Attention and Visual Feedback: The Bimanual Frame of Reference

Ken Hinckley,[1,2] Randy Pausch,[1] and Dennis Proffitt[3]

University of Virginia: Departments of Computer Science,[1] Neurosurgery,[2] and Psychology[3]
{kph2q, pausch, drp}@virginia.edu

Abstract

We investigate the synergy of the two hands for virtual object manipulation. We report results from an experiment which suggest that the two hands together provide sufficient perceptual cues to form a frame of reference which is independent of visual feedback. The same is not true for one hand moving in empty space. Our interpretation is that users may not have to constantly maintain visual attention when both hands can be involved in a manipulation.

Our results suggest that using two hands can provide more than just a time savings over one-handed manipulation. Two hands together provide the user with information which one hand alone cannot. Our results also suggest that using two hands can potentially impact performance at the cognitive level by changing how users think about a task. Since the user can potentially integrate subtasks controlled by each hand without an explicit cost to switch between subtasks, this encourages exploration of the task solution space. Finally, to illustrate why one might expect this to be true, we present a task analysis which helps to reason about the differences between one and two-handed interfaces.

CR Categories and Subject Descriptors: I.3.6 [Computer Graphics]: Methodology and Techniques - Interaction Techniques; H.5.2 [Information Systems]: Information Interfaces and Presentation - User Interfaces - Input devices and strategies.

Additional Keywords: Two-handed interaction, virtual manipulation, attention and feedback, frames of reference, virtual reality.

1 INTRODUCTION

The central hypothesis of this study is our belief that the combined action of the two hands can play a vital role for the interactive manipulation of virtual objects. There is much informal evidence to support this position. Most everyday manipulative tasks involve both hands: for example, striking a match; unscrewing a jar; sweeping with a broom; writing on a piece of paper; dealing cards; threading a needle; or painting on a canvas, where the preferred hand holds the paintbrush, the nonpreferred hand the palette.

A few virtual reality applications and demonstrations, both on the desktop [20][21][8] and in fully immersive situations [16][15][22], have recognized the design possibilities for the two hands, and for years Buxton [3][13] has argued that one can improve both the naturalness and degree of manipulation of interfaces by employing both hands. Yet, beyond a possibly improved efficiency of hand motion, there has been little formal evidence of precisely what advantages, if any, the two hands can bring to virtual manipulation.

To appear in:

ACM / SIGGRAPH 1997 Symposium on Interactive 3D Graphics

Moreover, just because a behavioral strategy is exhibited in the real world, this does not necessarily mean that it will be useful in a virtual environment. Virtual environments offer opportunities to violate the limitations of physical reality, and one only needs to mimic those qualities of physical reality which facilitate skill transfer or which form essential perceptual cues for the human participant to perform his or her tasks.

To establish the utility of the two hands in virtual environments, we need to formally demonstrate what we can do with two hands that we can't easily do with one, and address some questions of when, and why, a bimanual interface might offer some advantages. We do not claim to answer all of these questions, but the current study offers some data which suggests areas where involving both hands may have some advantages.

We present an experiment which suggests that the two hands together form a hand-relative-to-hand frame of reference. A frame of reference is a centered and oriented perceptual coordinate system which is specified by a center point plus three directional axes. An interesting property of the bimanual frame of reference is that the information can be encoded by the hands themselves, and as such does not necessarily rely on visual feedback. As an intuitive example, it is easy to touch your index fingers behind your head, but this action is clearly not guided by vision.

2 HYPOTHESES

Our experiment investigates the following specific hypotheses:

H1. The two hands together provide sufficient perceptual cues to form a frame of reference which is independent of visual feedback.

H2. When using just one hand, subjects can employ other body-relative cues (sense of joint angles, sense of torso midline) to make an unbiased estimate of a remembered hand position, but these cues are less precise. Thus, unimanual control is more dependent on visual feedback.

H3. The physical articulation of a task can influence cognitive aspects of performance, in terms of the task strategy used. Using two hands together encourages exploration of the task solution space, and this will allow subjects to get a better sense of what a good strategy is for the experimental task.

2.1 Cognitive Aspects of Performance

Hypothesis 3, which asserts that using two hands can influence cognitive aspects of performance, has previously been articulated by Buxton. Working with his Input Research Group, Leganchuk [14] has provided some preliminary evidence which suggests that "representation of the task in the bimanual case reduces cognitive load."

Leganchuk's experiment studied an "area sweeping" task in which subjects selected an area encompassing a target. This is similar to sweeping out a rectangle to select a set of targets in a graphics editing application. Using both hands allowed subjects to complete the task significantly faster than using just one hand. Furthermore, the difference in times could not be attributed to the increased time-motion efficiency alone. This was interpreted as evidence that the bimanual technique "reduces cognitive load."

Another way to investigate the hypothesis that bimanual control can influence cognitive aspects of performance is to take direct measures of cognition, such as quantifiable metrics of learning, memory, or transfer of skill. Leganchuk's strategy of taking differences between one and two-handed techniques relies on the assumption that differences beyond those clearly accounted for by increased time-motion efficiency can be attributed to differences in cognitive load. But if one can demonstrate a direct metric of cognition, this assumption does not have to be introduced.

3 THE EXPERIMENT

3.1 Subjects

Seventeen unpaid subjects (13 female, 4 male) were recruited from our psychology department's subject pool. One subject (male) was left-handed. No subjects had experience with 3D input devices or two-handed computer interfaces.

3.2 Task

The task consisted of two phases. In the primary phase, users attempted to align two virtual objects. The purpose of this phase was to engage the user in an initial task which would require moving and placing the hand(s) in the environment. The second phase consisted of a "memory test" where users tried to reproduce the placement of their dominant hand without any visual feedback.

We used the input devices from our neurosurgical visualization system (*fig. 1*), which were designed to allow neurosurgeons to explore cross-sections of volumetric data [8]. In the current experiment, the doll's head controls the orientation and depth of a target object (*fig. 2, left*). The doll's head uses only four degrees of freedom for movement: the up-down and left-right translations are constrained so that the target object always stays centered on the screen. The plate tool controls the position and orientation of a blue semi-transparent rectangle on the screen.

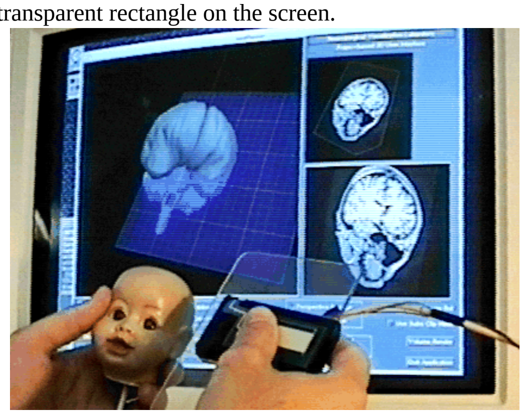

Figure 1: **Input devices from our neurosurgical visualization system (see also Color Plate 1).**

For the primary task, users were instructed to align and intersect the triangle and the plane so that they were coplanar (*fig. 2, right*). The triangle would highlight in yellow when the plane was aligned with it (the triangle appeared at a new random orientation for each trial). The interaction techniques used for this portion of the experiment were identical to those used in our neurosurgical visualization system, and have been informally evaluated and refined with hundreds of test users (including both physicians and non-physicians) prior to this experiment [8]. The precise mapping of input dimensions to virtual object motion is given the Appendix to this paper; the main idea of the two-handed interaction technique is to map motion of the preferred hand (holding the plate) relative to the nonpreferred hand (holding the doll's head) [10]. Our user testing as well as related theory on how people use two hands [7][11] suggests that a mapping which satisfies the property of preferred-hand relative to nonpreferred-hand reference should be the most natural mapping.

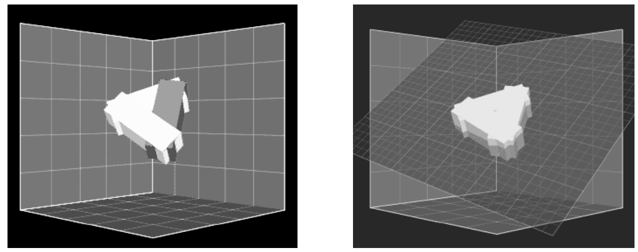

Figure 2: **Stimuli for the Primary task. The target object, an extruded triangle, is shown on the left. Users tried to move the plane so that it was roughly coplanar with the target object, causing it to highlight (right). The "stage" in the background served only as a perceptual aid and never moved. (See also Color Plate 2.)**

A footpedal was used as a clutch for the plate tool. When subjects held the pedal down, the plane could move freely relative to the target object. When the pedal was released, the plane would stay embedded in the target object. If the two were aligned when the pedal was released, this ended the primary task.

At the end of the primary task, the computer recorded the position and orientation of the preferred hand (which was always holding the plate tool). A dialog then appeared telling the subject to "Get Ready for Memory Test!" (*fig. 3*). For the memory test, subjects were instructed to put their preferred hand down on a mouse pad at the side of the work space, to close their eyes, and then to attempt to exactly reproduce the position and orientation of the plate tool without any visual feedback (*fig. 3*). At the end of each trial, the computer displayed the subject's best accuracy so far.

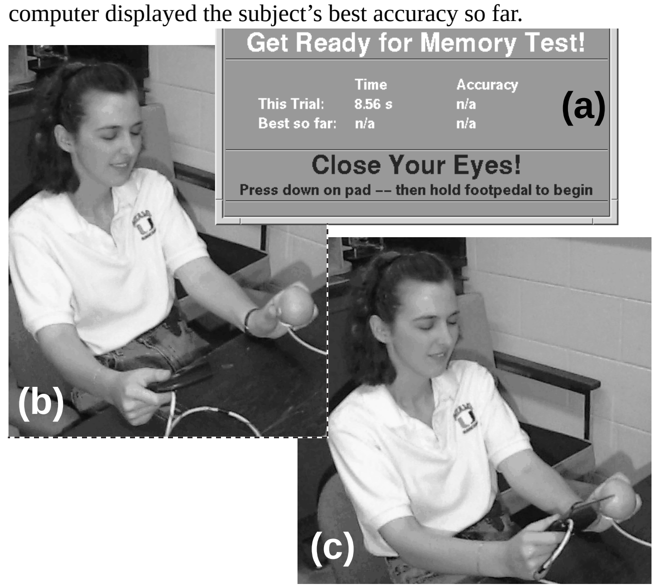

Figure 3: **The memory test. A dialog appeared telling subjects to get ready (*a*), after which they would close their eyes and place their dominant hand on a pad at the side of the working area (*b*). Subjects then attempted to reproduce their posture from the end of the primary task, without any visual feedback (*c*). (See also Color Plate 3.)**

3.3 Experimental Conditions

The experiment compared two conditions, a bimanual condition and a unimanual condition. In the bimanual condition, simultaneous motion of both input devices was possible. The doll's head was held in the nonpreferred hand, the plate in the preferred hand.

Since the distance between the two objects does not affect the alignment required for the primary task, it is always possible for the doll's head and the plate tool to make physical contact when completing the task. Subjects were instructed to use this technique in the bimanual condition, since we wanted to test how well subjects could use the nonpreferred hand as a reference. The bimanual condition is shown in figure 3.

In the unimanual condition, subjects were instructed to always keep their nonpreferred hand in their lap. Subjects were only allowed to grasp one device at a time, using only their preferred hand. There was a definite device acquisition time required when switching input devices, but subjects were instructed that time to complete the task was not important -- only their accuracy on the memory test mattered. For the memory test, the unimanual condition was identical to the bimanual condition, except that the nonpreferred hand was no longer available as a reference.

Clearly, both conditions utilized a space-multiplexed design, with a separate input device for each function, as opposed to a time-multiplexed design, where a single device controls multiple functions by changing modes. Brooks [2] reports that overloading a device with multiple functions can often cause confusion. Thus, we chose a space-multiplexed design for the unimanual condition because we did not want the possible issue of users becoming confused over which "mode" they were in to interfere with the experiment itself.

For the unimanual condition only, a second clutch footpedal was needed to allow subjects to rotate the doll's head and leave it "parked" at a particular orientation, thus allowing them to put down the doll's head and pick up the plate tool. Users had no difficulty in using the two pedals: there were no experimental trials where a user clicked the wrong pedal in the unimanual condition.

Originally, we had planned to use two footpedals in the bimanual condition as well, but in pilot studies we found this was problematic. If the footpedal is used to "park" the target object in the bimanual condition, the user is again moving relative to the environment, not relative to the reference frame specified by the nonpreferred hand. Once pilot subjects developed some experience with the task, they would essentially always hold down the second footpedal to maintain the doll's head as a reference. Thus in the bimanual case the second footpedal seemed to introduce confusion without adding any new or helpful capabilities.

This parallels our experience with usability tests of the original neurosurgical visualization system [8], where users tried to position a plane relative to a polygonal brain (analogous to the "target object" of this experiment). Freezing (or "clutching") the polygonal brain in place initially seemed like a useful thing to do, but we found that if the doll's head was clutched, it was no longer useful as a reference. I have watched many users clutch the head and then become confused as they subconsciously begin to move their nonpreferred hand to aid the action of the preferred hand, only to have no effect. In the current design of the interface, the polygonal brain is always allowed to move[1]. Bimanual manipulation seems to work best when there is no clutch for the base frame of reference (in the nonpreferred hand), which suggests a general design principle: always maintain the nonpreferred hand as a dynamic reference for the action of the preferred hand.

3.4 Experimental Procedure and Design

A within-subjects latin square design was used. Eight subjects (six female, two male) performed the unimanual condition first and nine subjects (seven female, two male) performed the bimanual condition first. Subjects performed 12 experimental trials for each condition.

During a practice session subjects were introduced to the equipment[2] and allowed to become familiar with it. We gradually introduced each element of the experimental procedure and made sure that subjects could perform the task before moving on. Practice sessions lasted 10-20 minutes (prior to the first experimental condition) and 5-10 minutes (prior to the second experimental condition). We did not have a fixed number of trials or set time limit for practice, but rather practiced with each subject until he or she felt completely comfortable with the equipment and experimental procedure.

4 RESULTS

Accuracy on the memory test was the only dependent measure. Accuracy was measured in terms of angle (shortest-arc rotation to align the remembered reference frame with the ideal reference frame) and distance (translational offset between the reference frames). Distance was also logged as single-axis offsets in the reference frame of the plate tool (offsets along the left-right axis, front-back axis, and up-down axis).

Table 1 reports the overall means obtained in each experimental condition. We performed an analysis of variance with repeated measures on the within-subjects factor of Condition (unimanual vs. bimanual). Condition was not a significant factor for angle but was highly significant for measures of distance (*table 2*).

Table 1: **Overall means obtained in each experimental condition.**

Accuracy metric	Bimanual	Unimanual
Angle (degrees)	11.2	10.4
Distance (mm)	16.6	32.6
Left-right distance (mm)	3.7	17.5
Up-down distance (mm)	8.3	14.2
Front-back distance (mm)	11.5	16.6

Table 2: **Significance levels for main effects.**

Accuracy metric	F statistic	Significance
Angle	$F_{(1,16)} = 0.74$	Not Significant
Distance	$F_{(1,16)} = 44.21$	$p < .0001$
Left-right distance	$F_{(1,16)} = 51.95$	$p < .0001$
Up-down distance	$F_{(1,16)} = 10.22$	$p < .006$
Front-back distance	$F_{(1,16)} = 12.04$	$p < .0035$

This evidence strongly supports hypothesis H1, suggesting that subjects were able to utilize the perceptual cues provided by the nonpreferred hand to reproduce their six degree-of-freedom posture independent of visual feedback. Subjects were significantly more accurate with both hands than with just one, supporting H2.

The analysis also revealed a significant Condition $\times$ Order interaction for Distance ($F_{(1,15)} = 9.09$, $p < .01$). It is often assumed that alternating the order of two conditions across subjects automatically controls for order effects caused by transfer of skill. But this is not true if there is a one way (or asymmetric) transfer of skill between the two conditions. A Condition $\times$ Order interaction is the statistical evidence for such an asymmetric transfer effect [19].

The means grouped by order (*table 3*) show this effect. When performing unimanual first, subjects' distance was 7% better on the subsequent bimanual condition than those subjects who completed bimanual first (16.0 mm vs. 17.2 mm). But when performing bimanual first, subjects performed 28% better on the subsequent

1. The interface can generate detailed still images for later reference using a "snapshot" command. Unlike clutching, this does not have the side-effect of interfering with further manipulation.

2. All experimental data was collected with a Polhemus FASTRAK [17].

unimanual condition than those subjects who completed unimanual first (27.6 mm vs. 38.2 mm). Our interpretation is that subjects learned a more effective task strategy in the bimanual condition, and were able to transfer some of this skill to the unimanual condition.

Table 3: **Means grouped by order of experimental conditions.**

	Bimanual First		Unimanual First	
Means by Order	**Bi**	**Uni**	**Bi**	**Uni**
Angle	11.5	9.8	11.0	11.0
Distance	**17.2**	**27.6**	**16.0**	**38.2**
Left-right dist.	4.4	13.8	3.0	21.6
Up-down dist.	7.8	12.1	8.9	16.6
Front-back dist.	12.3	15.0	10.6	18.4

Our qualitative observations also supported this position. When performing the unimanual condition first, subjects had a tendency to avoid using the doll's head: only 2 out of 8 of these subjects consistently reoriented the target object with the doll's head. Subjects would instead adapt the plate tool to the initial (randomly generated) orientation of the target object. But for 8 out of the 9 subjects who tried the bimanual condition first, during the unimanual condition they would re-orient the doll's head on essentially every trial. As one subject explained, during the bimanual condition she had learned that "instead of accepting what it gave me, I did better when I moved [the doll's head]."

All of this evidence supports H3, suggesting that bimanual control can affect performance at the cognitive level by influencing a subject's task-solving strategy. However, to definitively demonstrate that there is a bimanual to unimanual transfer effect, in future work we would like to perform a control study, comparing results from this experiment to subjects who perform two blocks of the unimanual condition or two blocks of the bimanual condition.

Finally, an analysis of signed distance errors supported H2: with just one hand, subjects could make unbiased estimates of a remembered hand position. By "unbiased" we mean that the means of the signed errors along each axis did not significantly differ from zero.

Table 4: **Means of signed distance errors.**

Accuracy metric	Mean (mm)	Std. Deviation
Unimanual		
Left-right distance	-1.1	16.6
Up-down distance	-5.8	11.6
Front-back distance	-4.6	11.2
Bimanual		
Left-right distance	-1.7	2.1
Up-down distance	-2.9	6.4
Front-back distance	+2.3	6.3

4.1 Qualitative Results

The unimanual condition seemed to impose two chief difficulties for users. First, when orienting the target object with the doll's head, the subject had to think ahead to the next step to anticipate which orientations of the target would be easiest for the action of the plate tool. Second, since the unimanual condition requires movement relative to the environment, the user had to remember approximately where he or she had "parked" the doll's head.

The bimanual condition avoids both of these difficulties. When using both hands, it is much easier to see what orientations of the target will be easy to select with the plane. As one subject commented, "it was easier to get them to come together, and faster too." Another noted that "two hands have much more flexibility for how you solve the problem." And since the doll's head is always part of the interaction, there is no need to remember where you "parked" it.

The bimanual condition does introduce the possibility of the two objects colliding. For example, if the plate is directly behind the doll's head and the subject needs to move it forward, he or she cannot do this directly since the doll's head is in the way. But the virtual plane is bigger than the physical plate, so one can solve this problem by (for example) holding the plate immediately to the right of the doll's head. On the screen, one sees the blue plane intersecting the target object even though the two input devices don't physically intersect. A couple of users initially found this to be confusing, but adapted after some practice.

We asked subjects about their task strategies and the cues they had used to perform the memory test. Subjects generally tried to hold as many of the variables constant as possible, and then memorized the rest. For example, subjects often kept the elbow and wrist angles fixed and would try to maintain an invariant hand posture with respect to the input device. Remaining variables such as the height and depth of the hand placement were then estimated from memory.

In the bimanual condition, subjects seemed to have an innate sense of where they had touched (either on the doll's head or on their hand)-- as one subject explained, "the touch knew the position"-- and many subjects thought of the angle of the plane as a separate thing to memorize. Subjects certainly also made use of the physical landmarks on the doll's head (such as the ears or the features of the face). Our impression is that without these landmarks, subjects probably would have been less accurate, but many subjects seemed to zero in on the exact spot even before physical contact was made. As one subject commented, "even before you touch the spot you know it."

In the Unimanual condition, the edge and surface of the desk served as a physical reference and most subjects tried to use this to their advantage, for example by resting their forearm against the desk and remembering the touch point. Subjects would often attempt to estimate the left-right placement of their hand using body-relative cues such as the torso midline or the positions of their legs. One subject commented that "because there was nothing to land on, your sort of lose your sense of balance." Another described her hand as "just floating in space, but using both hands gave you something else to reference."

5 DISCUSSION

Our results have clear design implications for the role of visual feedback and attention in human-computer interaction. Users maintain a fairly precise, body-relative representation of space which does not depend on visual feedback. A relatively inaccurate environment-relative representation of the space is also maintained. Our interpretation is that two hands and split attention go well together, opening up new possibilities for eyes-free interaction.

When using two hands, the user's attention does not necessarily have to constantly monitor the manipulation itself, and attention can be directed towards a secondary task, such as watching an animation or a representation of the manipulation from a second viewpoint. Unimanual control is more dependent on visual feedback and can therefore impede the user's ability to split attention between multiple tasks.

The Worlds-in-Miniature (WIM) interface metaphor [22] provides an example of these issues in action. The WIM provides the virtual reality user with a hand-held miniature representation of the immersive life-size world. Users interact with the WIM using both hands. The user's nonpreferred hand holds the WIM on a clipboard while the preferred hand holds a ball instrumented with some buttons. Moving a miniature object on the WIM with the ball moves the corresponding life-size representation of that object.

By using both hands, users can quickly develop a sense of the space represented by the WIM. We have even observed users manipulating objects in the WIM without looking at it (that is, users have manipulated objects while holding the clipboard below the field-of-view seen in the immersive display). This is convenient for tasks such as hanging a picture, where it is useful to manipulate objects in the WIM while attending to the 1:1 scale view to check the picture's alignment and to evaluate how it fits the architectural space.

This is perhaps the strongest qualitative evidence that interaction techniques based on hand-relative-to-hand manipulation can allow users to focus attention on their tasks without necessarily becoming distracted by the interface technology.

5.1 Task Analysis

One might argue that using two hands only adds complexity and makes an interface harder, not easier, to use-- after all, it is difficult to "rub your head and pat your stomach at the same time." Rubbing one's head while patting one's stomach is indeed an example of a difficult two-handed task because the subtasks assigned to each hand are completely independent subtasks. There are many compound tasks, however, such as navigation and selection in a text document or positioning and scaling a rectangle, which users perceive as integral attributes [12] that are aspects of a single cognitive chunk [5]. When designed appropriately, a two handed interface for integral compound tasks does not necessarily impose a cognitive burden, and can help users to reason about their tasks.

Figure 4 illustrates the task hierarchy for selecting a cutting-plane relative to a specific view of a target object, as required by the primary task for this experiment. Cutting relative to a view consists of two sub-tasks: viewing and cutting. Viewing can further be subdivided into orienting the target object and specifying a zoom factor, and so forth. At the lowest level, there are ten separate control parameters (yaw, pitch, roll, and zoom for the view; x, y, z, yaw, pitch, and roll for the cutting plane) that can be 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 user interface which would be difficult to use without training and practice. A two-handed interface with multiple degree-of-freedom input devices in each hand, however, reduces this entire hierarchy into a single transaction (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.

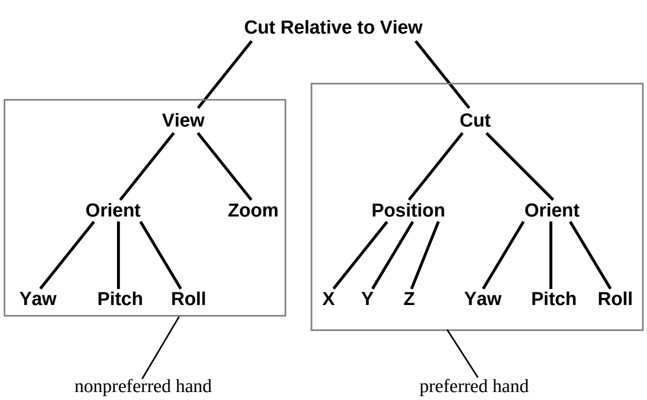

Figure 4: **Task hierarchy for selecting a cut relative to a specific view.**

This framework, suggested by Buxton's work on chunking and phrasing [5], is useful for reasoning about the differences between one and two-handed interfaces, and helps to suggest why the present experiment found a transfer effect from the bimanual to the unimanual interfaces tested. With a unimanual interface, View and Cut must be performed as purely sequential subtasks. There is also the need to switch back and forth between viewing and cutting, so this implies a third sub-task, that of changing modes. Changing modes might involve acquiring another input device (as in this study), speaking a voice command, or (in 2D interaction) moving a mouse to another region of the screen, but regardless, all of these mode switching techniques take a non-zero amount of time. This process can be modelled as a simple state diagram (*fig. 5*).

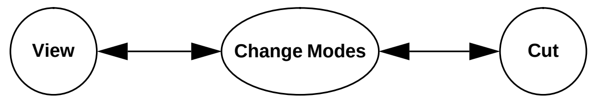

Figure 5: **State diagram for unimanual subtasks.**

A two-handed interface changes the syntax for this task. Under bimanual control, a new meta-task with a single *Cut Relative to View* state becomes possible. The simultaneous Cut Relative to View task is not the same thing as the serial combination of the sub-tasks. The simultaneous task allows for hierarchical specialization of the hands [7], and there is no cost (or need) to switch between *View* and *Cut* subtasks. Thus, there is the potential for bimanual control to impact performance at the cognitive level: it can change how users think about the task. Since the *View* and *Cut* subtasks can be integrated without cost, this encourages exploration of the task solution space. And since the user never has to engage in a *Change Modes* sub-task, there is no possibility for this extraneous sub-task to interfere with the main goal of viewing and cutting.

6 CONCLUSION AND FUTURE WORK

This study has provided some initial evidence which helps to support the claim that using both hands can help users gain a better sense of the space they are working in [6]. Immersive VR systems which use just one hand often do not offer users any physical reference points. Making use of both hands provides a simple way to increase the degree of manipulation and to let the user's own hand act as a physical reference. Another technique is to introduce a grounding object such as a drafting table; but even in this situation, using two hands plus the grounding object allows interesting design possibilities [1].

Our second high-level hypothesis is that in some cases using both hands can change the way users think about a task. This experiment also provided some initial evidence in favor of this hypothesis, suggesting that it may be easier to explore alternative strategies for problem solving when both hands can be used. This was reinforced by our qualitative observation that in the unimanual condition, subjects were more likely to "take what they were given" because they had difficulty anticipating which orientation of the target object would be easy to select with the plane.

The input devices used in this study were rich in tactile orientation cues and this certainly helped subjects to perform the experimental task more precisely. If we had used featureless spheres as input devices, for example, subjects probably would have had a less acute sense of the orientation of each device. We also believe that allowing contact between the two hands was a factor in the experiment, but not the only factor. When using two hands, subjects could often come quite close to the original position even before contact was established. Further study is necessary to determine if this differs significantly from moving a single hand relative to the environment.

As a thought experiment, one can imagine using a single hand to move the plate tool relative to a doll's head mounted on a pivot. This would be analogous to using one hand on a tablet fitted with a physical template, which works well [4]. But our experimental findings suggest that the dynamic role of the nonpreferred hand also led to a cognitive performance benefit in terms of task strategy chosen. The task syntax supported by moving one hand relative to a reference object on a pivot is quite similar to that required by our unimanual condition. As such we expect using the pivot with just one hand would have some of the same limitations: users might have difficulty anticipating what orientation of the pivot object would be most facile for the action of the plate tool.

7 REFERENCES

1. Abel, K., Alic, M., Moshell, J. M., "Rapid Construction of Model Building for Urban Combat Mission Rehearsal Databases (PolyShop[TM])." URL: http://www.vsl.ist.ucf.edu/~polyshop/polyshop.html
2. Brooks, F. P. Jr., "Grasping Reality Through Illusion: Interactive Graphics Serving Science," Proc. ACM CHI'88 Conf. on Human Factors in Computing Systems, pp. 1-11.
3. Buxton, W., Myers, B., "A Study in Two-Handed Input," Proc. CHI'86, pp. 321-326.
4. Buxton, W., Hill, R., Rowley, P., "Issues and Techniques in Touch-Sensitive Tablet Input," Computer Graphics, July 1985, pp. 215-224.
5. Buxton, W., "Chunking and Phrasing and the Design of Human-computer Dialogues," Kugler, H. J., ed., Information Processing '86, Proc. of the IFIP 10th World Computer Congress, Amsterdam, North Holland Publishers, 1986, pp. 475-480.
6. Fitzmaurice, G., Ishii, H., Buxton, W., "Bricks: Laying the Foundations for Graspable User Interfaces," Proc. CHI'95, pp. 442-449.
7. Guiard, Y., "Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model," The Journal of Motor Behavior, 19 (4), 1987, pp. 486-517.
8. Hinckley, K., Pausch, R, Goble, J., Kassell, N., "Passive Real-World Interface Props for Neurosurgical Visualization," CHI'94, pp. 452-458.
9. Hinckley, Pausch, Goble, Kassell, "A Survey of Design Issues in Spatial Input," UIST'94, pp. 213-222.
10. Hinckley, K., *Haptic Issues for Virtual Manipulation.* PhD Dissertation, University of Virginia, Department of Computer Science, Charlottesville, VA, January 1997.
11. Hinckley, K., Pausch, R., Proffitt, D., Patten, J., Kassell, N., "Cooperative Bimanual Action," CHI'97 Conference on Human Factors in Computing Systems, to appear.
12. Jacob, R., Sibert, L., "The Perceptual Structure of Multidimensional Input Device Selection," Proc. ACM CHI'92 Human Factors in Computing Systems Conference, 1992, pp. 211-218.
13. Kabbash, P., Buxton, W., Sellen, A., "Two-Handed Input in a Compound Task," CHI'94, pp. 417-423.
14. Leganchuk, A., Zhai, S. Buxton, W., "Bimanual direct manipulation in area sweeping tasks," 1995. http://www.dgp.toronto.edu/people/andrea/bimanual.html.
15. Mine, M., http://www.cs.unc.edu/~mine/chimp.html.
16. Multigen Inc., http://www.multigen.com/smart.htm.
17. Polhemus Navigation Sciences, Inc., (802) 655-3159.
18. Poston, T., Serra, L., "Dextrous Virtual Work," Communications of the ACM, 39 (5), pp. 37-45, 1996.
19. Poulton, E., Freeman, P., "Unwanted asymmetrical transfer effects" Psychlogical Bulletin, 66 (1), 1966.
20. Sachs, E., Roberts, A., Stoops, D., "3-Draw: A Tool for Designing 3D Shapes," IEEE CG&A, Nov. 1991.
21. Shaw, C., Green, M., "Two-Handed Polygonal Surface Design," Proc. UIST'94, pp. 205-212.
22. Stoakley, Conway, Pausch, "Virtual Reality on a WIM: Interactive Worlds in Miniature," CHI'95.

8 Appendix: Mapping of Input to Output

The two-handed interaction technique used in this experiment (and in our neurosurgical visualization system [8]) assumes a natural central object which is being manipulated. The key design principle is *not* to maintain a direct 1:1 correspondence between physical and virtual object motion, but rather it is to maintain the nonpreferred hand as a dynamic frame-of-reference.

Users do expect the real-world relationship between the input devices to be mirrored by their on-screen graphical representations. Simplifying control of the target object by centering it on the screen, however, requires a software mapping of its real-world position to its centered position by constraining the x and y translations (note that no such mapping is required for the *orientation* of the prop). Define the position of the doll's head in the real world as (H_{Rx}, H_{Ry}, H_{Rz}). If the center point of the screen is defined as (C_x, C_y), then the *virtual* constrained head position is given by (H_{Vx}, H_{Vy}, H_{Vz}) = (C_x, C_y, H_{Rz}).

When the user moves the plate tool relative to the doll's head, the user expects to see this relative motion mirrored on the screen. This implies that the virtual representation of the plate tool is drawn relative to the virtual position of the doll's head. That is, the virtual position of the plane is equal to the virtual position of the *head* plus the real-world offset between the head and the plate tool. Define the position of the plate tool as (P_{Rx}, P_{Ry}, P_{Rz}). The offset is:

$$\left(\Delta_x, \Delta_y, \Delta_z\right) = \left(H_{R_x}, H_{R_y}, H_{R_z}\right) - \left(P_{R_x}, P_{R_y}, P_{R_z}\right)$$

The virtual position of the plane is then given by:

$$\left(P_{V_x}, P_{V_y}, P_{V_z}\right) = \left(H_{V_x}, H_{V_y}, H_{V_z}\right) + \left(\Delta_x, \Delta_y, \Delta_z\right)$$

This mapping results in the following non-correspondence artifact: if the user holds the plate tool still and translates only the doll's head, the target object will remain centered and the virtual plane will move in the *opposite* direction. This violates the generally accepted design principle that an interface should always maintain a direct 1:1 correspondence between physical and virtual object motion. But it adheres to the design principle that the object in the nonpreferred hand (the doll's head) should form a base frame of reference relative to which the preferred hand moves. In hundreds of informal user trials with our neurosurgical visualization system, we have found that users almost never discover this artifact, because they typically hold and orient the doll's head in a relatively stable location while moving the plate tool relative to it. The net effect is that the interaction behaves as users expect it would; the mapping is the software embodiment of the principle that the nonpreferred hand sets the frame of reference while the preferred hand articulates its motion *relative to* the nonpreferred hand [7].

Centering the reference object also has some other subtle effects. Since the nonpreferred hand now defines a dynamic frame of reference relative to which all manipulation occurs, this means that the user is not forced to work relative to the screen itself or relative to some center point within the environment. Users are free to shift their body posture or to hold their hands in a natural working posture. There is also no need for a "homing" command to move the center point, since the nonpreferred hand automatically and continuously performs this function just by holding the doll's head.