary. 2D T2-weighted MR sections are also displayed, and the current section and axis can be chosen interactively. Perhaps most importantly, the user can walk around, bend over, and move his head to observe the visualization from different perspectives and with different scales. The model appears stationary, and we hypothesize that a user can build a cognitive model of anatomy or pathology more effectively this way than by moving a model around on a monitor while holding his head still.

3 Data Preparation

We display DT-MRI and T2-weighted images of a human brain. Both of the datasets are registered in a $256 \times 256 \times 40$ volume. Each voxel is $0.89 \times 0.89 \times 3.2$ mm. We use tricubic B-spline functions to interpolate the second-order diffusion tensor field, and then interpolate the eigenvector field by calculating the eigenvector of the interpolated tensor at any location [1].

4 Geometric Models of DT-MRI Data

Visualizing neural structures contained in DT-MRI data is our primary interest. Thus we generate streamtubes and streamsurfaces to represent linear and planar structures. There is also anatomical context in the visualization to help users navigate the complex models. To ensure interactivity, we select a subset of the possible geometric models, and decimate them for display. These operations are discussed below.

4.1 Streamtubes and Streamsurfaces

We visualize DT-MRI data by generating geometric models that correlate with the underlying structures of the dataset.

First, we distinguish between regions of linear and planar anisotropies in the DT-MRI dataset, and then we employ streamtubes and streamsurfaces, respectively, to visualize these two types of anisotropic diffusion. Fibrous structures in the brain, like white matter tracts, result in diffusion tensors with linear anisotropy. We represent them with streamtubes, which are thin tubes that extend from seed points following the direction of fastest diffusion. This direction is only defined in regions of relatively high linear anisotropy, so we only create streamtubes in those regions. A streamtube ends when it hits the data boundary, reaches a region of low linear anisotropy, or bends too much. Similarly, we generate streamsurfaces to follow planar structures in regions of planar anisotropy. We map colors onto the geometric models to present the degree of linear or planar anisotropy. More saturated colors represent higher anisotropy. Fig. 3 (color Fig. 1) shows a set of streamtubes (red) and streamsurfaces (green).

The process for choosing a small, but representative, set of streamtubes operates as follows. Initially we select a dense set of seed points in the volume so as to cover as many interesting features as possible. The resulting number of streamtubes is huge. For a $256 \times 256 \times 144$ volume we produced more than 150,000 streamtubes by setting a seed point in every voxel. Some of them run very close to other streamtubes. Some are very short and are likely the result of noise. We use a culling algorithm to pick a representative subset [5, 11]. Culling criteria include the length of the streamtube, the average anisotropy along it, and its similarity to other streamtubes. The culling process that produced Fig. 3 (color Fig. 1) resulted in 694 streamtubes in the scene. Color Fig. 2 shows an example where less-restrictive culling parameters resulted in 1,347 streamtubes. Streamsurfaces are culled analogously.

4.2 Anatomical Context

The T2-weighted image shows the boundaries between different types of tissues in the brain in a format familiar to medical researchers. To give users more context, we also generated an isosurface representing the ventricles from the T2-weighted image using the Marching Cubes algorithm [12]. The blue surface in color Fig. 1 shows the approximate shape of the ventricles. Surfaces of other anatomical structures, such as blood vessels or tumors, could also be generated from appropriate imaging datasets.

4.3 Performance Considerations

With the amount of information we have, we need to balance between fidelity and interactivity. User feedback indicates that users want to see more detail. At the same time, delays in producing images for display need to be short enough for good interactivity. Initially we put a total of more than 500,000 triangles into the scene, and the system hardly moved. We learned to be careful in choosing what to show. Visualizing neural structures is our primary interest, therefore streamtubes and streamsurfaces get high priority. Isosurfaces of various anatomical features get different

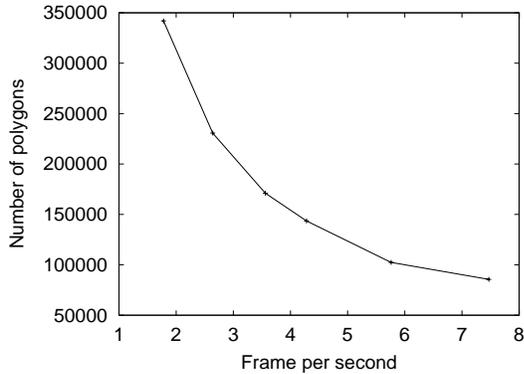


Figure 4: A graph of the number of polygons versus frame rate. This shows the fidelity versus interactivity tradeoff in the Cave.

priorities for different tasks. For example, to a neurosurgeon planning a brain tumor surgery, blood vessels around the tumor are important anatomical landmarks. Alternatively, to a student of neural science, the ventricles are extremely useful for identifying the neural structures around them.

Besides choosing what to put into the scene, we also decimate the chosen models to reduce the total number of polygons. Again, different applications require different decimation ratios for different kinds of geometric models. A cortical surface, for example, could be decimated to a smooth surface to provide a crude approximation of the skull, or its topological details could be preserved to provide a fine context for the neural fibers running beneath it.

We use the decimation algorithm designed by William Schroeder [13] and implemented in VTK [14] to reduce the size of the isosurface models, and use our own algorithm to reduce the size of the streamtube models. Our algorithm ensures that the distance between the trajectories of the decimated streamtube and the original one is less than a given threshold everywhere. By setting the threshold to 0.2 mm, we can achieve a decimation ratio around 85%.

We tested the frame rate of the system with various numbers of polygons in the scene; the results are shown in Fig. 4.

With the number of polygons below 100,000, the system moves at more than 5 fps, which is enough to perform most interactive tasks in our application. However, we have to sacrifice detail on these models. Users report the need for higher anatomical resolution. Further study on the application specific mesh decimation algorithms would help speed up the system and preserve the important detail.

5 Displaying 2D MR Sections

The doctors who first used our system suggested that viewing typical 2D T2-weighted MRI sections might help in identifying anatomical features in the 3D models. Displaying 2D sections along with geometric models has several advantages. Doctors have years of experience in analyzing these sections. Examining the geometric models along with a 2D section gives users a way to interpret and validate the geometric models. The display of a section complements the features chosen in the geometric models by displaying the T2-weighted image everywhere within the section; the user can examine a region in detail by sliding the slice back and forth in that region. The slicer can also display scalar datasets other than the T2-weighted one. We have used the T2-weighted image on the slicer to validate the ventricle models, the linear anisotropy dataset to validate and understand the streamtube models, and the planar anisotropy dataset to validate and understand the streamsur-



Figure 5: A user using natural body motions to see more detail by leaning over the model. See color Fig. 6.

face models. In the future, particularly for brain tumor surgery applications, we will use T1-weighted 2D sections, possibly acquired using a gadolinium-based contrast agent. The addition of 2D MR section does not slow down the system significantly.

6 Interaction

In the current interaction scheme, geometric models and sections stay still above the virtual table. The user walks around and adjusts his head position to observe the data from different perspectives. Fig. 5 (color Fig. 6) shows a user leaning toward the visualization to get a closer look at the geometric models around the ventricles.

Because only one user is tracked, other users see a distorted world, and finger pointing is ineffective for communication. We generated a yellow line akin to a laser pointer from the wand in the direction the wand is pointing. The line stops where it hits an object. The user can use the wand to indicate a region of interest to others. The user can also change the position of the slice or cycle through axial, coronal, and sagittal sections with the buttons of the wand (see color Figs. 3 to 5).

7 Feedback

Five users, three MDs and two senior medical students, have used the system and provided feedback. Their focus on the data and application areas suggests that our system is easy to use. They had each also viewed images of the geometric models on a workstation screen, and reported that they found it easier to understand the 3D structures in the virtual environment.

In a typical visit, it took a few minutes for a user who had never used our application to understand how the Cave works and how to interact with the application. When a typical user put on the tracked glasses and began navigating, he tended to walk around the visualization, stopping from time to time to examine it from a specific perspective. Even when he stood still, he tended to slightly move his head to get more depth and shape information. We concluded that stereo and interactivity together help a user better explore a visualization with complex geometric models.

Once a user found a region of interest, he usually moved the slicer back and forth around that region and talked about the possible structures represented by geometric models in the context of the 2D sections. The 2D sections were reported to be a very valuable additional tool in the virtual environment. All of our users were eager to get more detail into the visualization. They suggested

that more anatomical detail and streamtubes in the regions of interest might help their understanding. Users also made application-specific suggestions, which we discuss below.

7.1 Studying Changes in White-matter Structures after Gamma-knife Capsulotomy

One of the users, a doctor from Brown University-affiliated Butler Hospital, has been using DT-MRI data to study OCD patients. He is particularly interested in studying changes that occur after radiation surgery that ablates an important white matter region. These changes should show up in the DT-MRI data, and the connectivity changes, which are inherently 3D, will be very difficult to see in sections alone. This should help generate a number of specific hypotheses about the mechanism that makes this treatment modality successful. Patients are currently being identified to participate in the project.

7.2 Pre-operative Brain Tumor Surgery Planning

Two of the users, neurosurgeons from Massachusetts General Hospital (Harvard Medical School), would like to use the system to perform surgical planning tasks. They also want to study 3D relationships between the tumor and other important structures like vasculature and motor cortex and thus would also like to add the option to display other scalar-valued images, for example, T1-weighted or contrast agent enhanced 2D sections, which show the viable highly vascularized tumor regions. With the virtual environment, they believe that they can get a better spatial understanding of the critical anatomical structures in the vicinity of the targeted tumor. Together with other colleagues from MGH-Radiology, they are in the process of identifying surgical candidates to image with MRI/DT-MRI in the routine pre-operative schedule and to study in our virtual environment before surgery. Their long term goal is the combination of a virtual environment with their novel computational (spatio-temporal) 3D tumor growth and invasion models to develop an interactive patient-specific treatment planning device and intraoperative navigation tool.

8 Summary and Conclusion

Our long term goal is to provide an efficient and effective visualization of DT-MRI volume data together with other medical imaging modalities. As effectiveness can only be measured in the context of a specific application, we intend to continue to pursue the applications that we have described. As a first step, we have constructed a framework in which geometric models are put into an immersive virtual environment. The users move in the virtual environment and observe the visualization in the same way they would view a physical object in the real world. Our feedback shows that virtual reality display and interaction help the user better interpret the complex geometric models representing neural structures. Anatomy provides important context in the visualization. Interactive 2D MR sections provide a useful context and complement for the 3D models. The environment shows strong potential for understanding these complicated datasets and the underlying anatomy and pathology that they can detect.

9 Acknowledgment

Thanks to Andrew Forsberg and Daniel Acevedo for their help developing the virtual environment, to Dr. Susumu Mori at Johns Hopkins for the human brain DT-MRI dataset, to Mark Oribello for creating Fig. 2, and to all of the users who provided valuable feedback and suggestions. This work was partially supported by the Human Brain Project (NIDA and NIMH) and NSF (CCR-0086065). Opinions expressed in this paper are those of the authors and do not necessarily reflect the opinions of NSF.

References

- [1] P.J. Basser, S. Pajevic, C. Peirpaoli, Jeffrey Duda, and A. Al-droubi. In vivo fiber tractography using dt-mri data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.
- [2] Rong Xue, Peter C.M. van Zijl, Barbara J. Crain, Meiyappan Solaiyappan, and Susumu Mori. In vivo three-dimensional reconstruction of rat brain axonal projections by diffusion tensor imaging. *Magnetic Resonance in Medicine*, 42:1123–1127, 1999.
- [3] David M. Weinstein, Gordon L. Kindlmann, and Eric C. Lundberg. Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *IEEE Visualization '99*, pages 249–254, 1999.
- [4] Gordon L. Kindlmann and David M. Weinstein. Hue-balls and lit-tensors for direct volume rendering of diffusion tensor fields. In *IEEE Visualization '99*, pages 183–190, 1999.
- [5] Song Zhang and David Laidlaw. Elucidating neural structure in diffusion tensor mri volumes using streamtubes and streamsurfaces. In *Proceedings of ISMRM*, 2001.
- [6] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In James T. Kajiya, editor, *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 135–142, August 1993.
- [7] Andrew S. Forsberg, Robert M. Kirby, David H. Laidlaw, George E. Karniadakis, Andries van Dam, and Jonathan L. Elion. Immersive virtual reality for visualizing flow through an artery. In *Proceedings Visualization '00*. IEEE Computer Society Press, 2000.
- [8] Colin Ware and Glenn Franck. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics*, 15(2):121–140, April 1996. ISSN 0730-0301.
- [9] Kevin Montgomery, Michael Stephanides, Stephen Schendel, and Muriel Ross. A case study using the virtual environment for reconstructive surgery. In *IEEE Visualization '98*, 1998.
- [10] S. Taylor. *Computer-Integrated Surgery*. MIT Press, 1996.
- [11] Song Zhang. Visualizing diffusion tensor mr image using streamtubes and streamsurfaces. Master's thesis, Brown University, 2000.
- [12] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 163–169, July 1987.
- [13] William J. Schroeder, Jonathan A. Zarge, and William E. Lorensen. Decimation of triangle meshes. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 65–70, July 1992.
- [14] W. Schroeder, K. Martin, and W. Lorensen. *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*. Prentice Hall, 1996.

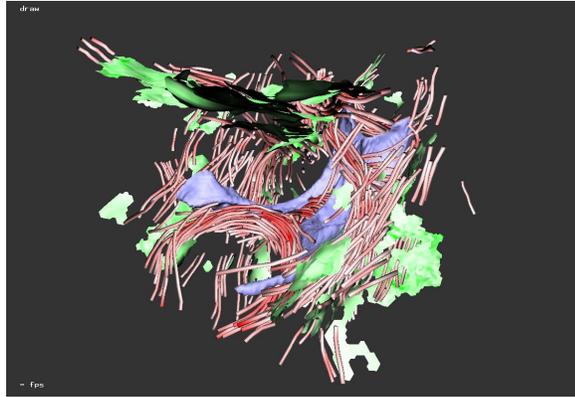


Figure 1: A geometric representation of a DT-MRI dataset. Red tubes represent linear structures, green surfaces represent planar structures, and the blue isosurface shows the surface of the ventricles. There are 694 streamtubes in this picture.

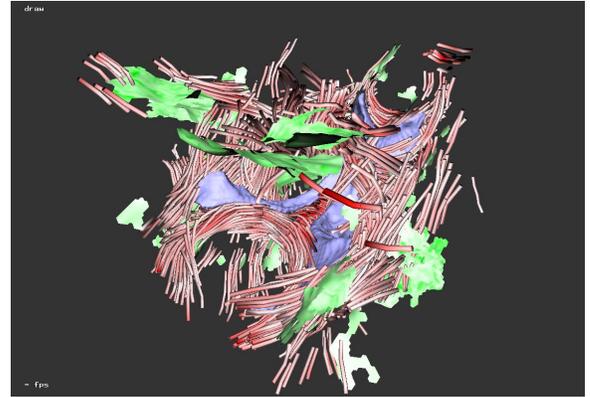


Figure 2: The use of less restrictive culling criteria results in 1,347 streamtubes in this figure. This more-detailed visualization is more difficult to understand using a static image than using our virtual environment.



Figure 3: Two users stand in a virtual room within the Cave and manipulate a virtual brain dataset. The environment gives them a better sense of the structure within the 3D data.

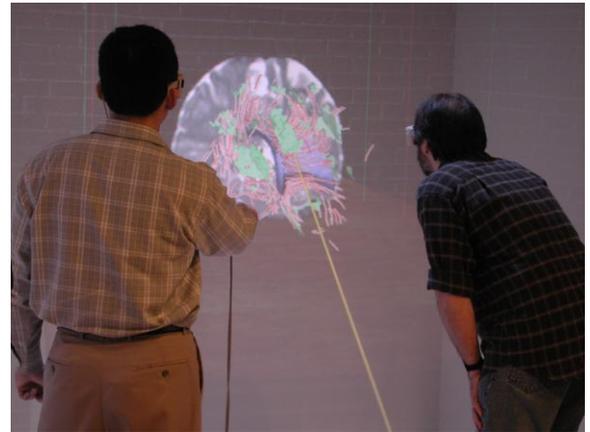


Figure 4: A user interacts with the virtual environment to change the orientation of the slicing planes.



Figure 5: A user views the model using a third orientation for the slicing planes. The yellow line is a “laser pointer” that helps the primary user communicate with colleagues.

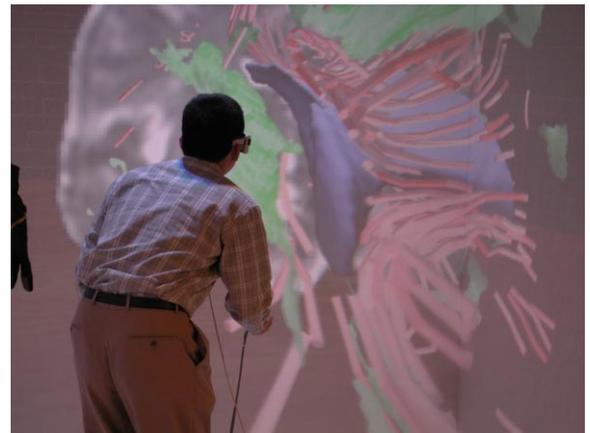


Figure 6: A user employs natural body motions to see more detail by leaning over the model. To the user, the model remains a constant size as it gets closer; to an outside viewer, it appears to get significantly larger.