

12. Iwata, H., "Artificial Reality with Force-feedback: Development of Desktop Virtual Space with Compact Master Manipulator," *Computer Graphics*, 24 (4), pp. 165-170.
13. Kabbash, P., MacKenzie, I. S., Buxton, W., "Human Performance Using Computer Input Devices in the Preferred and Non-Preferred Hands," *Proc. ACM INTERCHI'93 Conference on Human Factors in Computing Systems*, pp. 474-481.
14. Kaufman, A., Yagel, R., "Tools for Interaction in Three Dimensions," *Proc. 3rd International Conference on HCI*, Vol. 1 (Boston, MA, September 1989), pp. 468-475.
15. Kaufman, A., Yagel, R., and Bakalash, R., "Direct Interaction with a 3D Volumetric Environment," *Proc. 1990 Symposium on Interactive 3D Graphics*, *Computer Graphics*, 24 (2), pp. 33-34.
16. Kelly, P. J., Kall, B. A., eds., Computers in Stereotactic Neurosurgery, Blackwell Scientific Publications Inc., Cambridge, MA, 1992.
17. Kikinis, R., Cline, H. E., Altobelli, D., Halle, M. W., Lorensen, W. E., Jolesz, F. A., "Interactive Visualization and Manipulation of 3-D Reconstructions for the Planning of Surgical Procedures," *Visualization in Biomedical Computing 1992*, Robb, R. A., ed., *Proc. SPIE 1808*, October 1992, pp. 559-563.
18. McKenna, M., "Interactive Viewpoint Control and Three-dimensional Operations," *Proc. 1992 Symposium on Interactive 3D Graphics*, pp. 53-56.
19. Nielsen, J., Usability Engineering, Academic Press, Inc., San Diego, CA, 1993.
20. Norman, D., The Design of Everyday Things, Doubleday, New York, NY, 1990.
21. Osborn, J., Agogino, A., "An Interface for Interactive Spatial Reasoning and Visualization," *Proc. ACM CHI'92 Conference on Human Factors in Computing Systems*, pp. 75-82.
22. Pieper, S., Rosen, J., Zeltzer, D., "Interactive Graphics for Plastic Surgery: A Task-level Analysis and Implementation," *Proc. 1992 Symposium on Interactive 3D Graphics*, pp. 127-134.
23. Polhemus Navigation Sciences, Inc., P. O. Box 560, Colchester, VT 05446. (802) 655-3159.
24. Robb, R. A., "A Software System for Interactive and Quantitative Analysis of Biomedical Images," in Hohne, K. H., Fuchs, H., and Pizer, S., eds., 3D Imaging in Medicine, NATO ASI Series, 1990, Vol. F 60, pp. 333-361.
25. Sachs, E., Roberts, A., Stoops, D., "3-Draw: A Tool for Designing 3D Shapes," *IEEE Computer Graphics and Applications*, Nov. 1991, pp. 18-26.
26. Sears, A., Plaisant, C., Shneiderman, B., "A New Era for High Precision Touchscreens," in Hartson and Hix, eds., Advances in Human-Computer Interaction, Vol. 3, 1992, pp. 1-33.
27. Suetens, P., Vandermeulen, D., Oosterlinck, A., Gybels J., Marchal, G., "A 3-D Display System with Stereoscopic, Movement Parallax and Real-time Rotation Capabilities," *Proc. SPIE-Medical Imaging II: Image Data Management and Display (Part B)*, Vol. 914, Jan. 1988, pp. 855-861.
28. Snell, J. W., et al, "Segmentation of the Brain from 3D MRI Using a Hierarchical Active Surface Template," *Proc. of the SPIE Conference on Medical Imaging*, 1994 (to appear).
29. Verbex Voice Systems, Inc., 119 Russell Street, Littleton, MA 01460. 1-800-ASK-VRBX.
30. Ware, C., Arthur, K., Booth, K. S., "Fish Tank Virtual Reality," *Proc. INTERCHI'93 Conference on Human Factors in Computing Systems*, pp. 37-41.

The user should not have to switch modes or constantly pick up and put down separate devices to switch between 2D and 3D input. We believe that use of a touchscreen in combination with the props could provide a solution to this problem. For example, the trajectory prop could be moved freely in space to express 3D interactions, but it could be held against the touchscreen to perform interactions which are constrained to lie in a plane. Recent advances in touchscreen technology have greatly improved touchscreen usability, so the commonly held belief that all touchscreens are inherently inaccurate and useful only for the selection of large buttons is largely apocryphal [30].

There are also promising applications of our interface props for other tasks related to surgical planning. For example, there is a use for props in the operating room. If the real patient and the pre-surgical plan can be registered to one another, then the surgeon could preview and verify the surgical plan by indicating trajectories and cutting planes right in the operating room. Note that the design constraints for the interface change in several important ways:

- The real patient's head is present, so no artificial head prop is needed.
- The surgeon has only one hand available, so both hands cannot be used to manipulate the tools.
- Magnetic trackers cannot be used in the operating room due to the omnipresence of metallic objects. An optical, ultrasound, or mechanical 3D tracker which can track the props as well as the patient's head would be required.
- Constructing the props requires much more design effort, since the tool must be able to withstand sterilization.
- Positioning the computer monitor (which presents the graphical feedback for the prop) becomes a problem, since the patient and other equipment are in the way.

As another example, one could envision using the props as an interface to an MRI scanner. The physician or MR technician could position the cutting plane relative to the head prop to indicate the desired acquisition angle of the data set. The MRI scanner could then compute an appropriate pulse sequence and directly generate the raw data for an oblique slice. Alternatively, the scanner could acquire the data in real time as the user manipulates the cutting plane.

8. REFERENCES

1. Badler, N. I., Manoochehri, K. H., Baraff, D., "Multi-Dimensional Input Techniques and Articulated Figure Positioning by Multiple Constraints," *Proc. 1986 ACM Workshop on Interactive 3D Graphics*, pp. 151-170.
2. Brooks, F. P. Jr., Ouh-Young, M., Batter, J. J., Kilpatrick, P. J., "Project GROPE--Haptic Displays for Scientific Visualization," *Proc. ACM SIGGRAPH '90, Computer Graphics*, 24 (4), 1990, pp. 177-185.
3. Buxton, W., Myers, B., "A Study in Two-Handed Input," *Proc. ACM CHI'86 Conference on Human Factors in Computing Systems*, 1986, pp. 321-326.
4. Buxton, W., "Chunking and Phrasing and the Design of Human-computer Dialogues," in Kugler, H. J., ed., *Information Processing '86, Proc. of the IFIP 10th World Computer Congress*, North Holland Publishers, 1986, pp. 475-480.
5. Fitzmaurice, G. W., "Situated Information Spaces and Spatially Aware Palmtop Computers," *Communications of the ACM*, 36 (7), 1993, pp. 39-49.
6. Hewlett Packard, 175 Waltham Street, Waltham, MA 02254-9030.
7. Hill, D. L., Green, S. E., Crossman, J. E., Hawkes, D. J., Robinson, G. P., Ruff, C. F., Cox, T. C., Strong, A. J., Gleeson, M. J., "Visualization of Multimodal Images for the Planning of Skull Base Surgery," *Visualization in Biomedical Computing 1992*, Robb, R. A., ed., *Proc. SPIE 1808*, October 1992, pp. 564-573.
8. Hinckley, K., Pausch, R., Goble, J. C., Kassell, N. F., "Passive Real-World Interface Props for Neurosurgical Visualization," *Proc. ACM CHI'94 Conference on Human Factors in Computing Systems* (to appear), 1994.
9. Hohne, K., Bomans, M., Riemer, M., Schubert, R., Tiede, U., Lierse, W., "A Volume-based Anatomical Atlas," *IEEE Computer Graphics and Applications*, July 1992, pp. 72-77.
10. Hohne, K. H., Fuchs, H., Pizer, S., eds., *3D Imaging in Medicine*, NATO ASI Series, 1990, Vol. F 60.
11. International Technologies, Inc., 68 Albany Road, Warwick, RI 02888-2152. (401) 781-5595.

There are a number of clear advantages in terms of interface design which naturally follow from using our neurosurgical planning props:

- *Familiarity*: Manipulating real-world objects is a familiar task and exploits existing user skills.
- *Direct Actions*: The user’s actions correspond directly to the user’s task.
- *Obvious Use*: The appearance of a prop indicates its use; no ambiguous icons or widgets are needed.
- *Palpability*: Users are immediately and continuously aware of the physical existence of each prop.
- *No Tool Moding*: The obvious use and palpability of a prop make the user instantly aware of which tools are being used. Behaviors do not have to be “selected” because one input device is not being overloaded for widely varying tasks.
- *Feedback*: Props provide real-world visual, tactile, and kinesthetic feedback.
- *Two-handed Interaction*: Users will naturally use both hands to manipulate real objects. This also allows users to specify actions in terms of high-level “cognitive chunks” (fig. 7) that single props cannot as easily express.
- *Pragmatics*: Designers can provide familiar tools with physical constraints appropriate to the user’s tasks, rather than constructing an interface which is limited by the unfamiliar physical packaging of the underlying tracker technology.
- *New Tool Uses*: Users will apply tools which exist in the real world in unforeseen and creative ways.

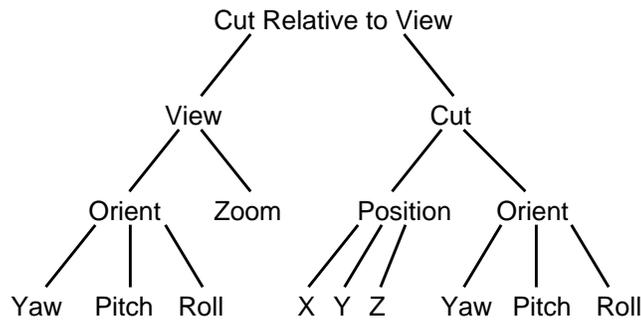


Figure 7: Task Hierarchy for Selecting a Cut Relative to a Specific View.

Figure 7 illustrates how the props based interface simplifies the apparent complexity of selecting a cutting-plane relative to a specific view of the polygonal brain. Cutting relative to a view consists of two sub-tasks: viewing and cutting. Viewing can further be subdivided into orienting the brain and specifying a zoom factor, and so forth. At the lowest level, there are ten separate parameters (yaw, pitch, roll, and zoom for the view; x, y, z, yaw, pitch, and roll for the cutting tool) being specified. In a sliders or knob-box implementation of this interface, the user would have to perform ten separate one-dimensional tasks to position the cutting plane relative to a view, resulting in a non-intuitive user interface. Using the props with both hands, however, reduces this entire hierarchy into a single transaction (or “cognitive chunk”) which directly corresponds to the task that the user has in mind. As a result the user perceives the interface as being much easier to use.

Neurosurgical planning is an ideal application for props because neurosurgeons work and think in terms of real objects in real space. After a cursory introduction, neurosurgeons can understand and use the interface we have described without difficulty and without training. We believe that no other medical imaging system, particularly those based on traditional input devices, can make this claim.

7. FUTURE WORK

So far we have not discussed the problem of designing the “total interface” for a medical imaging application. We have focused on 3D manipulation, but there are many tasks of interest which are inherently one or two-dimensional. It is important to recognize that using a 3D input device to control a task which is not three-dimensional can lead to an obtuse interface. Thus, there is a need for both 3D and 2D input devices in medical imaging applications, and the use of these devices must be seamlessly integrated in a hybrid 2D and 3D interface.

5.3. Balancing High Speed Interaction and Anatomical Detail

Fluid interaction is essential for the manipulation techniques we have described. Ideally we would like to provide interactive manipulation of the entire MRI data set, but generating a volume-rendered image at 25 frames per second is currently not feasible.

If the image takes too long for the computer to draw, the motion of objects on the screen looks choppy and may even disorient the user. Furthermore, the delay between the user's movements and the update of the image (often called latency or lag) is increased, resulting in a degraded quality of interaction. For good user performance, image updates of at least 15 to 20 frames per second with latencies under 100 milliseconds are needed.

Unfortunately, full-resolution 3D visualizations cannot be generated at these speeds using the display algorithms and mid-range workstations available today; update times of several seconds or minutes are more typical for commercially available volume renderers. This means that the computer must display the image at a reduced resolution while the surgeon manipulates the props; the detailed image can only be rendered after the surgeon has finished.

Our strategy to address this problem has been to render the patient's anatomy using polygons during motion, and then use an in-house volume renderer to generate the final detailed image. The current implementation of our system is capable of rendering polygonal models consisting of approximately 15,000 triangles at interactive frame rates. The volume renderer is capable of generating the detailed images in about five seconds. At present these tools have only been integrated at the level of a crude prototype.



Figure 6: A Neurosurgeon Specifying a Cutting Plane.

6. DISCUSSION

The notion of passively observing a natural user dialog in the real world has guided us throughout the design of our user interface. Our goal is an interface which is so obvious, the user will scarcely be aware that any explicit interface exists. We believe a medical imaging system based on passive observation of real-world interface props can come close to achieving this goal.

4.5. System Hardware Components

Our workstation is a Hewlett Packard HP735 configured with a CRX48Z hardware graphics accelerator [8]. The magnetic tracking system is the “FASTRAK” manufactured by Polhemus, Inc. [27]. Voice input is achieved using a Verbex “Speech Commander” speaker-dependent recognizer [34]. An International Technologies “Sense & Switch” module [14] is used for digital input and output. RS232 serial lines connect these peripheral devices to the workstation.

5. TOWARDS A SYSTEM FOR USE IN CLINICAL ROUTINE

The present work has focused on identifying and implementing a paradigm for 3D interaction with medical images. As it stands, the current tool cannot be used in the clinical routine, because it has not yet been integrated with other necessary subsystems (*fig. 5*). For example, the actual retrieval of patient data must be made invisible to the surgeon, the image data needs to be segmented before a polygonal model of the brain surface can be created, and a volume renderer is needed to generate a detailed image of the patient data. We have separately implemented prototypes of these major software subsystems; in the near future, a full integration of these subsystems will allow neurosurgeons to visualize and manipulate 3D data as part of their daily routine.

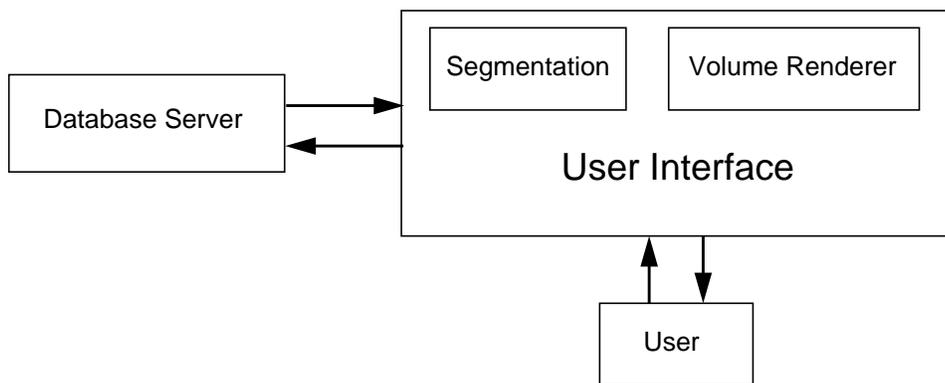


Figure 5: Components of a Neurosurgical Visualization System.

5.1. Accessing Patient Data from a Database

Surgeons should not be responsible for retrieving and maintaining the image data of their patients. Instead, an autonomous server should retrieve the image data. The surgeon can then select all the data for a specific patient from the database, without worrying about image formats or the physical location of the data. We are presently implementing such a database.

5.2. Generating a Polygonal Model of the Brain Surface

Before the surgeon can manipulate a polygonal model of the patient’s brain this model must be generated from the raw MRI image data. A polygonal representation of easily recognized structures such as the skin surface can be generated using crude segmentation techniques such as simple thresholding. The surgeon, however, wants to work with the brain and its substructures in addition to the skin surface, and identifying these structures from raw MRI data is a much more difficult problem, especially for brains which exhibit injury or pathology.

Some of our laboratory’s ongoing work in semiautomatic image segmentation algorithms is described elsewhere in these proceedings [33]. Snell’s active surfaces segmentation algorithm is particularly useful for interactive manipulation since its output consists of a polygonal mesh that can be rendered efficiently. Also, the density of polygons in the mesh can be specified by the user, allowing generation of meshes which are suited to a particular hardware platform’s graphics rendering capabilities.

Presently, segmentation of the image must be done in a separate step before interactive manipulation can begin, but it will soon be possible to perform this task within the framework of the interface. Since the actual segmentation process is almost entirely automated, very little interaction from the user will be necessary.

stage” (fig. 3, right side) where the polygonal brain could be immersed in any one of three static orthogonal cutting surfaces. Much to our surprise, this interface was also poorly received.

It was not until we informally compared the 3D stage interface with the cutting-plane prop selection method that we understood the problem. The neurosurgeon wants to select a cut relative to some specific view of the polygonal brain. The cutting-plane prop can express this concept, whereas the 3D pool-of-water and the 3D stage interfaces cannot. This illustrates why it is important to involve the real users of a system in the design of its interface: the untested intuition of the interface designer is almost always wrong.

4.3. Indicating Surgical Paths with a Trajectory Prop

A trajectory selection prop allows the surgeon to specify three-dimensional vectors and points. The trajectory prop is a stylus-shaped tool which is equipped with a tip switch. Moving the trajectory prop relative to the head prop specifies the position and orientation of a cylindrical virtual probe relative to the polygonal brain model.

In neurosurgery, a trajectory is defined as a three-dimensional path from the exterior of the head to a surgical target inside the brain. A linear trajectory is adequate for simple clinical cases, but often a nonlinear surgical path is required to avoid healthy brain tissue. The present prototype does not yet support nonlinear trajectory selection, although a solution using curves sketched in 3D (as done in 3-Draw [29]) can be envisioned.

A linear trajectory consists of a target point inside the brain volume and a vector from that point. The trajectory selection prop indicates the vector by its orientation relative to the head prop. The target of the trajectory is indicated by the intersection of a ray cast from the virtual probe and the brain model’s surface. When the trajectory prop’s tip switch is held against the head prop, the software enters a “constrained” mode which causes the tip of the virtual probe to be pegged to the intersection point. Since entering and exiting this constrained mode is controlled by a physical gesture (pressing the trajectory against the head prop), the user does not perceive this action as an explicit mode and the user cannot become “trapped” in the mode. The use of gestural phrasing has been advocated in previous work by Buxton [4].

The interior of the brain model can be selected by first bisecting the volume with the cutting plane to expose the contents of the volume, and then selecting a point on the exposed surface. Note that in this usage the plane not only exposes the interior of the data, but it also expresses constraint of the point indicated by the trajectory prop to a plane, without requiring an explicit mode to do so.

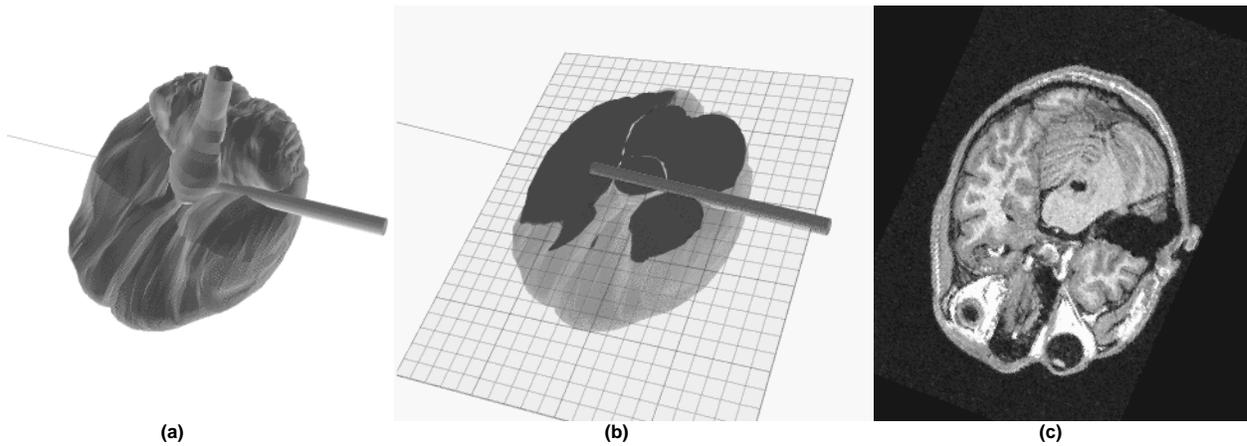


Figure 4: (a) Trajectory embedded in the brain and (b) exposed using the cutting plane. (c) The corresponding slice from the MRI data.

4.4. Two-handed Interaction

Note that our interface props are explicitly intended to be used with both hands. The coarse task of positioning the head is accomplished with the non-dominant hand, while the more precise tasks of specifying vectors and cutting planes are performed with the dominant hand. Also note that the interface allows the user to position their hands relative to one another, rather than forcing the user to move a single hand relative to an absolute 3D space. We have found that using such relative or *referential gestures* results in an interface which is much more natural than one which requires *absolute gestures*. See [10] for more information regarding two-handed interaction.

- Although many non-neurosurgeons have suggested using a more realistic head prop, not even one neurosurgeon has done so after operating the interface. In fact, when we suggest the idea, neurosurgeons flatly resist it.
- The ergonomics of a sphere are superior because it can comfortably be held at any orientation.

We currently believe that the best solution will combine the traits of the rubber ball and the doll’s head. For example, we will add tactile cues to the ball which indicate the orientation of the patient’s nose and ears.

4.2. Slicing the Patient Data with a Cutting-Plane Prop

We also provide a cutting-plane prop which can be used to specify arbitrarily oriented slices through the patient’s anatomy, allowing the surgeon to examine cross-sections of the volume or to dissect obstructing portions of the anatomy. For this prop, we use all six degrees of freedom that the tracker provides, thereby allowing the plane to be freely manipulated in space.

Note that the cutting-plane is used in conjunction with the head prop rather than as a separate tool. The user positions the head prop with their non-dominant hand while holding the cutting-plane up to it with their dominant hand. As a result, even though all six degrees of freedom are enabled when moving the plane, it does not seem difficult to control: your dominant hand can perform precise positioning tasks better than your non-dominant hand [16], and due to the inherent symmetry of the plane, all six degrees of freedom do not have to be controlled to perfection.

There are three distinct clinical uses for the cutting-plane prop as we have implemented it:

- **Volume Exploration:** The user can interactively sweep the cutting plane through the volume. In the present implementation, the grey-scale data from the MRI volume is only shown when the user stops moving the cutting plane. We plan to improve the efficiency of our algorithms to allow real-time interactive display of the data.
- **Volume Dissection:** Once the plane is selected, a portion of the volume can be permanently cut away.
- **Mensuration (Measuring Distances):** A grid pattern on the computer rendering of the plane can be used as a ruler.

We had not realized that the cutting-plane prop could be used to measure distances, but some surgeons started employing it for this purpose. When the user manipulates real objects in real space, new or unusual ideas can readily be expressed; the user is not artificially bound by an abstraction or a metaphor.

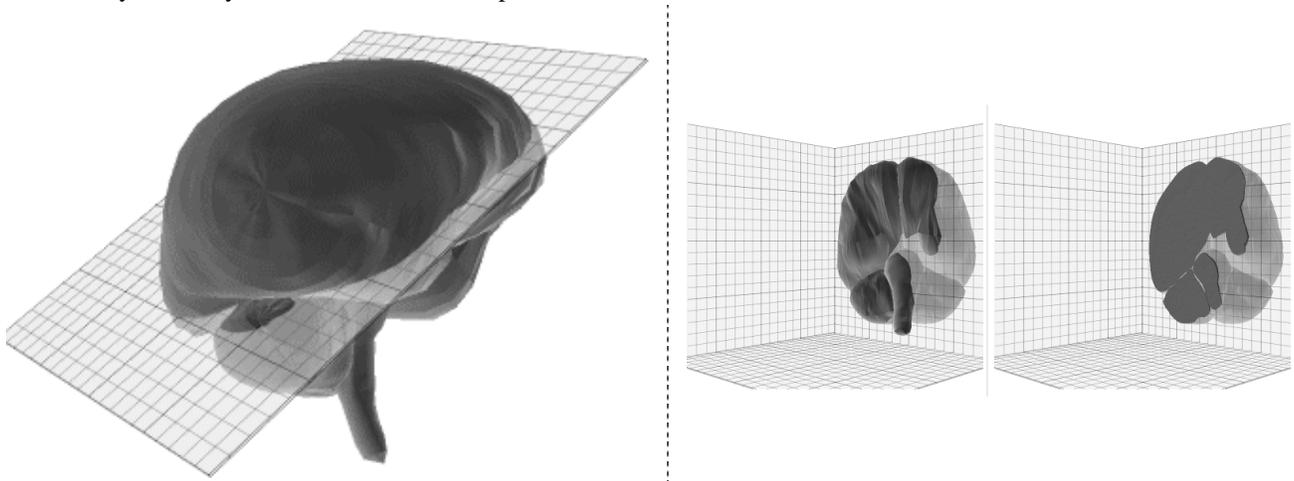


Figure 3: (Left) The Cutting-Plane Embedded in a Polygonal Brain. (Right) Selecting a Slice with the “3D Stage” and the Resulting Cross-Section.

We have also experimented with some alternative interfaces for the selection of cutting planes. Before we thought of using the cutting-plane prop, we implemented a 3D extension of Osborn’s pool-of-water interface [25] in which the polygonal brain model could be dipped in an imaginary pool of water, the surface of which defines the cutting plane. The stationary surface of the “pool” was parallel to the surface of the screen, so the brain model could be sliced by rotating the prop and moving it forwards or backwards to control its depth in the pool.

Although the engineers who developed the interface thought it seemed reasonable, the neurosurgeons reacted very negatively to it. We thought the major problem with the interface was that the cutting surface was parallel to the screen surface, making it difficult to perceive where the cut was relative to the rest of the brain model, so we implemented a variant known as the “3D

4. DESCRIPTION OF THE USER INTERFACE

We monitor the physician's gestures using the Polhemus FASTRAK magnetic tracking system, which can simultaneously report the position and orientation of up to four small, lightweight receivers (each approximately 1.5 cm^3) relative to a fixed magnetic transmitter. The transmitter constantly puts out a weak magnetic field. The receiver generates an electric signal as it is moved through the field. The Polhemus' processing electronics then amplify and analyze this signal to determine the real-world position and orientation of the receiver relative to the transmitter. We have embedded these small receivers into a set of tools which the surgeon can manipulate in combination with one another. The position and orientation of these props may then be coupled to the display of a graphical image on the computer screen.

Readers who might be familiar with older versions of the Polhemus technology should be aware that the Polhemus FASTRAK is a vastly superior device. The problems with older versions of the technology made it nearly unusable outside of the research laboratory, but we strongly believe that the FASTRAK is usable as it stands by real users doing real work.

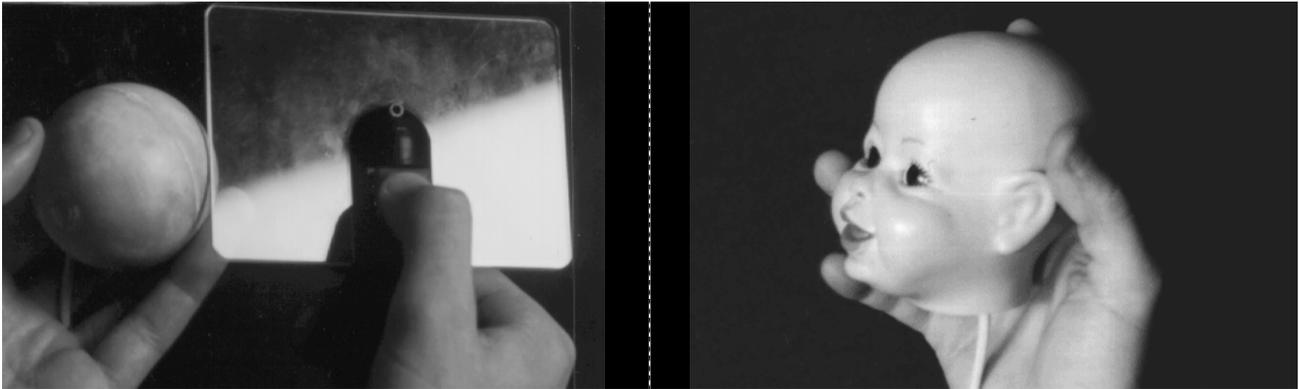


Figure 2: (Left) A User Holding the Head and Cutting-Plane Props; (Right) The Doll's Head Version of the Head Prop.

4.1. Viewing the Patient Data with a Head Prop

To manipulate the computer model of the patient's head, we provide the physician with a solid rubber ball that can be held comfortably in one hand. The ball can be turned to rotate a polygonal model of the patient's brain, and can simultaneously be moved towards or away from the user's body to control the zoom factor. We have also experimented with using a doll's head for the head prop, which provides additional tactile cues. Note that the ball is not devoid of tactile cues: the cable exiting the ball is roughly aligned with the spinal cord of the polygonal brain, and the ball's molding seam corresponds to the patient's interhemispheric fissure.

The tracking technology provides six degrees of freedom (the x , y , z position plus three rotation angles), but when manipulating the head prop all of this information is typically not needed. Instead, since it is rarely useful to move the polygonal brain left-right or up-down, by default we constrain the polygonal brain to be centered on the screen. This reduces positioning the head to a four-degree of freedom task, which simplifies the manipulation in a way that users find natural. In cases where it is needed, users are still able to use all six degrees of freedom by switching modes.

Once a desired rotation and zoom is achieved, the surgeon often wants to freeze the image in place on the screen. Thus some "clutching mechanism" to tell the computer to stop tracking the interface props is necessary. We currently use a foot pedal to clutch the head prop: whenever the foot pedal is held down, motion is enabled. Having a button directly on the head proved impractical for ergonomic reasons. We have also experimented with voice control of the clutch [10].

We originally had planned to provide the surgeon with a realistic skull-shaped prop, but we have retreated from this approach for the following reasons:

- Using a realistic head prop leads directly to false user expectations. For example, the neurosurgeon will sometimes hold the cutting plane up to the doll's eyes, expecting to see a cut directly through the orbits. But since neuroanatomy varies greatly between individuals, the morphology of the real-world prop and the polygonal brain model do not precisely correspond. As a result, the cut does not go exactly through the orbits. When using a more abstract form such as the ball, the neurosurgeon does not expect the prop to precisely match the brain model, so this usability problem does not arise.

In the human-computer interaction literature, previous publications have skirted around the general idea of employing user interface props, but no research we are aware of has treated this as an important theme in itself.

McKenna's discussion of interactive viewpoint control via head tracking [21] suggests that head tracking could be augmented with "tracked objects in real space which have matching computer representations." Both McKenna and Fitzmaurice [6] describe interfaces which track a miniature monitor in real space, allowing users to view an imaginary 3D landscape that surrounds them.

Badler's discussion of multi-dimensional input [1] asserts that using objects themselves as feedback is important because it allows "the computer to interact with the real environment controlled by the operator." This observation is essential: if the computer interacts with the user's real environment, the computer is forced to work on the user's own terms. The computer, rather than the human, transduces the input stream into an appropriate format.

The 3-Draw computer-aided design tool [29] employs two interface props. In 3-Draw, the user holds a stylus in one hand and a tablet in the other. The props are used to draw and view an object which is seen on a desktop monitor. 3-Draw is a good example of a props-based interface, but unlike the present work, 3-Draw is not immediately applicable to medical imaging applications.

Also note that a separate paper describing our work [10] discusses our interface design rationale and some of the implementation issues in greater detail. This paper also more thoroughly references related work in human-computer interaction.

3. DESIGNING AN INTERFACE FOR THE NEUROSURGEON

Designing high-quality user interfaces differs from traditional engineering problem-solving in that the engineer cannot, alone, provide a good solution. It is extremely important to involve the real users of an interface in the design process from the very beginning. In the user interface community this is referred to as user-centered design. Many excellent references exist which discuss this topic and related issues. Nielsen's and Norman's texts are a good place to start [23][24].

Like any user group, neurosurgeons have particular needs and demands which their user interfaces must meet. Neurosurgeons are driven by a single goal: deliver improved quality of patient care at a lower cost. They are extremely busy, frank, demanding, and generally not interested in computers. They do not hesitate to criticize, they often suggest good new ideas, and they provide concrete goals. We also should note that there is considerable economic incentive to solve the neurosurgeon's problems. Neurosurgery is time-consuming and expensive, so reducing the time needed to plan and execute surgical procedures by even a small percentage can quickly pay for a lot of technology. None of this technology, however, will do the surgeon any good if it is unusable, so in a sense the software usability problem is the most severe bottleneck preventing the ubiquitous use of advanced visualization software.

A neurosurgeon does not have the time or the patience to learn and re-learn the details of a needlessly complex user interface. The interface should be obvious and self-explanatory. The neurosurgeon also must cope with frequent distractions, so the interface should not employ equipment which is difficult to remove or put down, and it must not have explicit software modes which are easily forgotten.

In our experience many neurosurgeons will a priori refuse to use an interface which is mouse or keyboard driven. Not surprisingly, although there are several commercial 3D neurosurgical planning packages on the market, none have come into common usage in our facility because their user interfaces are not intuitive to the neurosurgeon. We believe that a well-designed interface which allows direct 3D interaction will allow neurosurgeons to get over this "usability hump" and make use of advanced 3D visualization software as part of their daily routine.

With the present hard-to-use technology, 3D visualization and planning software will only be used for complex clinical cases. However, we believe that even in simple clinical cases 3D manipulation could yield useful information to the surgeon. For example, surgeons will plan many procedures directly from 2D slices taken from Magnetic Resonance Imaging (MRI) and CT (Computed Tomography) volumes. However, the standard plane of acquisition is different for these modalities, which sometimes leads to an incorrect perception of the 3D location of a tumor or other target. If 3D volume exploration tools were in routine use, such classes of errors could largely be eliminated. We argue that 3D visualization tools will not come into routine use unless they are easy for the neurosurgeon to use.

Physicians interact with each other by voice and gesture, and that is how we believe they should interact with their computer systems. The guiding principle of our user interface has therefore been to “listen in” on the surgeon’s natural speech and gestures. The surgeon’s speech can be recognized using voice input technology, and his or her gestures can be monitored by having the surgeon manipulate specially constructed tools or “props” whose position and orientation are tracked by the computer. Our interface tracks the props using the Polhemus FASTRAK six-degree-of-freedom digitizer [27]. We find it helpful *not* to think of the props as explicit input devices. Rather, we prefer to consider them as familiar tools which help the surgeon to reason about the spatial task he is performing.

An interface which requires the neurosurgeon to wear an instrumented glove and make grabbing gestures to manipulate imaginary objects would not offer this style of interaction. No matter how realistic the on-screen graphics are, the user does not experience the visceral kinesthetic and tactile feedback which comes from grasping a real-world object. Although some progress has been made with force-feedback manipulators [2][15], the technologies are still awkward and tend to limit the range of possible motions. Providing a real-world object that approximates the virtual object is a low-tech solution to the problem which provides many of the advantages of force-feedback without the associated technological headaches.

Although parallels can be drawn between our three-dimensional interface and “virtual reality” surgical simulation interfaces which have recently become popular, we stress that we do *not* consider our interface to be a form of virtual reality. In particular, our system does not employ a head-mounted display, as the present display technology has poor resolution and is cumbersome to wear. Instead we use a standard computer monitor, the projected 3D graphics on which mirror the real-world configuration of the interface props. We find it more useful to refer to our system as a “spatial desktop interface.”

The prototype of our props-based interface for neurosurgical planning has proven successful and has elicited enthusiastic comments from users. In particular, we find that neurosurgeons who have never before seen the system can easily “get the hang of it” within about one minute of touching the props.

2. PREVIOUS WORK

There is a vast literature describing systems for the planning and visualization of surgical procedures, yet a surprisingly tiny fraction of this literature has specifically addressed the usability issues that arise in medical imaging. In particular, there are very few high-quality published accounts in the medical imaging literature which discuss three-dimensional interaction.

Although the medical imaging field has long pushed the “state of the art” in computer graphics, we believe that as a whole the field greatly lags behind the state of the art in human-computer interaction. We hope that the present research will help our colleagues in both medical imaging and human-computer interaction to recognize the importance of research dedicated to medical imaging software usability, as the field as a whole offers particularly difficult and rewarding challenges to user interface design.

Suetens et al. [31] describe a system which provides stereoscopy, head motion parallax, and real-time rotation capabilities. The computer renders the image on a standard CRT, but the user’s head is tracked, so that as the user moves his head relative to the CRT the viewpoint of the 3D object changes slightly. Rotation of the object being viewed is achieved by turning a stylus about its endpoint. More recent papers by Ware [35] and McKenna [21] give implementation notes and provide more detailed discussions of the usability concerns of similar motion parallax systems. Our system does not presently provide motion parallax, although we plan to experiment with this capability in the future. However, we believe that turning a miniature head model in one’s hand is a more natural and ergonomically sound way to specify 3D rotation than Suetens’ technique of turning a stylus about its endpoint.

Kaufman[17][18] describes an interface which allows 3D interaction with a volumetric environment. This system primarily uses the 3D input device for the selection of single 3D points. Kaufman’s system aids 3D point selection by using gravity functions, that is, software constraints of the user’s motion which aid the selection pre-defined points that are known to be of interest.

Other researchers, such as Hohne [12] and Robb [28], have recognized the importance of good user interfaces. Their systems provide reasonable interfaces based on traditional input devices. However, we believe it is possible to provide much more natural interaction with interfaces predicated upon direct spatial interaction.

A Three-Dimensional User Interface for Neurosurgical Visualization

Ken Hinckley, Randy Pausch, John C. Goble, Neal F. Kassell
University of Virginia, Charlottesville, VA 22903
kph2q@virginia.edu, pausch@virginia.edu, jcg7q@virginia.edu

ABSTRACT

We describe a three-dimensional user interface for pre-operative neurosurgical planning based on the physical manipulation of familiar real-world objects in free space. Using these *passive interface props*, neurosurgeons can apply their existing skills to specify spatial relationships in a natural and direct manner. The interface currently employs a head viewing prop, a cutting-plane selection prop, and a trajectory selection prop. Each prop is a simple real-world tool, the position and orientation of which is tracked by the computer. The behaviors associated with each prop serve as “interaction primitives” which can be composited to describe complex spatial relationships, resulting in a powerful, expressive, and conceptually simple user interface.

From the surgeon’s perspective, the interface is analogous to holding a miniature skull which can be “sliced” and “pointed to” using the cutting-plane and trajectory props. Our informal evaluation sessions have shown that with a cursory introduction, neurosurgeons who have never seen our interface can understand and use it without training.

1. INTRODUCTION

Modern imaging technologies such as MRI generate tremendous quantities of three-dimensional information, but in order for this information to be clinically useful, physicians must have intuitive methods to view, explore, and dissect the volumetric data. Unfortunately, traditional user interface techniques which utilize input devices such as mice and keyboards are not well-suited to these inherently 3D tasks. The resulting hard-to-use software has made physicians hesitant to embrace the personal use of computers as part of their daily routine.

Three-dimensional interaction with a computer does not necessarily have to be difficult. In the real world, many of our daily tasks require three-dimensional manipulation of real objects. We typically perform these tasks with little cognitive effort, with both hands, and with total confidence in our movements. We believe that a user interface for three dimensional manipulation and visualization of medical images can offer equally facile interaction.

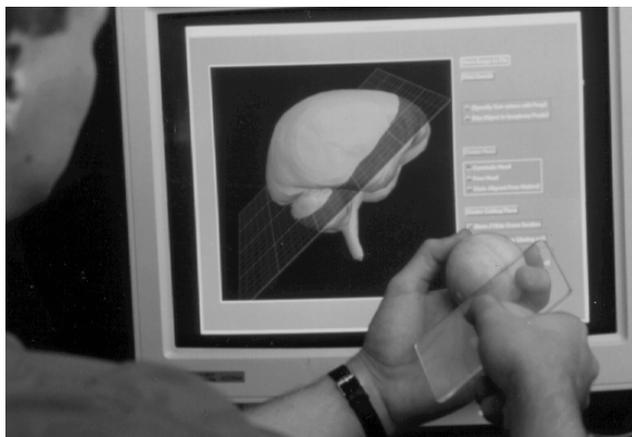


Figure 1: A User Specifying a Cutting-Plane with the Props.

We have specifically focused on the pre-operative planning of neurosurgical procedures, but we believe our techniques are extensible to other medical specialities. Neurosurgery occurs in three dimensions and deals with complex three-dimensional structures; the neurosurgeon works and thinks in terms of real objects in real space. By providing direct spatial interaction, our user interface allows the neurosurgeon to work and think in these same terms.