Passive Real-World Interface Props for Neurosurgical Visualization

Ken Hinckley[1,2], Randy Pausch[2], John C. Goble[1] and Neal F. Kassell[1]

University of Virginia
Departments of Neurosurgery[1] and Computer Science[2]
kph2q@virginia.edu (804/243-0310), pausch@virginia.edu,
jcg7q@virginia.edu, neal@msmail.neuro.virginia.edu

ABSTRACT

We claim that physical manipulation of familiar real-world objects in the user's real environment is an important technique for the design of three-dimensional user interfaces. These real-world *passive interface props* are manipulated by the user to specify spatial relationships between interface objects. By unobtrusively embedding free-space position and orientation trackers within the props, we enable the computer to passively observe a natural user dialog in the real world, rather than forcing the user to engage in a contrived dialog in the computer-generated world.

We present neurosurgical planning as a driving application and demonstrate the utility of a *head viewing prop,* a *cutting-plane selection prop*, and a *trajectory selection prop* in this domain. Using passive props in this interface exploits the surgeon's existing skills, provides direct action-task correspondence, eliminates explicit modes for separate tools, facilitates natural two-handed interaction, and provides tactile and kinesthetic feedback for the user. Our informal evaluation sessions have shown that with a cursory introduction, neurosurgeons who have never seen the interface can understand and use it without training.

KEYWORDS

three-dimensional interaction, gesture input, two-handed interaction, haptic input, neurosurgery, visualization

INTRODUCTION

In our everyday lives, we are constantly confronted with tasks that involve physical manipulation of real objects. We typically perform these tasks with little cognitive effort, with both hands [5], and with total confidence in our movements. We believe that for many applications a three-dimensional user interface can offer equally facile interaction.

Our application domain is the pre-operative planning of neurosurgical procedures. Neurosurgery occurs in three dimensions and deals with inherently three-dimensional structures. The neurosurgeon works and thinks in terms of real objects in real space; a three-dimensional user interface should allow the neurosurgeon to work and think in these same terms. As one surgeon put it, "*I want a skull I can hold in my hand.*"

We propose a 3D interface which permits the user to manipulate familiar objects in free space. These *passive interface props* act as tools which help users reason about their tasks. By unobtrusively embedding Polhemus FASTRAK six degree-of-freedom trackers [24] within the props, we enable the computer to observe the user's gestures. This results in a human-computer dialog where the system *watches* the user [23], in contrast to the traditional approach where the user generates input tokens in a contrived dialog.

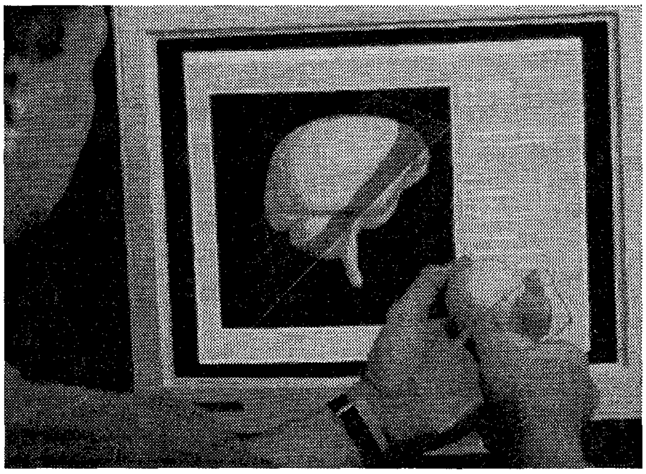

Figure 1: **A User Selecting a Cutting-Plane with the Props.**

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.

Compared to "3D widgets" [6][10], a props-based interface offers several advantages. There is no need to make a widget's behavior explicit or to make the user realize the widget is an active interface component. The appearance of the props indicates their use and their *palpability* makes users immediately and continuously aware they exist. Drawing a widget without cluttering the scene becomes trivial, since there *is no* widget. Also, for casual users such as surgeons, manipulating a real tool is familiar and natural, whereas an abstract widget, no matter how well designed, is not.

CHI94-4/94 Boston, Massachusetts USA

In the domain of neurosurgical planning, our props-based interface has proven successful and has elicited enthusiastic comments from users. With a cursory introduction, neurosurgeons who have never before seen the interface can understand and use it without training.

PREVIOUS WORK

Previous researchers have skirted around the general idea of passive interface props, but no research we are aware of treats this topic as an important theme in itself.

Tognazzini's discussion of stage magic applied to human-computer interaction [29] gives an excellent introduction to the "passive props" approach:

> The magician's tools should be disguised to look like objects in the real world. "If these things are common things, objects with which the spectator is familiar, this spectator will accept them in terms *as he knows them*. He will assume the device to be the same as the common article with which he is acquainted."--Fitzkee [7].

In 2D user interfaces, techniques such as the "desktop metaphor" are the magician's tools, but in 3D, these techniques are less satisfactory. In 3D the "magician's tools" literally are physical props from the real world.

Schmandt's pioneering work [26] suggests that real objects can themselves be used for feedback. Schmandt describes a workspace which allows users to reach into the space that the image appears to occupy. A hand-held wand is seen through a half-silvered mirror, upon which the computer graphics are projected. A white spot on the wand itself acts as a real-world cursor. The major drawback is a lack of correct occlusion depth cues.

In flight simulators, the user's entire environment is a mock-up of a real object, the aircraft cockpit. Simulators mimic an existing technological artifact. Our goal is not to mimic existing artifacts, but rather to provide three-dimensional real-world tools which allow a neurosurgeon to control advanced visualization software.

Several virtual reality programs make use of props. For example, a real golf club can be used to play virtual reality golf [30], or a real flashlight can be used to move virtual cameras and spotlights [22]. These systems each employ only a single prop, rather than multiple props which the user manipulates using two hands.

McKenna's discussion of interactive viewpoint control via head tracking [19] suggests that head tracking could be augmented with "tracked objects in real space which have matching computer representations." Both McKenna and Fitzmaurice [8] describe interfaces which track a miniature monitor in real space, allowing users to view an imaginary 3D landscape that surrounds them.

Badler [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 [25] 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. Unfortunately, the paper describing 3-Draw does not primarily focus on the issues raised by the 3D interaction techniques. The present paper contributes (1) an example of how props can be used in another problem domain, (2) an exposition of the underlying design philosophy, and (3) an explicit discussion of some issues raised by props-based interaction techniques.

NEUROSURGEONS AND THEIR NEEDS

Neurosurgeons are driven by a single goal: deliver improved patient care at a lower cost. They are 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.

A user interface for neurosurgical planning must have rapid learning and re-learning times. There may be several days between clinical cases which require sophisticated planning tools. Also, the surgeon must cope with frequent distractions, and therefore must be able to quickly detach from the user interface, both physically and cognitively. Thus, the interface must not employ devices that will be difficult to put down, and it must not have explicit modes that are easily forgotten.

We should stress that the neurosurgeon's existing visualization and planning tools are almost exclusively two dimensional. This is an artifact of historical technological limitations rather than preference; the three-dimensional tools we describe will allow neurosurgeons to view and explore the individual patient's anatomy in ways that previously have not been possible.

PROPS FOR NEUROSURGICAL VISUALIZATION

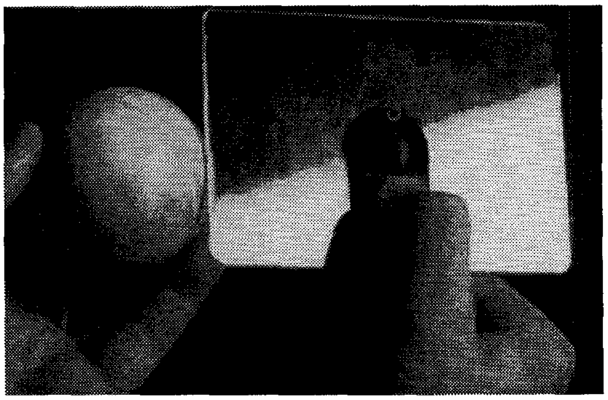

Figure 2: **A Close-up of the Head and Cutting Plane Props**

Viewing Patient Data with a Head Prop

We provide the surgeon with a *head prop* for manipulating the individual patient's head data. The prop is a small rubber sphere which can be held comfortably in one hand. We have also tried using a small doll's head, which provides additional tactile cues. Note that the ball is not devoid of tactile cues: the ball's seam corresponds to the patient's inter-hemispheric fissure, and the ball's electrical cable is aligned with the patient's spinal cord.

Rotating the prop causes a polygonal model of the patient's brain to rotate correspondingly on the screen. The user can control the image zoom factor by moving the prop towards or away from his or her body. Since moving the object left-right or up-down is typically not useful, we have found it helpful to constrain the (x, y) position of the polygonal brain to the center of the screen. This simplifies the task and users find it natural.

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:

- 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. This includes the surgeon who originally said he wanted a *skull* he could hold in his hand.
- The ergonomics of a sphere are superior because it can comfortably be held at any orientation.
- Using a realistic head prop leads directly to false user expectations. For example, users 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 virtual head 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 user does not expect the prop to precisely match the brain model, so this usability problem does not arise.

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.

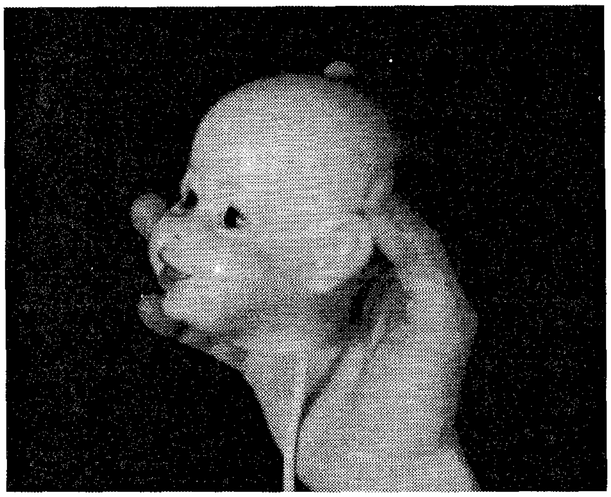

Figure 3: **Doll's Head Version of the Head Prop.**

Slicing the Patient Data with a Cutting-Plane Prop

We also provide a *cutting-plane prop* which specifies the position and orientation of an arbitrary slice through the patient's anatomy. The prop itself is a rectangular plate with a housing for the tracker (*fig. 2*). Users can spread their fingers across the plate to get a direct haptic sense of how it is oriented in space. The appearance of the cut-plane prop differentiates it from the head prop and makes its purpose immediately obvious.

Note that the cut-plane prop is used in concert with the head prop rather than as a separate tool. The user holds the cut-plane against the head to indicate a slice through the brain data. The reader can easily approximate this interface. Seat yourself in a chair with armrests. Grasp a ball in one hand and a small book in the other. While supporting your elbows with the armrests, hold the book up to the ball, and orient each as deemed necessary. This is all that our interface requires for 3D manipulation.

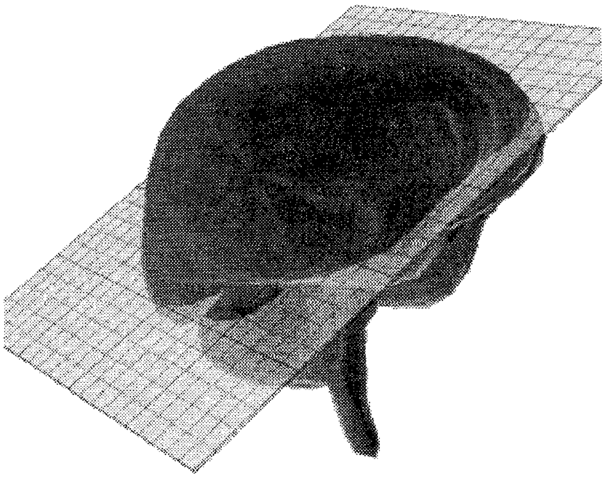

Figure 4: **The Cutting-Plane Embedded in a Polygonal Brain.**

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.
- Volume Dissection: Once the plane is selected, a portion of the volume can be permanently cut away.
- Measuring Distances: A grid pattern on the computer rendering of the plane can be used as a ruler.

We had not realized that the cut-plane prop could be used as a ruler, but much to our surprise some users 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.

Indicating Surgical Paths with a Trajectory Prop

A *trajectory selection prop* allows the surgeon to specify 3D vectors and points. The current prototype is a stylus-shaped tool equipped with a tip switch. Moving the trajectory prop relative to the head prop specifies the position and orientation of a cylindrical virtual probe (*fig. 5*) 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 cases, but often a *nonlinear* trajectory is required to avoid vasculature or healthy brain tissue. The present prototype does not yet support nonlinear trajectories, although a solution using curves sketched in 3D (as done in 3-Draw [25]) can be envisioned.

A linear trajectory consists of a target point inside the brain and a vector to that point. The trajectory 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 user holds the trajectory prop's tip switch 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.

Points which lie on 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 case 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.

Two-handed Interaction

Traditionally, two-handed input has been viewed as a technique which allows the user to perform two subtasks in parallel [5]. For 3D input, however, we believe two-handed interaction is of even greater importance. Previous work [9] has shown that people often express spatial manipulations using two-handed gestures. Furthermore, we have found that using both hands for 3D interaction has a number of other advantages:

- Users can effortlessly move their hands relative to one another or relative to a real object, but it requires a conscious effort to move a single hand relative to an abstract 3D space. (We note that the designers of 3-Draw [25] have made similar observations.)
- Use of two hands provides physical support. One-handed three-dimensional input can be fatiguing, but if the hands can rest against one another or against a real object, fatigue can be greatly reduced.
- The user can express complex spatial relations as a single cognitive chunk. For example, users can manipulate our interface props with two hands to specify a cut relative to a particular brain orientation in a single gesture. Not only does this make the interaction parallel (as opposed to being sequentially moded), but it also results in an interface which more directly matches the user's task.

In our specific interface, the use of two hands helps the user to avoid constantly picking up and putting down individual props. Also, the user's non-dominant hand is well suited to the head prop since its constrained movement requires less precision to control [15].

IMPLEMENTATION ISSUES

Real-World / Graphics Correspondence

Users expect the real-world relationship between the props to be mirrored by their on-screen graphical representations. Simplifying control of the head prop by centering it on the screen, however, requires a software mapping of its real-world position to its centered position (note that no such mapping is required for the *orientation* of the prop). This implies that the on-screen position of the plane is equal to the *mapped* position of the head prop plus the real-world (x, y, z) delta between the head prop and the cut-plane prop. In other words, the plane is drawn relative to the mapped position of the head prop.

This can result in the following artifact: if the user holds the cut-plane prop still and translates only the head prop, the polygonal brain will remain centered and the virtual plane will move in the *opposite* direction. Interestingly, users rarely expose this artifact, because they typically

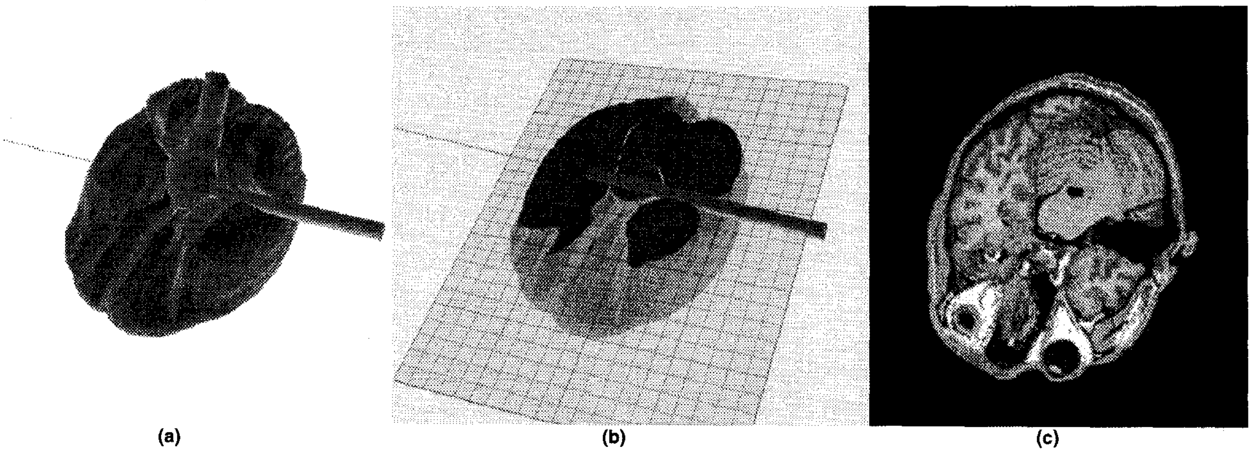
(a) (b) (c)

Figure 5: (a) Trajectory embedded in the brain and (b) exposed using the cutting plane. (c) The corresponding slice from the MRI data.

hold the head prop still and move the cut plane relative to it; as a result the interaction is quite natural.

Clutching Mechanisms

A mechanism is needed to tell the computer to "stop watching" a particular prop. This allows the surgeon to freeze the image in a desired configuration and put down the props. We use a foot pedal to "clutch" the head prop and a thumb button (mounted to the acrylic plate) to clutch the cut-plane prop. The foot pedal behaves like a gas pedal: hold it down when you want to move. Similarly, the cut-plane prop only allows motion while the thumb button is being held down. Sellen [27] has shown such tension can reduce mode errors. In our interface fatigue is not a problem, but we note that avoiding fatigue is a more important design issue than avoiding mode errors.

We have experimented with voice control of the clutch. Saying "move *<prop>*" enables motion, while saying "stop *<prop>*" disables motion. Since the user is engaged in a real-time manipulation task, the time to speak and recognize a voice command causes an irritating delay. It is not clear if this problem would persist with a more sophisticated voice recognizer than our low-cost unit [31]; the delay introduced by speaking the command might itself prove intolerable. Under some conditions voice input can also interfere with short term memory [16], which poses another possible difficulty.

The Disappearing Object Problem

It is possible to position the cutting plane such that it slices away the entire object (*fig. 5, top*), which sometimes leads to confusion. We currently draw a wireframe wherever the object has been cut away (*fig. 5, bottom*), but this solution is not ideal because the wireframe obscures the cross-section and its depth is ambiguous. We have avoided transparency for technical reasons, but transparent surfaces could offer alternative solutions.

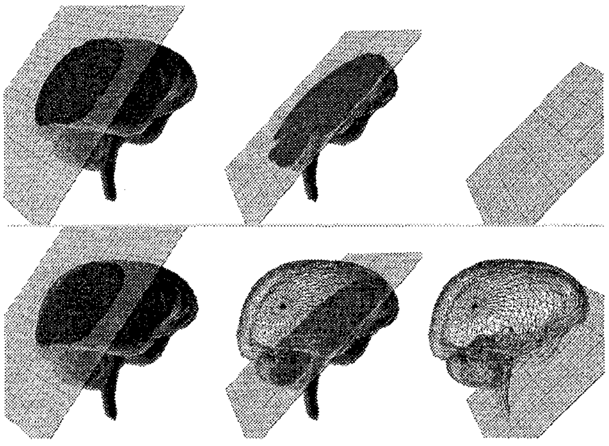

Figure 5: **The Disappearing Object Problem and Wireframe Solution**

System Components

We track the interface props using the FASTRAK system manufactured by Polhemus [24]. Its small and lightweight tracking devices can easily be embedded within our interface props. Other system components include:

- Hewlett-Packard 735 graphics workstation [11].
- "Sense & Switch" digital I/O module [14].
- Verbex Speech Commander voice recognizer [31].

INFORMAL INTERFACE EVALUATION

We have implemented an alternative cutting-planes interface based on Osborn's "pool-of-water" technique [21]. A static stage consisting of three orthogonal cutting surfaces (*fig. 6*) is drawn on the screen. The user indicates the cut by positioning the polygonal brain such that it intersects one or more of the three surfaces.

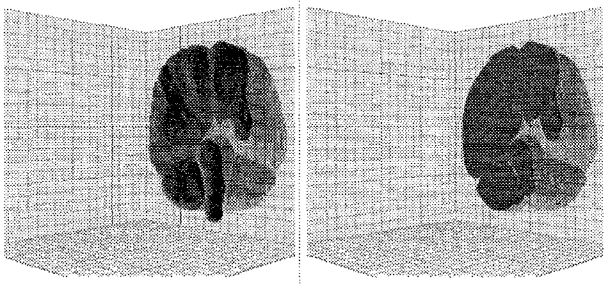

Figure 6: **Specifying a Cut with the "3D Stage" and the Resulting Cross-Section**

We have informally compared this "3D stage" interface to our props-based approach. We find that neurosurgeons uniformly and decisively prefer specifying the cutting plane using the props. We believe neurosurgeons prefer the props-based method because it more closely models the task they have in mind. For example, one user commented that he had to "first imagine the 3D slice" when using the stage whereas with the prop he could "move the plane to suit the cut." The surgeon wants to select the cut relative to a specific view of the brain. The cut-plane prop can express this concept, whereas the 3D stage cannot.

Most user complaints resulted from the ergonomic design of the cutting plane prop. The original prop had a handle and an ill-placed, hard-to-press thumb button. We have constructed a new prop which eliminates the handle and replaces the thumb button with a membrane switch. The membrane switch responds to a light touch, provides a low vertical profile, and has a large activation area. Several companies [13][28] offer these switches.

Despite the initial ergonomic difficulties, users were very enthusiastic about the props-based interface. All of the approximately 15 neurosurgeons who have tried the interface were able to "get the hang of it" within about one minute of touching the props; many users required considerably less time than this.

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 system based on passive observation can come close to achieving this goal.

Interface props provide an excellent mechanism for observing users. If a surgeon wants to describe a cutting plane to another person, the surgeon uses real objects to gesturally indicate the plane's location. Giving the surgeon specially instrumented props to express this idea allows our system to observe the dialog while also helping the surgeon to reason about the task.

Although the present interface props are specialized for neurosurgery, this does not preclude the addition of generic tools. Specialized props can be used for common tasks, whereas a generic prop whose graphical representation changes depending on the current mode could be employed for less common tasks. However, we encourage designers to use specialized props whenever appropriate. Although using application-specific props may limit generality, the haptic feedback from a physical tool provokes a visceral and enthusiastic response from users which should not be underestimated.

We have described several benefits that naturally follow from using passive real-world interface 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 chunks 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.

Neurosurgical planning is an ideal application for props because neurosurgeons work and think in terms of real objects in real space. 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.

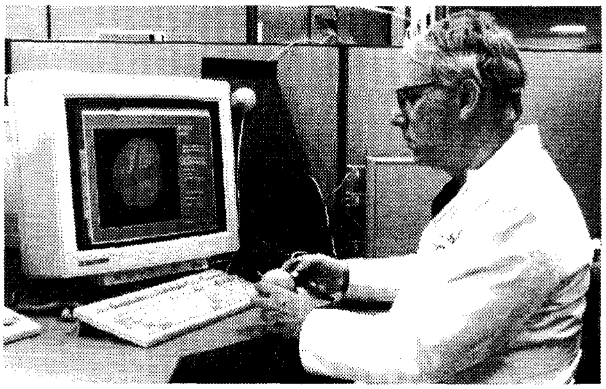

Figure 7: **A Neurosurgeon Specifying a Cutting-Plane.**

FUTURE WORK

One limitation of the props-based approach is that the user cannot easily express constrained motions. One solution is to provide additional explicit controls, such as an analog slider which moves the on-screen cutting plane along its normal. Alternatively, the system might "notice" when the user is making minute adjustments and switch into a "fine adjust" motion phase.

We currently assume the user's head is fixed in space, but other researchers [19][32] have demonstrated that head tracking can be helpful. We plan to experiment with this capability in the near future.

Our 3D display space appears to be *inside* the monitor, but the props are held in a space *outside* the monitor. One could build a work station (similar to Schmandt's [26]) which unifies the work and display spaces. The system would project the graphics onto an *opaque* mirror, behind which the user holds the props. The images appear to be behind the mirror, and the props *are* behind the mirror, so the work and display spaces correspond. Since only computer-generated graphics are seen, correct occlusion cues can be maintained. Ideally, this setup would allow the user's kinesthetic and visual perceptions of the objects to agree completely.

ACKNOWLEDGEMENTS

We wish to thank the Department of Neurosurgery for their support and guidance. We also wish to thank Marc Pilipuf and Bob Bryant for help with building the props, John Snell for the polygonal brain surfaces, William Campbell for the photography, and Matthew Conway for his comments on passive observation.

REFERENCES

1. Badler, N., Manoochehri, K., Baraff, D., "Multi-Dimensional Input Techniques and Articulated Figure Positioning by Multiple Constraints," Proc. 1986 ACM Workshop on Int. 3D Graph., 151-170.

2. Bier, E., Stone, M., Pier, K., Buxton, W., DeRose, T., "Toolglass and Magic Lenses: The See-Through Interface," SIGGRAPH '93, pp. 73-80.
3. Bolt, R., Herranz, E., "Two-Handed Gesture in Multi-Modal Natural Dialog," UIST '92, pp. 7-13.
4. Buxton, W., "Chunking and Phrasing and the Design of Human-Computer Dialogues," Information Processing '86, Proc. of the IFIP 10th World Computer Congress, pp. 475-480.
5. Buxton, W., Myers, B., "A Study in Two-Handed Input," Proc. CHI'86, pp. 321-326.
6. Conner, D., Snibbe, S., Herndon, K., Robbins, D., Zeleznik, R., van Dam, A., "Three-Dimensional Widgets," Proc. 1992 Symposium on Interactive 3D Graphics, pp. 183-188, 230-231.
7. Fitzkee, D., Magic By Misdirection, Lee Jacobs Productions, P. O. Box 362, Pomeroy, OH, 45769-0362, 1975.
8. Fitzmaurice, G., "Situated Information Spaces and Spatially Aware Palmtop Computers," Comm. of the ACM, 36 (7), 1993, pp. 39-49.
9. Hauptmann, A., "Speech and Gestures for Graphic Image Manipulation," Proc. CHI'89, pp. 241-245.
10. Herndon, K., Zeleznik, R., Robbins, D., Conner, B., Snibbe, S., van Dam, A., "Interactive Shadows," Proc. UIST'92, pp. 1-6.
11. Hewlett-Packard Company, 3000 Hanover St., Palo Alto CA 94304.
12. Hinckley, K., Pausch, R., Goble, J. C., Kassell, N. F., "A Three-Dimensional User Interface for Neurosurgical Visualization," Proc. SPIE Conf. on Medical Imaging, 1994 (to appear).
13. Interlink Electronics, (805) 484-8855.
14. International Technologies, Inc., (401) 781-5595.
15. Kabbash, P., MacKenzie, I. S., Buxton, W., "Human Performance Using Computer Input Devices in the Preferred and Non-Preferred Hands," Proc. INTERCHI'93, pp. 474-481.
16. Karl, L., Pettey, M., Shneiderman, B., "Speech-Activated versus Mouse-Activated Commands for Word Processing Applications: An Empirical Evaluation," Intl. J. Man-Machine Studies, 1993.
17. Kelly, P. J., Kall, B. A., eds., Computers in Stereotactic Neurosurgery, Blackwell Scientific Publications Inc., Cambridge, MA, 1992.
18. Liang, J., Green, M., "JDCAD: A Highly Interactive 3D Modeling System," 3rd International Conference on CAD and Computer Graphics, Beijing, China, Aug. 1993, pp. 217-222.
19. McKenna, M., "Interactive Viewpoint Control and Three-dimensional Operations," Proc. 1992 Symposium on Interactive 3D Graphics, pp. 53-56.
20. Nielsen, J., "Noncommand User Interfaces," Communications of the ACM, 36 (4), pp. 83-99.
21. Osborn, J., Agogino, A., "An Interface for Interactive Spatial Reasoning and Visualization," Proc. CHI'92, pp. 75-82.
22. Pausch, R., Shackelford, M. A., Proffitt, D., "A User Study Comparing Head-Mounted and Stationary Displays," Proc. IEEE Symposium on Research Frontiers in Virtual Reality, Oct. 1993.
23. Pausch, R., Vogtle, L., Conway, M., "One Dimensional Motion Tailoring for the Disabled: A User Study," Proc. CHI'92, pp. 405-411.
24. Polhemus Navigation Sciences, (802) 655-3159.
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. Schmandt, C. M., "Spatial Input/Display Correspondence in a Stereoscopic Computer Graphic Work Station," Proc. SIGGRAPH '83, Computer Graphics, 17 (3), 1983, pp. 253-262.
27. Sellen, A. J., Kurtenbach, G. P., Buxton, W. A. S., "The Role of Visual and Kinesthetic Feedback in the Prevention of Mode Errors," Proc. IFIP INTERACT'90, pp. 667-673.
28. Spectra Symbol, (801) 972-6995.
29. Tognazzini, B., "Principles, Techniques, and Ethics of Stage Magic and Their Application to Human Interface Design," Proc. INTERCHI'93, pp. 355-361.
30. University of North Carolina, Dept. of Computer Science, Chapel Hill, NC. Unpublished virtual reality demonstration.
31. Verbex Voice Systems, Inc., 1-800-ASK-VRBX.
32. Ware, C., Arthur, K., Booth, K. S., "Fish Tank Virtual Reality," Proc. INTERCHI'93, pp. 37-41.
33. Wellner, P., "Interacting with Paper on the DigitalDesk," Communications of the ACM, 36 (7), 1993, pp. 87-97.