
Global Illumination of White Matter Fibers from DT-MRI Data

David C. Banks¹ and Carl-Fredrik Westin²

¹ UT/ORNL Joint Institute for Computational Sciences
and Harvard Medical School dbanks@cs.utk.edu

² Harvard Medical School westin@bwh.harvard.edu

Summary. We describe our recent work in applying physically-based global illumination to fiber tractography. The geometry of the fiber tracts is derived from diffusion tensor magnetic resonance imaging (DT-MRI) datasets acquired from the white matter in the human brain. Most visualization systems display such fiber tracts using local illumination, a rendering technology provided by the video card on a typical desktop computer. There is indirect evidence that the human visual system perceives the shape of a fiber more quickly and more accurately when physically-based illumination is employed.

1 Introduction

This paper describes current work at the Laboratory for Mathematics in Imaging (LMI) at Harvard Medical School to incorporate global illumination into the visual data analysis pipeline for displaying 3D imagery of the human brain. The LMI works within the Department of Radiology at Brigham and Women's Hospital and has a particular emphasis on developing tools and techniques to analyze diffusion-tensor magnetic resonance imaging (DT-MRI). These tools are used by our collaborators to identify clusters of white-matter fiber tracts and to distinguish between fibers that infiltrate a tumor and fibers that pass around a tumor. In these tasks, an expert user examines a static image or a dynamic 3D display of geometry reconstructed from brain scans. The fiber tracts in these displays are long thin tubes, suggestive of the actual fiber bundles that innervate the brain's cortical surface (the gray matter). Water diffuses preferentially along the direction of the fibers, yielding orientation information in the DT-MRI data from which polygonal geometry is synthesized to create a 3D scene. The scene may contain many thousands of tubes representing fiber tracts. In the case of a patient with a brain tumor, surgical resection of the affected tissue destroys both the tumor and any fibers within it. As a result, the cortical regions connected through these fibers will lose their function. The treatment outcome may therefore depend on the relative pose of the fibers and the tumor; this information may be of great importance to the patient.

The geometric complexity of the fibers surrounding, and perhaps infiltrating, a tumor produces a 3D scene that is sometimes difficult to comprehend. The visual

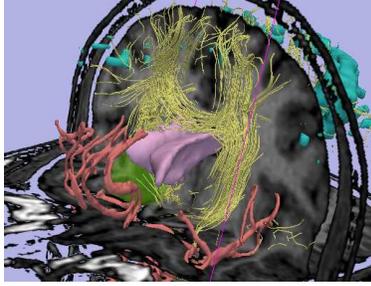


Fig. 1. Brain tumor (ganglioglioma) surrounded by blood vessels and infiltrated by fiber tracts. Visualization produced by *3D Slicer* using local illumination. Tumor represented in green; fiber tracts represented in yellow.

determination of whether a fiber passes through, or only near, the tumor surface can be a difficult task. Figure 1 illustrates the complex geometry of a tumor (green) surrounded by blood vessels (pink) and infiltrated by fiber tracts (yellow). The current state of the practice in DT-MRI visualization is to render the white matter fiber tracts as colored curves or tubes with either no illumination or with local illumination as provided by the underlying graphics hardware. We are investigating how to make physically-based global illumination of fiber tracts available to surgeons and patients.

2 Illumination

Illumination can be understood in different ways. It can be considered as a scattering process in which photons interact at the quantum level with molecules near the surface of an object. It can be considered as a mathematical equation for light transport. It can be considered as an algorithm implemented in hardware or software to produce an image. And it can be considered as the resulting perceptual phenomena that the human observer responds to. These latter two approaches (thinking of illumination as an algorithm or as a perceptual effect) are discussed below.

2.1 Algorithms for Illumination: Local and Global

Many 3D rendering techniques have been devised to improve the visual quality of rendered scenes via incremental changes within the basic polygon-rendering pipeline, where a polygon is illuminated without consulting other geometry in the scene; this technique is called local illumination. The one-polygon-at-a-time rendering approach is well suited for implementation in hardware, thus allowing a user to interactively manipulate a 3D scene containing thousands of locally-illuminated fibers.

Since the mid-1980's, efforts have been made to formulate the rendering process as a solution to the equation for light transport. In the 1980's, work at Cornell [4] [7] produced the first images in which radiosity techniques generated realistic images of scenes containing diffusely reflecting surfaces, while other work approximated the solution to the governing integral equation for light transport by using Monte

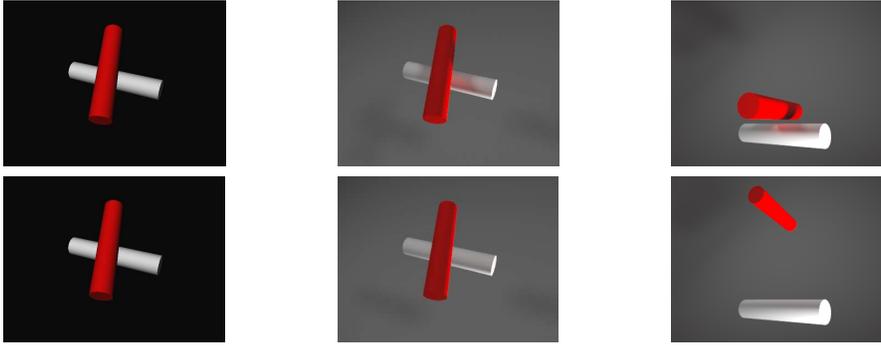


Fig. 2. Ambiguous poses of two tubes. Top row: tubes are closely spaced. Bottom row: tubes are far apart.

Carlo techniques [5] [10]. Interpreting renderers as light-transport solvers has led to increasingly realistic images that are sometimes indistinguishable from photographs. This strategy, which requires that the illumination at one point in a scene must account for illumination at other points in the scene, is called global illumination. Generally speaking, global illumination is performed in software rather than via graphics hardware and is therefore too slow to be used in an interactive visualization system.

The difference between local illumination and global illumination can be appreciated in Figure 2, which shows scenes containing a red tube and a white tube. In the scene shown on the top row, the tubes are close together, within a tube-radius of each other. In the scene shown on the bottom row, the tubes are far from each other, separated by several radii. When the tubes are viewed from above, the two poses are nearly indistinguishable with illumination (left images). But from the same viewpoint one can distinguish the two poses when the images are rendered using physically-based illumination (middle images): shadows and inter-reflections indicate when one tube is close to the other. When viewed from the side (right images), of course, the two poses can be distinguished immediately. Figure 3 shows a more complex scene containing many cylinders; note the shadows cast on and cast by the cylinder near the middle when global illumination is computed.

Two distinct threads have thus emerged in 3D graphics. One thread includes hardware-accelerated real-time video games and real-time visualization systems that

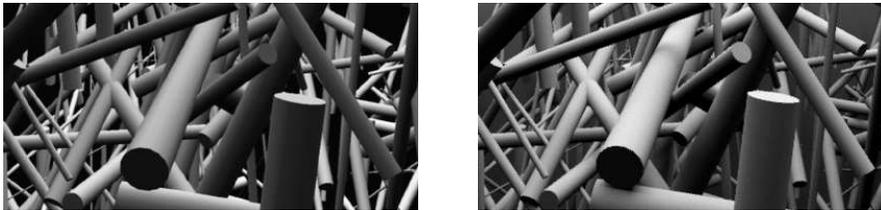


Fig. 3. Synthetic scene containing randomly placed cylinders. Left: local illumination. Right: global illumination.

use local illumination, where physical realism is sacrificed in favor of interactivity. The other thread includes studio films where interactivity is sacrificed in favor of realistic appearance of globally-illuminated 3D characters. The visual analysis of 3D scientific and medical data has generally followed the hardware-based thread. Visualization software tools have historically relied on the default capability of the graphics card to render images on the screen via local illumination (typically by implementing the OpenGL standard). The statement is sometimes repeated within the visualization community that “scientists don’t use global illumination.” But the scientific community may not have access to global illumination within the visualization tools they use. Unless the scientists are motivated to re-implement a physically-based light-transport code for data analysis, they are unlikely to develop an opinion regarding the value of global illumination or to demand it in their data analysis software. Thus a stable state currently exists in which visualization software generally relies on local illumination for its rendering since global illumination is not demanded by users; meanwhile, users generally do not render their scenes with global illumination since it is not available from their visualization tools. The absence of demand for global illumination in DT-MRI visualization software tools is both a cause for and the result of an absence of its supply, as illustrated below.



A basic unresolved question for those who design 3D visualization tools for displaying the complex geometry arising from brain data is: in order to make 3D shapes of lesions, fibers, blood vessels, and cortex perceptually evident to the human viewer (*e.g.* the surgeon), does physically-based rendering offer an improvement over the standard “local illumination” provided by graphics hardware? If so, then surgeons would demand this improvement and developers of visualization tools would be motivated to offer global illumination as a rendering option in their software systems.

Our work inserts global illumination into the visualization pipeline for analyzing DT-MRI data, giving neurosurgeons and radiologists the opportunity for the first time to analyze complex 3D scenes rendered with physically realistic illumination. Figure 4 illustrates the difference between local illumination and global illumination for a scene composed of several thousand fibers, represented as polygonal tubes, surrounded by a cut-away of a surface approximating the cortex. The brain geometry is situated within a rectangular box, exhibiting shadows and inter-reflections as a result of physically-based global illumination. Informal comments from colleagues (including radiologists and surgeons) indicate that the globally illuminated display of fibers “looks more 3D,” “looks more real,” “looks prettier,” or “looks like you could touch it.” Although there may be intrinsic value in the esthetic appeal of global illumination, it has not yet been shown whether global illumination of fibers provides measurable improvements in shape inference. We are currently engaged in designing a user study to measure differences in recognition and performance that may arise as a result of local versus global illumination of fibers.

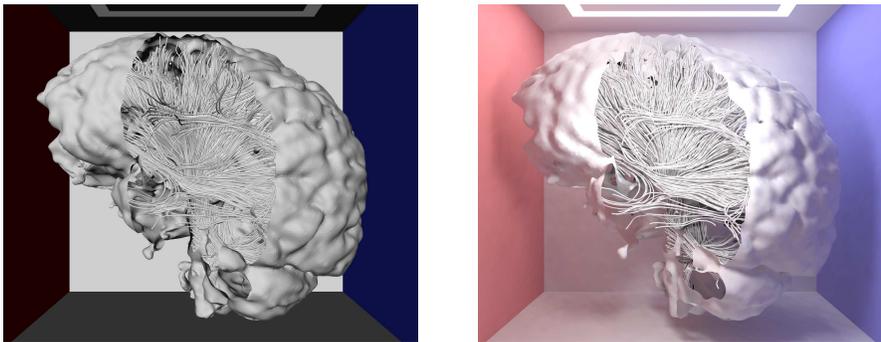


Fig. 4. White matter fibers within a cut-away approximation of the cortex. Left: local illumination. Right: global illumination. Brain dataset courtesy of Gordon Kindlmann at the Scientific Computing and Imaging Institute, University of Utah, and Andrew Alexander, W. M. Keck Laboratory for Functional Brain Imaging and Behavior, University of Wisconsin-Madison.

2.2 Perception of Illumination: Shadows

Although there is evidence that global illumination enhances 3D shape perception in general, this hypothesis is the subject of debate. The issue of whether global illumination adds perceptual value to a scene is sometimes characterized by whether shadows enhance perception of 3D shape. From a physics standpoint, this characterization is awkward: when we are situated in an unlit room, every object is “in shadow” from sunlight, being illuminated indirectly by skylight through the window, and yet fully shadowed 3D shapes are perceived all around us within the room.

While the physics of light transport proceeds without regard to the phenomenological qualities a human observer may attribute to the perceived scene, an artist may “add” specular highlights to shiny surfaces, “add” shading to rounded objects, and “add” shadows onto the painted canvas, and therefore conceive categories (highlights, texture, shadow) and subcategories (shadow is attached, grounding, or cast) of perceptual phenomena. The distinction is important when considering how physically-based illumination may improve the perceptual quality of a rendering of fiber tracts. A 3D visualization system can be designed to numerically solve the equation of light transport to greater or lesser degrees of accuracy, or else it can be designed to add perceptual phenomena to a greater or lesser degrees of completeness within artistically conceived categories. Historically, the latter artistic approach has led to disputes about whether or not shadows should be “added” to a 3D visualization.

There is conflicting evidence about the role that shadowing (as a phenomenological effect) plays in the perception of geometric attributes of a 3D scene. Madison [12] and Mamassian [13], for example, found that cast shadows provide a positive depth cue when present in a scene. But shadows can be problematic in a complex 3D scene. For example, shadows from tree branches on a sunny day can ruin the appearance of human faces in an outdoor photograph, in part because these shadows act as camouflage and in part because they interfere with the face’s natural shading that would be exhibited were the tree not interposed between the face and the sun.

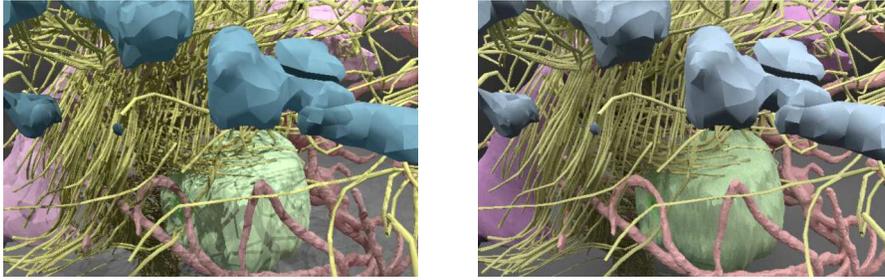


Fig. 5. Shadows in globally illuminated ganglioglioma dataset (used in Figure 1) resulting from luminaires of different sizes. Tumor represented in green; fibers represented in yellow. Left: Small-area luminaire produces sharp shadows. Right: Large-area luminaire produces soft shadows.

This a problem may be eliminated by using area lights rather than a single point light; Figure 5 illustrates the difference in appearance between a small-area and a large-area light source. The small-area luminaire produces sharp shadows cast by the fiber tracts and blood vessels onto the surface of the tumor; these shadows may hinder visual determination of which fibers penetrate the tumor.

The human visual system is somewhat insensitive to inaccuracies of shadows in a synthetic rendering: Ostrovsky [17] found that inconsistent shading of objects is slow and/or difficult to detect, while Jacobson [8] found subjects could not distinguish between a stripe-textured cube shadowed by leaves or a leaf-textured cube shadowed by stripes. But if 3D shadow and shape perception occurs early in visual processing, then the cognitive task of inferring a globally consistent light source or determining the source of a shadow may not correlate with the early mental formation of a geometric model influenced by the presence of shadows. The issue of whether shadow-processing is performed by an early-level system has not been resolved, but experiments by Rensink [19] suggest that shadow-processing is indeed performed during the early stages of visual search. Since visual search is an essential component of the surgeon's task during visual data analysis, display technologies that exploit early-level visual processing may improve the surgeon's spatial perception of the complex geometric interplay between lesions, fibers, blood vessels, and cortex.

3 Visual Integration of Glyphs and Curves

Figure 2 illustrates the difference between local and global illumination for only a single pair of straight tubes. Our actual concern is to apply global illumination to many long curvilinear tubes, each constructed explicitly as a polygonal mesh as in Figure 4, or else implicitly as a set of disconnected glyphs as in Figure 6.

The direction of a fiber is generally assumed to align with the principal eigenvector of the diffusion tensor acquired by DT-MRI (at least when the principal eigenvalue is much larger in magnitude than the other two eigenvalues); an overview of the process can be found in the summary by Westin [23]. The centerline of the tube follows the integral curve of the principal eigenvector. Vector integration of the

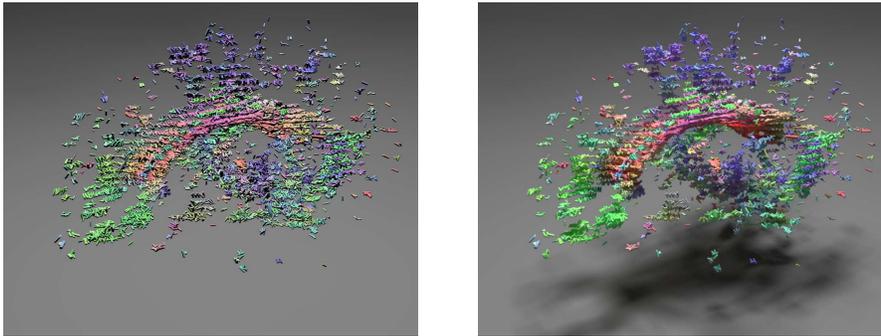


Fig. 6. Superquadric glyphs derived from DT-MRI data. The extent and orientation of the glyphs invite the visual system to integrate their small shapes into larger coherent structures. This visual integration process may be enhanced by global illumination. Left: local illumination. Right: global illumination.

principal eigenvector $v_1(x)$ at point $x = (x[0], x[1], x[2])$ produces a streamline curve $c(t)$ defined by the equation

$$c(t) = \int_0^T v_1(c(t)) dt$$

which can be approximated by standard numerical techniques for vector integration. Fiber tracts are expected to align, subject to some confidence bounds, with these glyphs. Two basic strategies are employed to visualize the fiber tracts: either (1) display discrete glyphs representing either the principal eigenvector or the entire tensor, or (2) display the actual streamlines, whose geometry is expected to follow the tracts.

In strategy (1), the glyphs may be uniformly placed at gridpoints of the 3D volume, or may be packed at nonuniform densities [11]. A glyph may take the shape of an ellipsoid or superquadric to incorporate tensor information about all three diffusion directions.

Glyphs and color. In examining the display of individual 3D glyphs, the surgeon is required to visually integrate the underlying vector field and infer fiber pathways, presumably by estimating streamlines through the glyphs. It is an open question whether illumination contributes to this visual-integration task. Beaudot [2] recently investigated the role of color and curvature on the visual task of linking oriented glyphs across a display in order to infer the presence of a 2D path. This 2D task is simpler than, but similar to, the task of visually linking 3D tensor glyphs to infer fiber tracts. Beaudot measured reaction time where subjects pressed a mouse button after determining whether a path did or did not exist within a field of densely packed 2D oriented glyphs. The study found reaction times increased (worsened) by 100ms when the path was curved rather than straight, and increased by 100ms when the glyph/background combinations possessed saturated values of red/green or blue/yellow rather than achromatic shades of gray varying in luminance.

The current state of 3D visualization of DT-MRI data exhibits, generally speaking, the worst case scenario for path following. Most glyph-based visualization tools display fully saturated glyphs representing curved paths. Figure 6 (left) shows an

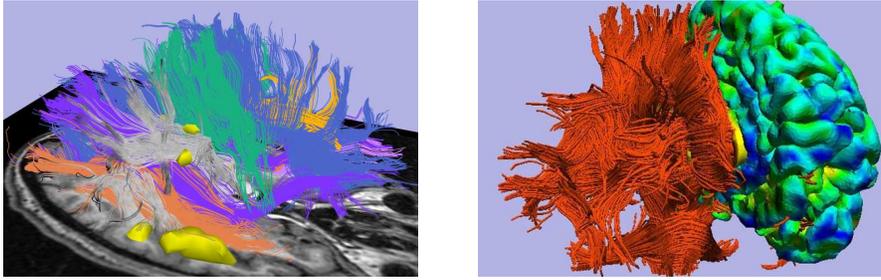


Fig. 7. Representative examples of current state-of-the-art visualization of fiber tracts derived from DT-MRI data. Most of the fibers are occluded by other fibers, requiring the visual system to perform amodal completion of the hidden fiber trajectories. Left: saturated constant-hue curves. Right: polygonal tubes with local illumination (image by H. J. Park).

example image representative of glyphs rendered by a typical visualization tool. An informal sampling of presentations at the recent International Society of Magnetic Resonance in Medicine meeting (ISMRM 2006) revealed that nearly all of the 3D images of diffusion MRI data used fully saturated glyphs/streamlines to represent curving paths, a coloring scheme that Beaudot found to increase visual processing time, or else employed luminance changes resulting from local illumination only; Figure 7 illustrates these two widely-used techniques. Although Beaudot did not measure spatial error in visual path integration, it is plausible that increased visual processing time indicates increased complexity for perception and thus correlates with increased spatial error by the visual system. The presence of shadows resulting from physically based illumination may provide information through the luminance channel to assist the visual system in path-following.

Shape completion. Whether DT-MRI data is displayed using 3D glyphs or 3D curves, the complexity of the resulting geometry practically ensures that most of the glyphs/curves are occluded by others. The visual system is adept at “amodal completion” [14], allowing us to perceive a partially hidden object as complete. Completion seems to result from both local cues (e.g., extrapolating shape direction across the discontinuous edges of the occluder) and global cues [21], and may require 75–200 ms of processing time. What factors influence the visual task of perceiving a 3D curve amidst the clutter of other fragments of geometry in the scene, and how can they be measured?

In experiments by Plomp *et al.* [18], an eye-tracking system was used to avoid the extra processing time involved in a motor task (*e.g.* pressing a mouse button). Subjects were asked to visually search for target shapes that were partly occluded while their eye movements were being recorded. The premise of the study was that completion of an occluded figure correlates with longer gaze duration when the eye fixates on the occluded object. They found that subjects demonstrated longer gaze durations looking at targets than non-targets, about 450–600 ms versus 350–450 ms on average. They found further that a less familiar target (hexagon versus circle) added to the gaze duration, by an increment of about 50 ms. They propose that familiar objects are completed more quickly when occluded. Although no experiments have yet been conducted to determine the effect of illumination on amodal comple-



Fig. 8. Fibers automatically clustered into fascicles. Left: local illumination. Right: global illumination.

tion of fiber tracts surrounded by occluders, it is plausible that 3D tubes subjected to physically-based lighting (as opposed to local illumination or no illumination) are more familiar from daily experience and might therefore be completed more quickly and accurately by the visual system.

Amodal completion of fibers operates at a coarse spatial scale when the surgeon or radiologist analyzes clusters (fascicles) of tracts. In this case, not only do individual fibers occlude each other, but entire functional groups of fibers hide each other. Figure 8 illustrates the situation, where several fascicles have been automatically grouped (by a technique described by O’Donnell [15] [16] and Brun [3]) and then color coded.

4 Results and Acknowledgments

The process of upgrading a visualization software system from local illumination to global illumination can be attacked in various ways. In our case, we wished to introduce global illumination into the *Slicer* visualization tool developed by Harvard and MIT [6] (*Slicer* can be downloaded from <http://slicer.org>). *Slicer* was designed to use VTK, the Visualization Toolkit [20], which manages geometry and renders it using OpenGL (and therefore local illumination). We used VTK’s ability to export geometry in the Open Inventor [22] scene-graph file format so that the 3D scene could be fed into a stand-alone global illumination tool that implements photon mapping [9]. This pipeline of tools allows the 3D scenes from *Slicer* to pass through a post-processing step that generates 2D images of the globally illuminated scene [1]. On a typical desktop computer, the renderer requires 1-100 minutes to generate an image, depending on the number of photons used, the density estimation algorithm applied, and the size of the final image. The initial results of incorporating global illumination into the *Slicer* pipeline can be seen in Figures 1 and 8.

The National Alliance for Medical Image Computing (NA-MIC), which supports and distributes *Slicer*, uses the *Teem* library. *Teem* a group of tools for manipulating and displaying scientific and medical data; it can be found on the Web at <http://teem.sourceforge.net>. We converted *Teem*’s superquadric glyphs into Open Inventor format and rendered the resulting scene graph to produce the images in Figure 6.

Within the computer graphics research community, a considerable amount of effort is being applied to develop real-time systems for global illumination. Our proof-of-concept demonstration offers a preview of what may one day become the standard for 3D visualization systems for analyzing DT-MRI datasets acquired from the human brain.

This work was supported by NSF award CCF 0432133 and NIH NIBIB NAMIC U54-EB005149, NIH NCRR NAC P41-RR13218, NIH NCRR mBIRN U24-RR021382, and R01 MH 50747. The authors express their appreciation to Kevin Beason (Rhythm and Hues Studios), Yoshihito Yagi (Florida State University), and Israel Huff (University of Tennessee) for assistance with rendering and generation of test scenes; Gordon Kindlmann (Laboratory of Mathematics in Imaging, Harvard Medical School) for his DT-MRI dataset and his glyph dataset; Lauren O'Donnell for her expertise with *Slicer* and VTK, and for use of the clustered-fiber dataset; Ion-Florin Talos for use of the ganglioglioma dataset; and Susumu Mori (JHU) for diffusion MRI data (R01 AG20012-01, P41 RR15241-01A1).

References

1. David C. Banks and Kevin Beason. Pre-computed global illumination of mr and dti data. In *International Society for Magnetic Resonance in Medicine 14th Scientific Meeting (ISMRM '06)*, page 531, 2006. Seattle, WA, 6-12 May 2006, Abstract 2755.
2. William H. A. Beaudot and Kathy T. Mullen. Processing time of contour integration: the role of colour, contrast, and curvature. *Perception*, 30:833–853, 2001.
3. A. Brun, H. Knutsson, H. J. Park, M. E. Shenton, and C.-F. Westin. Clustering fiber tracts using normalized cuts. In *Seventh International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI '04)*, pages 368–375, 2004. Rennes - Saint Malo, France.
4. Michael F. Cohen and Donald P. Greenberg. The hemi-cube: A radiosity solution for complex environments. In *Computer Graphics (Proceedings of SIGGRAPH 85)*, pages 31–40, August 1985.
5. Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. In *Computer Graphics (Proceedings of SIGGRAPH 84)*, pages 137–145, July 1984.
6. D. Gering, A. Nabavi, R. Kikinis, N. Hata, L. Odonnell, W. Eric L. Grimson, F. Jolesz, P. Black, and W. Wells III. An integrated visualization system for surgical planning and guidance using image fusion and an open mr. *Journal of Magnetic Resonance Imaging*, 13:967–975, 2001.
7. Donald P. Greenberg, Michael Cohen, and Kenneth E. Torrance. Radiosity: A method for computing global illumination. *The Visual Computer*, 2(5):291–297, September 1986.
8. Jayme Jacobson and Steffen Werner. Why cast shadows are expendable: Insensitivity of human observers and the inherent ambiguity of cast shadows in pictorial art. *Perception*, 33:1369–1383, 2004.
9. Henrik Wann Jensen. *Realistic Image Synthesis Using Photon Mapping*. AK Peters, 2001.

10. James T. Kajiya. The rendering equation. In *Computer Graphics (Proceedings of SIGGRAPH 86)*, pages 143–150, August 1986.
11. Gordon Kindlmann and Carl-Fredrik Westin. Diffusion tensor visualization with glyph packing. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization / Information Visualization 2006)*, 12(5), September–October 2006.
12. C. Madison, W. Thompson, D. Kersten, P. Shirley, and B. Smits. Use of interreflection and shadow for surface contact. *Perception and Psychophysics*, 63:187–194, 2001.
13. P. Mamassian, D. C. Knill, and D. Kersten. The perception of cast shadows. *Trends in Cognitive Sciences*, 2:288–295, 1998.
14. A. Michotte, G. Thines, and G. Crabbe. Les compléments amodaux des structures perceptives. In G. Thines, A. Costall, and G. Butterworth, editors, *Michotte's Experimental Phenomenology of Perception*, pages 140–167, Hillsdale, NJ, 1991. Lawrence Erlbaum Associates.
15. Lauren O'Donnell, Marek Kubicki, Martha E. Shenton, Mark E. Dreusicke, W. Eric L. Grimson, and Carl-Fredrik Westin. A method for clustering white matter fiber tracts. *American Journal of Neuroradiology (AJNR)*, 27(5):1032–1036, 2006.
16. Lauren O'Donnell and Carl-Fredrik Westin. White matter tract clustering and correspondence in populations. In *Eighth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI '05)*, Lecture Notes in Computer Science, pages 140–147, 2005. Palm Springs, CA, USA.
17. Y. Ostrovsky, P. Cavanagh, and P. Sinha. Perceiving illumination inconsistencies in scenes. Technical report, Massachusetts Institute of Technology, Cambridge, MA, USA, November 2001.
18. Gijs Plomp, Chie Nakatani, Valérie Bonnardel, and Cees van Leeuwen. Amodal completion as reflected by gaze durations. *Perception*, 33:1185–1200, 2004.
19. Ronald A. Rensink and Patrick Cavanagh. The influence of cast shadows on visual search. *Perception*, 33:1339–1358, 2004.
20. Will Schroeder, Ken Martin, and Bill Lorensen. *Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*. Kitware, 2006. 4th edition.
21. A. B. Sekuler, S. E. Palmer, and C. Flynn. Local and global processes in visual completion. *Psychological Science*, 5:260–267, 1994.
22. Josie Wernecke and the Open Inventor Architecture Group. *The Inventor Mentor: Programming Object-Oriented 3D Graphics with Open Inventor*. Addison-Wesley Professional, 2004.
23. C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualization of diffusion tensor MRI. *Medical Image Analysis*, 6(2):93–108, 2002.

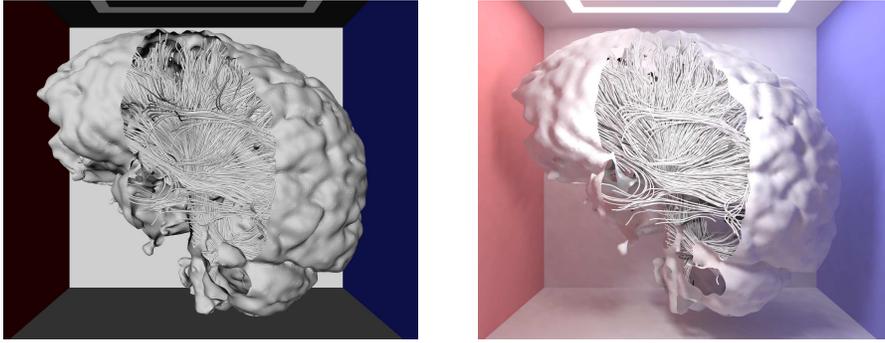


Fig. 4. White matter fibers within a cut-away approximation of the cortex. Left: local illumination. Right: global illumination. Brain dataset courtesy of Gordon Kindlmann at the Scientific Computing and Imaging Institute, University of Utah, and Andrew Alexander, W. M. Keck Laboratory for Functional Brain Imaging and Behavior, University of Wisconsin-Madison.

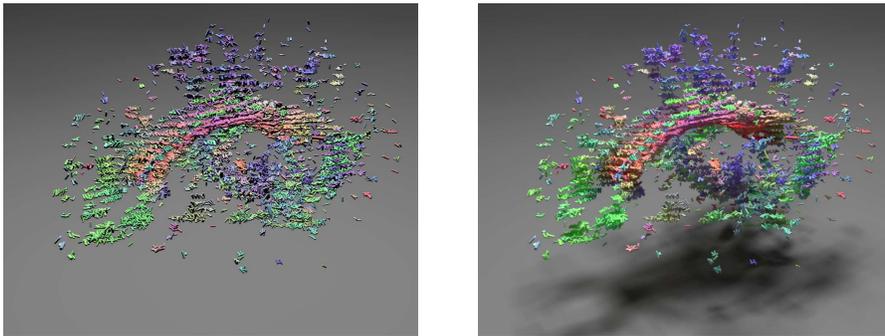


Fig. 6. Superquadric glyphs derived from DT-MRI data. The extent and orientation of the glyphs invite the visual system to integrate their small shapes into larger coherent structures. This visual integration process may be enhanced by global illumination. Left: local illumination. Right: global illumination.



Fig. 8. Fibers automatically clustered into fascicles. Left: local illumination. Right: global illumination.