Two-Handed Spatial Interface Tools for Neurosurgical Planning

John C. Goble, Ken Hinckley, Randy Pausch, John W. Snell, and Neal F. Kassell
University of Virginia

A computer-based neurosurgical planning system lets neurosurgeons easily manipulate 3D data using their everyday skills from handling tools with two hands. It has been tested in actual surgical procedures.

Despite tremendous innovations in medical imaging and virtual reality technologies, real-world applications of virtual reality in clinical medicine remain scarce. Liability and medical licensing concerns, physician training issues, and the catastrophic effects of system failure have slowed virtual reality's diffusion into routine medical care. Moreover, technology's role in medicine—a climate dominated by cost containment and increased demands on the health care delivery system—is undergoing unprecedented scrutiny.

Nevertheless, the potential applications of virtual reality in the medical environment continue to motivate research and development in both academic and commercial settings. Prototype application areas include medical training, minimally invasive diagnosis and surgery, and remote therapeutic intervention techniques. However, most physicians will not use a system that does not convincingly improve their ability to safely and efficiently deliver medical services of the highest quality. Therefore, how can we introduce the technologies of virtual reality into medicine? How can we satisfy this requirement for safety and utility with rapidly evolving hardware and software?

Our approach has been to exploit the tools and techniques developed in the virtual reality community to build safe and reliable systems for the preoperative planning of complex neurosurgical procedures. In this spirit, we have designed and implemented a computer-based neurosurgical planning system, called Netra. The system includes a three-dimensional interface tailored toward surgeons' skills. We have used high-performance graphics workstations, high-speed networking hardware, and six-degree-of-freedom magnetic tracking devices to deliver safe, efficient surgical care to patients with neurological diseases. In fact, Netra has been used for various precision, computer-assisted surgical procedures. Neurosurgeons use Netra to plan precision biopsies, laser-guided tumor resections, surgery for Parkinson's Disease and other motor disorders, and surgical implantation of electrode arrays for epilepsy. Nevertheless, the techniques we've developed for the planning and simulation of neurosurgical interventions may also have important applications beyond the realm of medicine.

THE APPLICATION DOMAIN: NEUROSURGERY

Neurosurgery is inherently a three-dimensional activity. It deals with complex structures in the brain and spine that overlap and interact in complicated ways. The neurosurgeon must visualize these structures and understand the consequences of a proposed surgical intervention to both the pathology (the "target") and the surrounding, viable tissues. Surgical instruments can then be reliably guided to these targets with millimeter accuracy, through methods known as *stereotactic techniques*.[1]

In planning surgery, neurosurgeons traditionally use two-dimensional slices from scanning techniques such as computed tomography (CT) and

0018-9162/95/$4.00 © 1995 IEEE

magnetic resonance imaging (MRI). These 2D slices are further restricted to appear in planes orthogonal to canonical axes through the patient's head. These standard views, known as the axial, coronal, and sagittal planes, represent the frame of reference the physician uses in learning and understanding the anatomy. But many structures within the brain—and several surgical paths to these structures that are clinically useful—are oblique to these canonical views. Development of a 3D anatomical model from these 2D slices remains challenging, even for experienced neurosurgeons.

Surgeons have become increasingly interested in computer-based surgical planning systems that let them quantify and visualize the 3D information available from medical imaging studies. By more effectively using this information and letting the surgeon quickly and intuitively access it, computer-based visualization and planning systems can positively impact both cost of care and patient outcome.

A key element in our system, Netra, drawn directly from the technologies normally associated with virtual reality, is a 3D interface based on the surgeon's ability to manipulate real-world tools with two hands. This *props-based interface*[2] gives the neurosurgeon natural access to complicated planning software (see Figure 1).

A typical computer-assisted surgery includes the following elements:

- *Stereotactic frame placement.* A frame providing a reference coordinate system is attached to the patient's skull before imaging. This frame carries fiducial indicators, which permit precise localization of pathology, and mechanical guides, which are used during the actual procedure.
- *Medical image acquisition.* Due to its superior soft-tissue resolution and capability for true 3D data acquisition, we principally depend on MRI for anatomic information. We may supplement this data with digital subtraction angiography images, which superbly capture the brain's vascular structures. CT is sometimes employed for bony detail.
- *Image segmentation and classification.* We employ various computer algorithms to delineate pathology and other surgically important brain structures. These may include major blood vessels, cranial nerves, and other structures that the surgeon wishes to avoid.
- *Surgical planning.* The surgeon uses planning tools to select the surgical targets and a path to those targets that produces the least possible damage to viable tissue. This is the surgical trajectory. Due to the surgical instrumentation, our trajectories are linear, but piecewise linear and completely arbitrary trajectories are possible.
- *Implementation.* Surgeons can modify the plan intraoperatively using a workstation located in the operating room. They can select additional targets or modify the proposed trajectory based on information not visible in the medical images.

Because the brain changes with time, preparatory steps must be performed during a three- or four-hour window on the morning of surgery. Since the principal neurosurgeon might be caring for several other patients, the actual time available for planning can be as little as fifteen minutes. (See Goble et al.[3] for a detailed description of how our computer-based planning systems and the asynchronous transfer mode (ATM) network infrastructure that we have employed handle these real-time constraints.)

The user interface to the planning system must let the surgeon work quickly. The morning of surgery is perhaps the least optimal time for a surgeon to be fussing with a maze of slider bars and command prompts. Furthermore, the surgeon must cope with frequent distractions and be able to quickly detach from the user interface, both physically and cognitively. Therefore, the interface must not employ devices that are difficult to put down or use explicit modes that are easily forgotten.

RELATED WORK

Several commercially available packages support 3D neurosurgical planning. A representative system is the Viewing Wand,[4] made by ISG of Toronto. Although the Viewing Wand is actually intended for intraoperative rather than preoperative use, it does include surgical planning software that incorporates 3D planning capabilities. The user interface to the system has a mechanical six-degree-of-freedom arm, which the surgeon can use in the operating room to indicate surgical trajectories relative to the actual patient. Our interaction techniques, on the other hand, focus on providing tools that are used outside the operating room before surgery begins.

In another exciting approach to the planning problem, images can be acquired during surgery. A prototype magnetic resonance imager of open design, through which the surgeon can access the patient and conduct contemporaneous imaging, has recently been installed at the Brigham and Women's Hospital in Boston. This approach permits direct feedback of trajectory information—since the surgeon's tools are visible in the image—and lets surgeons view changes occurring in the brain during the surgery. However, this imaging device is extremely expensive, and most surgical tools must be redesigned to prevent imaging artifacts. Moreover, this approach does not focus on user interface issues. (Kikinis et al.[5] describe early results obtained with this instrument.)

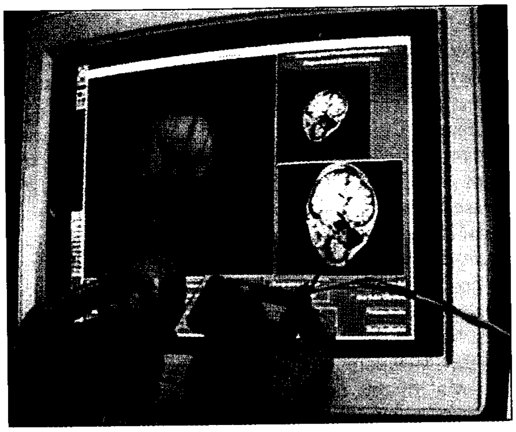

Figure 1. A user views a cross section of a brain via the two-handed interface tools.

Even outside the medical imaging area, few virtual reality or desktop virtual reality interfaces let the user interact with real-world tools using two hands. The 3-Draw computer-aided design tool[6] requires the user to hold a stylus in one hand relative to a tablet held in the other. The user moves the stylus relative to the tablet to draw 3D curves, and the user can rotate the tablet to view the object being drawn. However, the tasks supported by our neurosurgical planning system differ substantially from those supported by 3-Draw.

MEDICAL USER INTERFACES

Software usability is crucial to getting neurosurgeons to actually employ advanced visualization software in the clinical routine. The manipulative capabilities of common input devices such as mice and keyboards are poorly matched to the volumetric manipulation and visualization tasks of the neurosurgeon. Rather than typing commands or moving sliders with a mouse, the neurosurgeon is accustomed to seeing a real patient's head and operating on it with surgical instruments. In fact, regarding an idealized interface, one surgeon remarked, "I want a skull I can hold in my hand."

Netra does not track the user's head, generate stereoscopic images, or require the user to wear a head-mounted display.

We have, therefore, pursued a 3D interface based on the physical manipulation of real-world, handheld tools, which we call *props*. The interface employs a head-viewing prop, cross-sectioning prop, and trajectory-selection prop, to rotate, cross-section, and point to a computer-generated 3D image of the patient's head (which we call the virtual head). The head-viewing prop is a small doll's head that can be comfortably held in one hand; when the user rotates the doll's head, the virtual head rotates correspondingly. The cross-sectioning prop is a clear plastic plate (see Figure 2) that can be held to the doll's head to indicate a desired cross section through the virtual head. Finally, the trajectory-selection prop is a handheld stylus for pointing to the doll's head to indicate a desired approach to a surgical target in the virtual head (see Figure 3).

Many virtual reality systems are characterized by their use of a stereo head-tracked, head-mounted display. Netra, however, does not track the user's head, generate stereoscopic images, or require the user to wear a head-mounted display. Most physicians believe the current head-mounted display technology is too encumbering and too limited in resolution for adequate viewing of complex medical data.

Therefore, in Netra, users manipulate virtual objects seen on a standard workstation monitor by moving the props with their hands. Since many people associate the phrase "virtual reality interface" with immersing head-mounted displays, we often characterize our system as a *spatial desktop interface*—spatial because it involves moving six-degree-of-freedom input sensors in free space, desktop because it uses a standard monitor on the user's desk. Nevertheless, our system includes the qualities of real-time viewpoint change, real space, and real interaction via direct manipulation—characteristics that researchers have used to informally define virtual reality.

Users can change the graphics displayed on the screen by moving the props, but they do not have to keep their hands in a fixed position relative to the monitor or the graphics. Other desktop systems, such as Deering's virtual lathe,[7] require users to hold a hand up to a miniature lathe, projected stereoscopically in front of the monitor. This forces them to keep their hand in a small, fixed volume of space, which can be fatiguing with prolonged use. Surgeons using Netra's props-based interface can move their hands in whatever working space they find comfortable, since the center of the working volume is defined by the doll's head held in the user's nondominant hand. Hence, users can shift their body posture over time to work in whatever position they find comfortable, thereby reducing fatigue.[8]

An underlying hypothesis in Netra's design was that the visual and haptic (related to the sense of touch) cues of the spatial desktop interface would let surgeons directly transfer skills involved in everyday bimanual tasks to manipulate virtual 3D medical images, with little or no need for training. Our informal evaluations of over 50 physicians and more than 1,000 nonphysicians have shown that with a cursory introduction, people who have never before seen the interface can understand and use it within about one minute of touching the props. This suggests that our underlying hypothesis is sound and that interacting with virtual objects via two-handed manipulation of props works well not only for neurosurgeons but also for a wide range of potential users.

We have developed our 3D surgical planning system with the intention of making the computer a digital tool that fits into the surgeon's existing paradigm and complements their abilities. Our work, therefore, has been heavily collaborative, relying on the advice and opinions of neurosurgeons to provide goals and specifications throughout the design process.

INTERFACE OVERVIEW

To monitor prop locations, we use a commercially available six-degree-of-freedom tracking system, the Fastrak, which was manufactured by Polhemus Navigation Systems. Each prop is instrumented with a small magnetic receiver, which generates a signal in response to pulsed magnetic waves emitted by a nearby magnetic-field transmitting box. Fastrak processes the signals generated by each receiver to determine its location (x, y, z) and orientation (yaw, pitch, roll) relative to the transmitting box. Fastrak returns this information to the host computer via an RS-232 serial port connection.

As the user rotates the doll's head, the virtual head (on-screen image) is automatically updated to match its orientation. Surgeons can control the image zoom factor by moving the doll's head toward or away from their body. (Of course, the zoom factor of the image is actually computed from the distance between the sensor and the transmitter, since the software has no information about the user's position.)

Since translating the head from left to right or up and down is typically not useful, we have constrained the (x, y) position of the polygonal brain to the center of the screen. This reduces a task's dimensions, and surgeons find it natural.

The virtual head is updated approximately 15 times per second, so from the user's perspective, the motions of the doll's head and the virtual head appear to be tightly

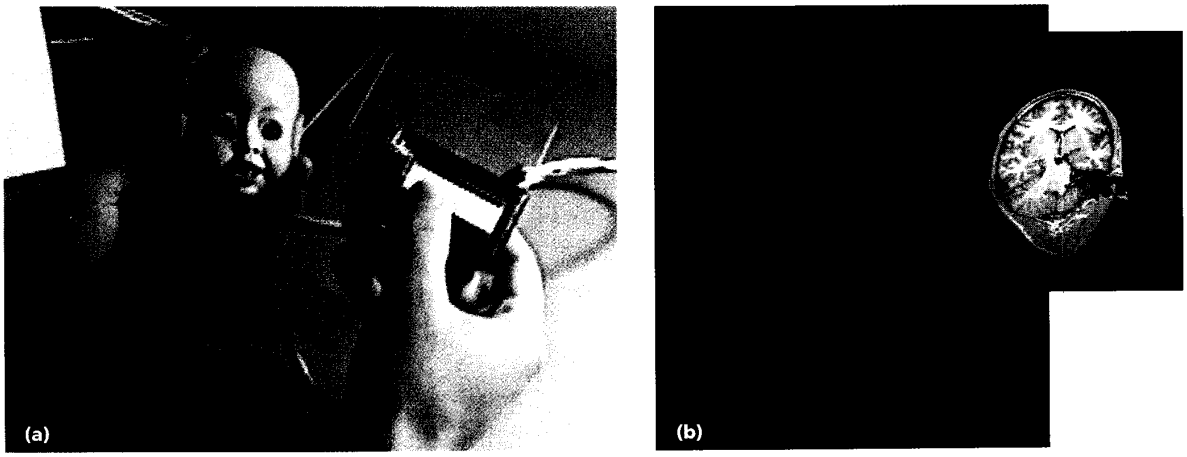

Figure 2. Cross-section selection: (a) user indicates a desired cross section by holding a clear plastic plate to the doll's head; (b) computer shows a corresponding virtual tool intersecting the virtual head, along with a cross section of the volumetric brain data (inset).

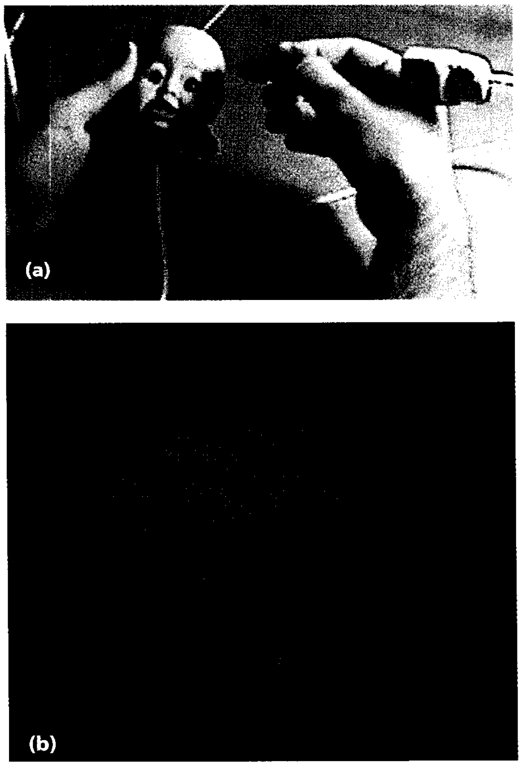

Figure 3. Trajectory selection: (a) user selects a trajectory by pointing to the doll's head with a stylus; (b) computer shows a corresponding virtual probe intersecting the virtual head.

coupled. Ideally, as surgeons rotate the doll's head, they would like to see a detailed volume-rendered image of the patient's head (see Figure 4). Unfortunately, a sufficiently detailed model, volume rendered with transparency and other attributes, requires approximately three seconds to render. This is too slow for the user to maintain cognitive

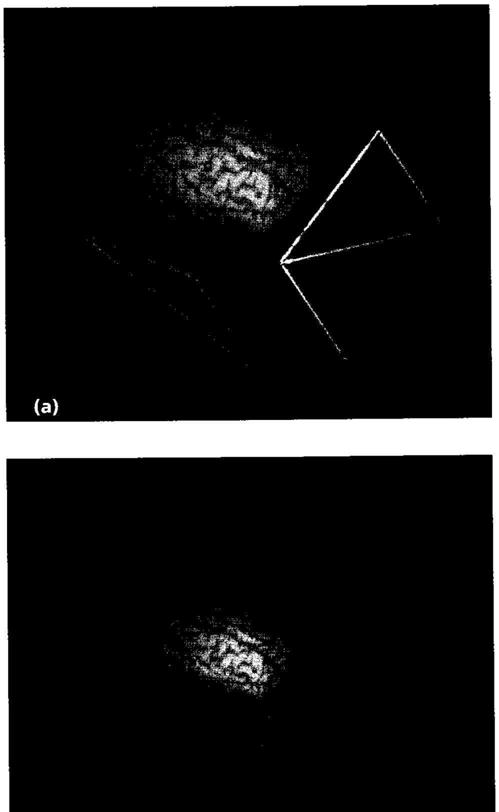

Figure 4. Image transformation from computer to stereotactic-frame coordinate system: (a) the image data's *N*-shaped fiducial patterns specify its coordinate system relative to the stereotactic frame's coordinate system, allowing presurgical plans to be transferred to the operating room; (b) 3D rendering of the frame, and the positioning arc.

Figure 5. This sequence shows the user positioning the cross-sectioning prop, over a period of a few seconds, to expose the optic nerves. The optic nerves are very difficult to visualize if the surgeon has only a series of 2D orthogonal slices available for viewing.

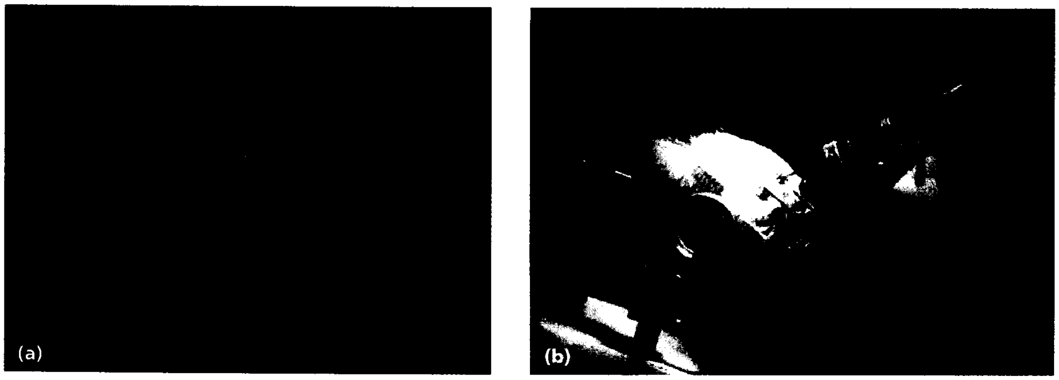

Figure 6. A needle biopsy: (a) the proposed surgical plan; (b) the actual procedure.

coupling between movement of the interface props and the resulting movement of the patient's head data on screen.

To allow interactive rotation of the virtual head, we render an image of a polygonal approximation to the patient's brain (as shown in Figure 3b), which our hardware graphics accelerator can accomplish very quickly. When the surgeon stops moving or issues an explicit command, the system automatically generates the detailed volume-rendered image. In this way, the surgeon can manipulate the 3D data very quickly and still be able to inspect it at full resolution, as desired.

The cross-sectioning prop (a clear plate) is used in concert with the head prop (the doll's head), rather than as a separate tool. As the surgeon moves the plate relative to the doll's head, the computer shows the corresponding cross-section of the patient's MRI volume. This lets the surgeon explore the entire volumetric data set quickly and at any orientation. Structures such as the optic nerves, which have been difficult to visualize based on the paradigm of viewing orthogonal slices on film, can now be easily found and inspected (see Figure 5).

The trajectory selection prop is also used in concert with the head prop. The user points the stylus to the doll's head, thereby positioning and orienting a cylindrical virtual probe relative to the polygonal brain model. The stylus doubles as a tool for selecting points on the exposed brain surface, since the computer can calculate the intersection of the surface data with the 3D vector formed by the stylus.

One could argue that using two hands to operate the interface only adds complexity and makes it harder to use. However, we maintain (as demonstrated by Kabbash[9]) that using two hands does not necessarily impose a cognitive burden, and can help users reason about their tasks. The work of Guiard[10] is enlightening in this regard. Guiard argues that humans use the right and left hands to control frames of reference that are serially linked. For right-handers, the left hand specifies a base frame of reference relative to which the right hand expresses a second active frame of reference. Thus, in terms of human perception, the hands do not operate independently and in parallel, but in conjunction by assuming asymmetric roles.

Our two-handed interface assigns the base frame of reference to the doll's head and the active frame of reference to the cross-sectioning plate. Since the surgeon's task is to specify a cutting plane relative to a particular view of the virtual head, our frames-of-reference assignment matches the surgeon's perceptual model of the task. This results in an easily understood two-handed interface.

3D IMAGE SEGMENTATION

Effective visualization and manipulation of 3D head images requires segmenting the volume data into anatomically meaningful volumes of interest. These volumes define surfaces that can be efficiently rendered as 3D perspective views; they can further be used to compute quantitative data properties such as volume, surface area, and compactness. Location of blood vessels, for instance, is important in selecting trajectories that minimize vascular damage to the patient.

Our semiautomatic segmentation algorithm incorporates a priori anatomical knowledge. This permits accurate definition of the brain's boundaries and major substructures from 3D MRI, despite low-contrast areas in the data and individual anatomical variations. We represent this anatomical knowledge in two formats: A brain surface model encodes the gross brain surface, while voxel (3D pixel)-based models encode typical brain shape, gray/white composition, and ventricle shape. The entire segmentation process requires about 10 to 15 minutes on a Hewlett-Packard 9000/735 workstation, including all necessary user interaction.[3, 11]

TRANSFERRING PRESURGICAL PLANS TO THE OPERATING ROOM

We use a commercially available stereotactic frame system known as the Leksell microstereotactic system.[12] The coordinate system defined by this frame is called Leksell space. A positioning arc is attached to the frame. By mounting a biopsy needle to this arc (see Figure 4), for example, a surgeon can localize any point in the head to within about one millimeter.

Proceeding from a presurgical plan to the actual patient in the operating room requires transforming the coordinate system I of the volumetric image data to the coordinate system L of Leksell space. This is accomplished by imaging the patient on the morning of surgery with a Leksell frame that has been attached to the patient's head and fitted with a special fiducial system. The fiducial box contains N-shaped tubes filled with a contrast agent, which can be clearly seen as N-shaped patterns in the volumetric image data (see Figure 4a). The N-shaped patterns unambiguously specify the relationship—which can be algorithmically derived—of the I and L coordinate systems.

Figure 6 depicts a needle biopsy, where the simulated plan shown on the left demonstrates the degree of realism available to the neurosurgeon during planning.

WE HAVE WORKED CLOSELY WITH THE NEUROSURGEONS throughout the design of this three-dimensional neurosurgical planning system and associated 3D user interface. Our focus has always been to support the neurosurgeon's preoperative clinical needs. In this context, we have tested Netra by using it to plan actual surgical procedures involving real patients.

The 3D user interface exploits the neurosurgeon's existing skills for manipulating tools with two hands by providing real-world tools, or props, which can be moved relative to one another in free space. This lets neurosurgeons control complex visualization software using skills they already have for manipulating physical objects with two hands.

Neurosurgeons have been enthusiastic about Netra's 3D planning capabilities and the ease of use of the props-based interface. The interface, coupled with powerful computer and network capabilities and robust algorithms for tissue classification, has encouraged our surgeons to develop new techniques, which they would not have previously attempted because of their inability to visualize and simulate the surgical trajectories required. In fact, various components of the Netra system have been used for approximately 40 patients in clinical cases involving epilepsy, brain tumors, or neurological-based motor disorders.

Our experience, though preliminary, indicates that sophisticated, technology-based systems such as Netra can significantly contribute to the surgical management of patients with certain neurological diseases. We also have anecdotal evidence that our neurosurgeons are now attempting more difficult tumor resections, which they would not have performed without precision trajectory information. We will continue to gather patient outcome information and compare our results with those obtained through traditional surgical approaches.

A key element in the acceptance of Netra by practicing physicians has been their active participation in the software design. Virtual reality applications must safely and efficiently solve real-world problems. Successful introduction of virtual reality techniques into medicine depends on careful collaboration and an appreciation of an evolving technology's limitations. ▮

References

1. P.J. Kelly and B.A. Kall, eds., *Computers in Stereotactic Neurosurgery*, Blackwell Scientific Publications, Cambridge, Mass., 1992.
2. K. Hinckley et al., "Passive Real-World Interface Props for Neurosurgical Visualization," *Proc. ACM CHI 94 Conf. Human Factors in Computing Systems*, ACM, New York, 1994, pp. 452-458.
3. J. Goble et al., "A Real-Time System for 3D Neurosurgical Planning," *Proc. Visualization in Biomedical Computing 1994*, Int'l Soc. for Optical Eng., Bellingham, Wash., 1994, pp. 552-573.
4. J. Drake, J. Rutka, and J. Hoffman, "The ISG Viewing Wand," *Neurosurgery*, Vol. 34, 1994, pp. 1,094-1,097.
5. R. Kikinis et al., "Image Guidance Techniques for Neurosurgery," *Proc. Visualization in Biomedical Computing 1994*, Int'l Soc. for Optical Eng., Bellingham, Wash., 1994, pp. 537-540.
6. E. Sachs, A. Roberts, and D. Stoops, "3-Draw: A Tool for Designing 3D Shapes," *IEEE Computer Graphics and Applications*, Vol. 11, No. 6, Nov. 1991, pp. 18-26.
7. M. Deering, "High-Resolution Virtual Reality," *Computer Graphics*, Vol. 26, No. 2, July 1992, pp. 195-202.
8. K. Hinckley, et al., "A Survey of Design Issues in Spatial Input," *Proc. ACM UIST 94 Symp. User Interface Software and Tech.*, ACM, New York, 1994, pp. 213-222.

9. P. Kabbash, W. Buxton, and A. Sellen, "Two-Handed Input in a Compound Task," *ACM CHI 94 Conf. Human Factors in Computing Systems*, ACM, New York, 1994, pp. 417-423.
10. Y. Guiard, "Asymmetric Division of Labor in Human-Skilled Bimanual Action," *J. of Motor Behavior*, Vol. 19, No. 4, 1987, pp. 486-517.
11. J.W. Snell et al., "Model-Based Boundary Estimation of Complex Objects Using Active Surface Templates," *Pattern Recognition*, in press.
12. L. Leksell, "A Stereotactic Apparatus for Intracerebral Surgery," *Acta Chirurgica Scandinavica*, Vol. 99, 1949, pp. 229-233.

John C. Goble *is an assistant professor in the Department of Neurosurgery at the University of Virginia and director of the Neurosurgical Visualization Lab. The lab builds hardware and software tools to support academic neurosurgeons' clinical and research needs. He has an ScM in radiological physics from Johns Hopkins University and a PhD in computer science from Rensselaer Polytechnic Institute.*

Ken Hinckley *is a doctoral student at the University of Virginia, where he obtained a masters degree in computer science in 1993. His research focuses on enhancing user interfaces for manipulation and visualization of complex volumetric data, with a special interest in interface issues for medical imaging applications.*

Randy Pausch *is an associate professor of computer science at the University of Virginia. His primary interests are human-computer interaction and undergraduate education. He received a BS in computer science from Brown University in 1982 and a PhD in computer science from Carnegie Mellon University in 1988. He is a National Science Foundation Presidential Young Investigator and a Lilly Foundation Teaching Fellow. He is also a member of the ACM and IEEE.*

John W. Snell *is an assistant research professor in the Department of Neurosurgery at the University of Virginia. His research interests focus on the segmentation of medical imagery using knowledge-based approaches. He received a BS in biomedical engineering from Duke University in 1987 and a PhD in biomedical engineering from the University of Virginia in 1993.*

Neal F. Kassell *is a distinguished professor and vice chair of the Department of Neurosurgery at the University of Virginia. Kassell's clinical practice specializes in patients with neurovascular disease, including aneurysms and arteriovenous malformations. His research interests focus on cerebral vasopasm and subarachnoid hemorrhage. He also has a strong interest in computing and information technology. He received undergraduate and medical degrees from the University of Pennsylvania and completed his neurosurgery training at the University of Western Ontario.*

Readers can contact the authors at the Neurosurgical Visualization Laboratory, University of Virginia, 853 West Main St., Suite 102, Charlottesville, VA 22903; e-mail {goble, kph2q, pausch, jws2f}@virginia.edu.