THE PERCEPTUAL STRUCTURE OF MULTIDIMENSIONAL INPUT DEVICE SELECTION

Robert J.K. Jacob
Linda E. Sibert

Human-Computer Interaction Lab
Naval Research Laboratory
Washington, D.C.

jacob@itd.nrl.navy.mil
sibert@itd.nrl.navy.mil

ABSTRACT

Concepts such as the logical device, taxonomies, and other descriptive frameworks have improved understanding of input devices but ignored or else treated informally their pragmatic qualities, which are fundamental to selection of input devices for tasks. We seek the greater leverage of a predictive theoretical framework by basing our investigation of three-dimensional vs. two-dimensional input devices on Garner's theory of processing of perceptual structure in multidimensional space. We hypothesize that perceptual structure provides a key to understanding performance of multidimensional input devices on multidimensional tasks. Two three-dimensional tasks may seem equivalent, but if they involve different types of perceptual spaces, they should be assigned correspondingly different input devices. Our experiment supports this hypothesis and thus both indicates when to use three-dimensional input devices and gives credence to our theoretical basis for this indication.

KEYWORDS: Input devices, interaction techniques, gesture input, Polhemus tracker, perceptual space, integrality, separability.

INTRODUCTION

In studying interaction techniques, each new piece of hardware that appears raises the question *What tasks is this device good for, and how should it be incorporated into interface designs?* Such questions are typically answered specifically for each new device, based on the intuition and judgment of designers and, perhaps, on empirical studies of that device. Greater leverage would be achieved if such questions could be answered by reasoning from a more general predictive theoretical framework, rather than in an *ad hoc* way. We provide one example of how the answer to this question can be derived from a theoretical framework (as, for example, Card, Moran, and Newell [5] have done for pointing devices).

1992 ACM 0-89791-513-5/92/0005-0211

We begin by posing the question for the three-dimensional position tracker, such as the Polhemus 3SPACE or Ascension Bird trackers. While directly answering the question *What is a three-dimensional tracker good for?* we will also try to shed light on the next level question, i.e., *How should you answer questions like* What is a three-dimensional tracker good for?

Concepts such as the logical input device (discussed further below) provide descriptive models for understanding input devices, but they tend to ignore the crucial pragmatic aspects of haptic input by treating devices that output the same information as equivalent, despite the different subjective qualities they present to the user. Taxonomies and other frameworks for understanding input devices have tended to hide these pragmatic qualities or else relegate them to a "miscellaneous" category, without further structure.

Instead, we draw on the theory of processing of perceptual structure in multidimensional space [8,9]. The attributes of objects in multidimensional spaces can have different dominant perceptual structures. The nature of that structure, that is, the way in which the dimensions of the space combine perceptually, affects how an observer perceives an object. We posit that this distinction between perceptual structures provides a key to understanding performance of multidimensional input devices on multidimensional tasks. Hence two three-dimensional tasks may seem equivalent, but if they involve different types of perceptual spaces, they should be assigned to correspondingly different input devices.

The three-dimensional position tracker can be viewed as a three-dimensional absolute-position mouse or data tablet; it provides continuous reports of its position in three-space relative to a user-defined origin. (In fact, the Polhemus and Ascension devices also report their orientation, but we have focused on position only in the present study.) The device thus allows a user to input three coordinates or data values simultaneously and to input changes that cut across all three coordinate axes in a single operation. (A mouse or trackball allows this in only two dimensions.) Such a device is obviously useful for pointing in three-space, but it is also applicable in many other situations that involve changing three values simultaneously. We concentrated on task spaces that do not directly map to three-dimensional physical space, and on manipulation rather than pointing tasks, in order to frame a balanced comparison. Specifically, we considered two tasks that both involve three degrees of freedom, i.e., that require adjusting three variables. For comparison with the three-dimensional tracker, we used a conventional mouse (for two of the three variables in the tasks) and then provided a mode change button to turn the mouse temporarily into a one-dimensional slider for the third variable.

A naive view of these two alternatives suggests that the three-dimensional tracker is a superset of the two-dimensional mouse, since it provides the same two outputs plus a third. Thus the three-dimensional tracker should always be used in place of a mouse (assuming ideal devices with equal cost and equal accuracy), since it is always at least as good and sometimes better. Our intuition tells us that this is unlikely—but why? The goal of this research is to develop a firmer foundation from which to draw such judgments.

APPROACH

To do this, we extend Garner's theory of processing of perceptual structure [9], first developed with fixed images, to interactive graphical manipulation tasks and thereby use it to shed light on the selection of multidimensional input devices. Our objectives are thus twofold:

- To answer the original question, *What is a three-dimensional tracker good for?* In so doing, we hope to provide by example a step toward putting the study of multidimensional input devices on a firmer theoretical footing.
- To extend Garner's theory of perceptual space from perception of fixed stimuli to interaction techniques.

Our hypothesis is that the structure of the perceptual space of the interaction task should mirror that of the control space of the input device.

BACKGROUND

The present work builds on two separate threads of research. One is the understanding of input devices in human-computer interaction, from the logical device concept through more recent taxonomies and tools. The other is the theory of perception of relationships between dimensions of a multidimensional space.

Study of Input Devices

A number of frameworks have been proposed to organize knowledge about input devices. An early abstraction is that of the logical device found in device-independent graphics packages based on standards such as ACM's Core Graphics System [10]. Fundamental user actions form the basis for the logical device equivalences. The idea is to separate the interaction device from the code needed to handle it in order to gain flexibility and facilitate rapid prototyping. However, experience shows that a system configured with a joystick is quite different from the same system with a trackball, although the logical device approach considers both equivalently as locators.

Foley, Wallace, and Chan's taxonomy of interaction techniques improves upon the logical device concept [7]. It maps elementary interaction tasks to the devices that can perform those functions but adds a middle layer that makes explicit the fact that there are many interaction techniques with which to perform a given elementary task. However, because the taxonomy is organized by task, a device can appear at more than one leaf in the techniques trees, hiding the structure of the device space as well as the qualitative differences that differentiate devices.

Buxton [4] calls these qualitative differences pragmatic attributes. He developed a taxonomy that organizes continuous input devices into a two-dimensional space whose dimensions are property and the number of dimensions sensed. In this way, the similarities and differences in the structure of the devices are highlighted. A tablet, light pen, and two-dimensional joystick are two-dimensional and sense position, but they differ from a two-dimensional trackball because it senses motion. This approach can thus point out that substituting a trackball for a joystick is incorrect; but it does not explain why it results in an awkward interface.

Mackinlay, Card, and Robertson [6,13] expand Buxton's taxonomy into a methodology consisting of three parts: a three-dimensional design space of most input devices; a functional mapping of information from the raw transducers of an input device into the semantics of the application; and evaluation techniques for comparing alternative designs in terms of expressiveness and effectiveness. This scheme, along with a defined set of connectors, allows continuous and discrete devices to be described and simple input devices to be combined into complex controls. Their approach both furthers our understanding of the structure of input device space and recognizes that human performance issues are important to understanding how a device actually works in a given situation. Although some relatively straightforward human factors issues are handled formally, such as matching the size of the domain and range of a value, the more subtle pragmatics of input device usage and task characteristics are still handled by a set of specific rules.

Bleser [3] developed a device taxonomy and input model that explicitly incorporates the physical attributes of input devices, including the notion of the physical separability of input degrees of freedom, and knowledge about task requirements. The taxonomy and model are used in an interactive design tool to suggest one or more natural interaction techniques based on a description of an interaction task [2]. The model depends on a set of heuristic rules and a pattern matching procedure rather than a more general, theoretical framework, but it highlights the need for such information in the design process.

Processing of Perceptual Structure

A multidimensional object is characterized by its attributes. A red circle has size, color, shape, and location, and these attributes define a perceptual space. Garner [9] observed that relationships between the attributes of an object can be perceived in two ways that differ in how well the component attributes remain identifiable. Some attributes are *integrally* related to one another—the values of these attributes combine to form a single composite perception in the observer's mind, and each object is seen as a unitary whole; while other attributes are *separably* related—the attributes remain distinct, and the observer does not integrate them, but sees an object as a collection of attributes. For example, value and chroma are perceived integrally, while size and lightness are perceived separably [11]. The horizontal and vertical positions of a single dot in the middle of an outline square are integral [9], while color and shape are separable [12]. This leads to two classes of perceptual space, one whose coordinate axes are perceived integrally, and one, separably, although there is really a continuum rather than a sharp dichotomy. There are two operational methods to measure integrality or separability. First, subjective judgments of similarity between stimuli in an integral space (as measured in a direct scaling experiment) obey a Euclidean distance metric, as if the observer is responding to the overall composite distance between the points in the space; while similarity judgments in a separable space obey a city-block metric, as if the observer perceives the similarity along each axis separately. Second, in a classification task, where an observer groups objects into classes to maximize perceived difference between classes and perceived similarity within classes, integral objects are grouped by their overall similarity, while separable objects are grouped by their individual attributes.

We can extend the notion of integral and separable attributes to interactive tasks. In an interactive graphical computer system, the user varies attributes of a fixed visual object over time. Interaction is thus simply movement through the perceptual space to vary the attributes of an object. Integral movement is Euclidean and cuts across the dimensions defined by the attributes; separable movement is city-block and moves parallel to the axes of the space.

We can view the control spaces of input devices in a similar way. Input devices with more than one degree of freedom can be characterized as integral or separable based on whether it is natural (or possible) to move "diagonally" across the dimensions. With an integral device, movement is in Euclidean space and cuts across all the dimensions of control. A separable device constrains movement to a stair-step pattern; movement occurs along one dimension at a time.

METHOD

Our hypothesis is that the structure of the perceptual space of an interaction task should mirror that of the control space of its input device. To examine it, we considered two interactive tasks, one set within an integral space and one in a separable one, and two input devices, one with integral dimensions and one, separable. This yields a two by two experiment, with four conditions. We expect performance on each task to be superior in the condition where the device matches that task in integrality/separability. That is, the interaction effect between choice of task and choice of device should far exceed the main effects of task or device alone.

For the integral three-attribute task in the experiment, the user manipulates the *x-y* location and the size of an object to match a target, since location and size tend to be perceived as integral attributes, as observed in studies of sets of fixed stimuli [9]. For the separable task, the user manipulates the *x-y* location and color (lightness or darkness of greyscale) of an object to match a target, since location and color are perceived separably [9]. The difference in perceptual structure between these two tasks is in the relationship of the third dimension (size or greyscale) to the first two (x and y location); in all cases, the x and y attributes are integral.

For the integral device condition, we use a Polhemus tracker, which permits input of three integral values. For the separable condition, we use a conventional mouse, which permits two integral values, to which we added a mode change to enable input of a third—separable—value. Our hypothesis predicts that the three degree of freedom input device will be superior to the two degree of freedom (plus mode change) device only when the task involves three integral values, rather than in all cases, as with the naive hypothesis mentioned above.

Design

We used a repeated measures design, with each subject performing all four experimental conditions. Forty subjects (26 men and 14 women) were assigned randomly to four groups for different presentation orders. Order was counterbalanced for practice and fatigue. Within each group, one task is performed with both devices before the second task is presented, to emphasize device differences within task and control for variability caused by practice with a device. Subjects were clerical, administrative, and technical personnel from the Information Technology Division of the Naval Research Laboratory who volunteered to participate without compensation.

Stimuli and Apparatus

A stimulus is a pair consisting of a user-controllable object and a target object. The user is asked to adjust the moveable object to match the target on each trial. For the size task, the moveable object is a solid square whose intensity is 50% of the greyscale range and whose location and size are adjustable (see Figure 1). The target object is a black outline square showing the desired size and position. The range of possible sizes varied from 0.7 to 6.2 inches on a side.

For the greyscale task, the moveable object is a square of size 2.8 inches (the midpoint of the size range), containing an embedded circle 1.5 inches in diameter whose color is adjustable and an outer region displaying the target color (see Figure 2). Color varies between 10% and 90% of the greyscale range. The target object is a black outline square the same size as the moveable object, which presents the position to be matched; the outer area of the moveable object gives the color to be matched. In both tasks, the moveable object is translucent so it never obscures the target.

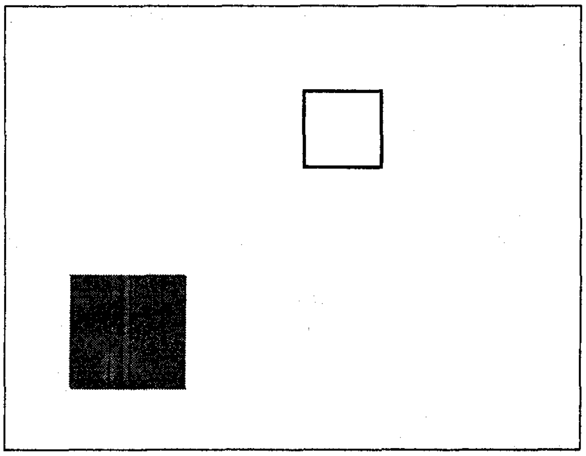

Figure 1. Stimulus for the integral (size) task. The outline square is the target. The user adjusts the location and size of the solid grey square to match the target.

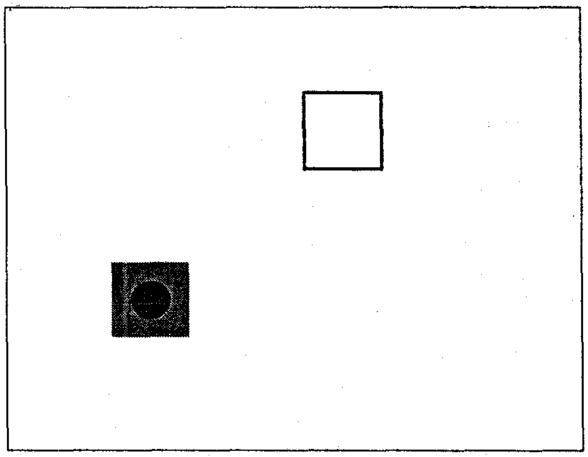

Figure 2. Stimulus for the separable (greyscale) task. The outline square gives the target location, and the outer area of the solid square gives the target color. The user adjusts the location of the solid grey square to match the target outline and the color of the inner circle on the grey square to match that of the outer area.

The maximum possible space of stimuli is a cube 13.75 inches on a side. However, not all of these targets can be used, because we require that the entire target square fit within the 13.75 inch by 11 inch screen and because we reserve colors near black and white for the target outline and screen background. Since these restrictions reduce the stimulus space differently for each task, we further adjusted the two spaces by matching their centroids, aligning the ranges of size and greyscale, equating the extent of hand motion needed to traverse them, and adjusting the monitor gamma correction to enhance greyscale range.

The stimulus set for each of the four conditions was randomly generated, but constrained by a script so that the distances in three-space from one trial to another were the same across conditions. Corresponding trials thus have equal distances, facilitating comparison across conditions. In addition, we required that each stimulus differ from the preceding one by at least 0.5 inches in each dimension, to avoid degenerate trials, which the user could complete without exercising all three dimensions of motion. Each stimulus trial is a three-dimensional point (x,y,z), and the software is identical except for how the z dimension is displayed on the screen. The x and y dimensions of both tasks are mapped as position, but z is mapped as either size or greyscale color.

The three-dimensional tracker used in the experiment was the Polhemus 3SPACE magnetic tracker. It consists of a transmitter that generates electromagnetic fields and a wand housing three orthogonal coils that sense the fields. The position and orientation of the wand is transmitted to a host computer. Only the three position values were used in this experiment. The Polhemus is an absolute device with the origin located at what would be the forward end of an armrest of the chair. The source was permanently fixed under the subject's chair. The control-display ratio for the Polhemus was one inch of device movement to one inch of screen movement. The data from the Polhemus was filtered with a small moving average filter to help smooth out device response. To improve the performance of the Polhemus, we eliminated as much metal as possible from the surrounding area. We used wooden furniture, and the metal mouse pad and its table were removed when not in use. The Polhemus was calibrated in one of two standard ways (one for left handed use, one for right) so that the operating area was consistent across subjects and its axes were orthogonal. Movement in the plane parallel to the screen moved the cursor in x and y. Moving the Polhemus toward screen made the object either bigger or darker, away made it smaller or lighter.

The mouse used was the standard optical mouse supplied with the Silicon Graphics workstation. It is a relative device, and the mouse pad was located under the preferred hand of the subject. The x and y mouse coordinates, combined with a mode change into a one-dimensional slider, provided the three dimensions of control for the experiment. We chose this control strategy because it is similar to those currently used in applications with multiple variables. Movement over the optical pad corresponded to movement in x and y. Holding any button down and moving the mouse toward the screen made the object either bigger or darker, away made it smaller or lighter.

The computer was a two-processor (16 MHz) Silicon Graphics Iris workstation, model 4D/120G. The program was divided into two processes, which ran concurrently on the two processors. One continuously monitored the Polhemus over a serial port and fed data into an event queue in shared memory, while the other drew the images and supervised the experiment. Most other system processes and all network daemons were eliminated. This architecture was effective in greatly reducing the often-observed lag in response to movements of the Polhemus. Position data from the mouse or Polhemus were recorded continuously throughout the experiment, approximately every 20 milliseconds. The monitor was a 19-inch Hitachi running under a gamma correction of 2.7.

Procedure

The subject was seated in a chair without arms, in front of a monitor on a desk 29.5 inches high, with either the mouse (placed on a small, moveable table) or the Polhemus in the subject's preferred hand and a 5.5 inch square flat button located on another small table under the non-preferred hand. The chair was located 36.25 inches away from the desk and the subject's eyes were approximately 56 inches from the monitor. This distance allowed enough room for the subject to manipulate the Polhemus and reduced the interference between the electromagnetic fields of the monitor and the Polhemus. The experimenter sat at a terminal in the rear of the room. The experiment was conducted in a special purpose laboratory designed for such work [1].

Each trial consisted of changing the position and either size or greyscale color of a moveable object using one of the two devices until it matched the position and size or color of the target. The subject pushed the button located under the non-preferred hand when the match was considered good enough. For each condition, the subject was presented first with 33 practice trials, followed by 88 experiment trials. The trials were subdivided into sets of 11, with the first in each set not scored because it measured the time to home the moveable object from an uncontrolled starting position. The home position for the first trial in each set was located in the middle of the screen and was either mid-sized or mid-colored. Within a set, as soon as the subject pushed the button indicating a target had been matched, the trial was ended and another target presented. An instruction screen separated the sets of trials and a subject determined when to start a set. The subjects were instructed to rest as long as they needed between sets of trials.

The subjects were encouraged to ask questions during practice but not during the experimental trials. They were instructed that accuracy and speed were of equal importance. Subjects completed a short questionnaire at the conclusion of the experiment. Each subject took approximately 1.5 hours to complete the experiment, which was run during March and April, 1991.

RESULTS

The data were analyzed in two ways: the first and simpler one measures overall time per trial; the second combines speed and accuracy into a single measure. The mean time per trial (time from appearance of the target until the subject pressed the button) for each of the four experimental conditions is shown in Figure 3 and graphed with dotted lines in Figure 4. The data suggest that neither task nor device alone produces as large an effect as the interaction of the two: that is, matching the integrality/separability of the device to that of the task gives superior performance. To evaluate this observation, we performed a repeated-measures analysis of variance on these data. First, we checked the effect of order of presentation of the four conditions, found no significant effect for experimental group ($F(3,36) = 1.09$, $p > .30$), and thus aggregated the groups in further analyses. There were significant effects for both task ($F(1,39) = 7.40$, $p < .01$) and device ($F(1,39) = 7.80$, $p < .01$) and a highly significant effect for interaction between task and device ($F(1,39) = 69.34$, $p < .0001$), as predicted.

Task:	Device: Integral (Polhemus)	Device: Separable (Mouse)
Integral (Size)	4981 *(2065)*	6274 *(2518)*
Separable (Grey)	5357 *(1613)*	4838 *(1269)*

Figure 3. Mean time per trial in msec. *(and standard deviation).*

The second analysis measures the time required to reach a fixed accuracy criterion on each trial (where accuracy is overall Euclidean distance to target in three-dimensional space). This combines speed and accuracy into a single measure and removes the effect of individual subjects' "personal" accuracy criteria for terminating trials. To facilitate this analysis, we had recorded the position of the mouse or Polhemus approximately every 20 ms. during each trial. This enabled us to simulate retroactively an experiment in which the subject would have been required to reach a certain accuracy criterion, which would then automatically terminate the trial. Since a subject may briefly, inadvertently pass through a point that corresponds to good accuracy, retroactive analysis allows us to correct this by measuring the time until the subject reached our criterion for the *last* time during a trial. Figure 5 and the solid lines in Figure 4 show the results of this analysis, where the criterion was set at the 75th percentile of the final accuracies actually achieved over all trials, conditions, and subjects (which was 0.11 inches in three-dimensional Euclidean distance to the target). The results are similar to those in Figure 3. (The choice of 75th percentile accuracy is not critical; analysis with other criteria gave similar results.) An analysis of variance, with repeated measures, shows a significant effect for task ($F(1,39) = 5.78$, $p < .05$) but not device ($F(1,39) = 1.91$, $p > .2$), and, again, a highly significant effect for interaction between task and device ($F(1,39) = 52.07$, $p < .0001$).

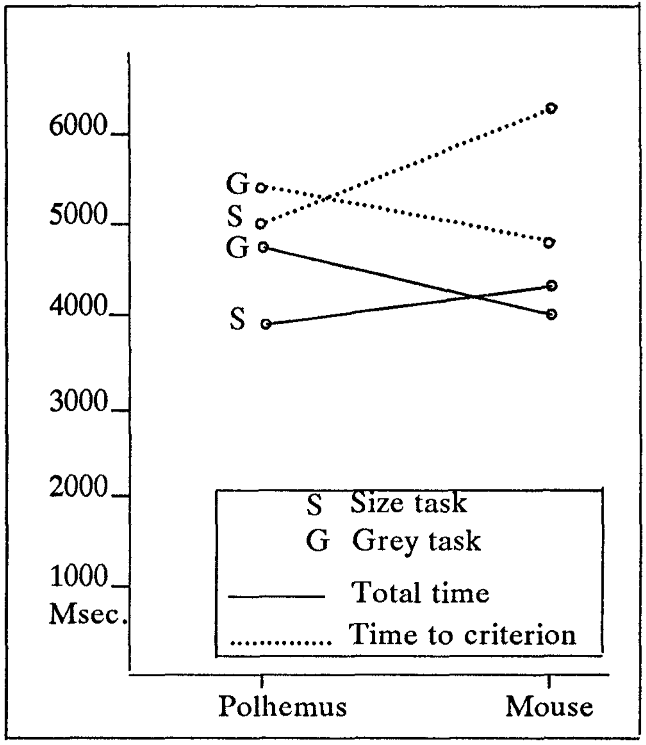

Figure 4. Graph of mean times in msec., illustrating interaction effect. This figure contains two graphs superimposed on the same axes; the two dotted lines show the results for total time per trial (as in Figure 3); the two solid lines, time to criterion (as in Figure 5). Lines marked **S** show performance on the integral (size) task; **G**, the separable (grey) task.

Our questionnaire asked how subjects felt about the experiment and for their prior experience with various input devices. Subjects preferred the mouse to the Polhemus and thought they performed faster and more accurately with it, but significantly more so on the greyscale task than on the size task (t test at $p < .05$), as expected. Overall, however, they preferred the mouse on both tasks, despite the fact that they performed worse with it on the size task. They considered the Polhemus more tiring for both tasks and the mouse a little easier to learn for both tasks. The input device familiarity question suggests why the mouse was easier to learn: 37 of the 40 subjects used one daily, while only two had ever used a Polhemus or other three-dimensional tracker. What was surprising is that the Polhemus did not seem to be hard to learn even though the device was new to most subjects, and some subjects were very enthusiastic about the Polhemus. Data from the practice trials show that subjects developed good performance with the Polhemus relatively quickly. While it is an inherently less stable device and might initially cause strain, some highly practiced users in pilot runs reported that the more they used the Polhemus, the less tiring and more natural it became.

Task:	Device: Integral (Polhemus)	Separable (Mouse)
Integral (Size)	3892 *(1430)*	4320 *(1323)*
Separable (Grey)	4739 *(1298)*	4000 *(893)*

Figure 5. Mean time in msec. to reach criterion accuracy *(and standard deviation),* where criterion was the 75th percentile accuracy attained over all the trials.

DISCUSSION

The results converge toward the conclusion that neither device is uniformly superior to the other in performance. Instead, we find significantly better performance in the experimental conditions where the task and device are both integral or both separable and inferior performance in the other two conditions. These results support our extension of the theory of perceptual space to interaction techniques, which predicts that the integral task (size) will be performed better with the integral device (Polhemus) and that the separable task (greyscale) will be performed better with the separable device (mouse).

Design rules for input device selection thus cannot be written from an examination of the input device space alone but must match the perceptual structure of the task space with that of the control space of the device used to perform the task.

APPLICATION

How might these results be used in designing controls for zooming and panning of a geographic display? Zooming and panning, taken together, involve three degrees of freedom. The most common design uses a mouse or trackball for two-dimensional panning and a separate control for zooming. We claim that a user typically does not really think of zooming or panning operations separably, but thinks rather of integral operations like "focus in on *that* area over there." The space is thus Euclidean, like that of the size task in the experiment, and, therefore, making the user do the two separately violates perceptual compatibility. It would be more natural to permit a user to make a gesture that performs the overall operation he or she had in mind, using an integral three-dimensional input device. The user moves the puck around in a volume directly in front of the display screen. Moving it in the *x* or *y* direction parallel to the display surface causes panning; moving it perpendicular to the display (directly toward or away from it) causes zooming. The user typically moves the puck in all three dimensions simultaneously, resulting in some combination of zooming and panning and directly reaches the view of interest. We have successfully demonstrated a mockup of this application.

CONCLUSIONS

This work shows how understanding of multidimensional input devices can be placed on a firmer theoretical footing. It introduces in a more formal way two important considerations in input device selection, which heretofore have not been so treated:

- the nature of task space—more specifically, the way its dimensions compose perceptually;
- the pragmatics of the input device—again, specifically, the way its dimensions compose.

Early research on input devices ignored these considerations; more recent work has begun to address them, but still lacks a theoretical framework that can support formal reasoning about multidimensional input device selection.

We therefore contribute:

- An answer to the original question *What is a three-dimensional tracker good for?* That is, when you want to vary three integrally-perceived attributes simultaneously. We found that the Polhemus is neither uniformly better nor worse than the mouse-based approach, but that it depends on the perceptual space of the task.
- An extension to the theory of perception of structure from fixed stimuli to interactive manipulation.
- The demonstration of an approach to answering questions like *What is a three-dimensional tracker good for?* based on extending a perceptual theory rather than on *ad hoc* testing or expert judgment. We derived our answer to the original question from the theory and then verified it experimentally.

ACKNOWLEDGMENTS

We thank Jim Ballas and Astrid Schmidt-Nielsen for help and advice, particularly in experimental design and data analysis; Jeff Brown, Robert Carter, Connie Heitmeyer, Dan McFarlane, Preston Mullen, and Stan Wilson, for all kinds of help with this research; and our NRL colleagues who took time from their own work to serve as experimental subjects. This work was sponsored by the Office of Naval Research.

REFERENCES

1. L.B. Achille, "Considerations in the Design and Development of a Human Computer Interaction Laboratory," NRL Report 9279, Naval Research Laboratory, Washington, D.C. (1990).

2. T.W. Bleser and J.L. Sibert, "Toto: A Tool for Selecting Interaction Techniques," *Proc. ACM UIST'90 Symposium on User Interface Software and Technology* pp. 135-142, Addison-Wesley/ACM Press, Snowbird, Utah (1990).

3. T.W. Bleser, "An Input Device Model of Interactive Systems Design," Doctoral Dissertation, The George Washington University (May 1991).

4. W. Buxton, "Lexical and Pragmatic Considerations of Input Structures," *Computer Graphics* **17**(1) pp. 31-37 (1983).

5. S.K. Card, T.P. Moran, and A. Newell, *The Psychology of Human-Computer Interaction*, Lawrence Erlbaum, Hillsdale, N.J. (1983).

6. S.K. Card, J.D. Mackinlay, and G.G. Robertson, "The Design Space of Input Devices," *Proc. ACM CHI'90 Human Factors in Computing Systems Conference* pp. 117-124, Addison-Wesley/ACM Press (1990).

7. J.D. Foley, V.L. Wallace, and P. Chan, "The Human Factors of Computer Graphics Interaction Techniques," *IEEE Computer Graphics and Applications* **4**(11) pp. 13-48 (1984).

8. W.R. Garner and G.L. Felfoldy, "Integrality of Stimulus Dimensions in Various Types of Information Processing," *Cognitive Psychology* **1** pp. 225-241 (1970).

9. W.R. Garner, *The Processing of Information and Structure*, Lawrence Erlbaum, Potomac, Md. (1974).

10. GSPC, "Status Report of the Graphics Standards Planning Committee," *Computer Graphics* **11** (1977).

11. S. Handel and S. Imai, "The Free Classification of Analyzable and Unanalyzable Stimuli," *Perception and Psychophysics* **12** pp. 108-116 (1972).

12. S. Imai and W.R. Garner, "Structure in Perceptual Classification," *Psychonomic Monograph Supplements* **2**(9) (1968). Whole No. 25.

13. J.D. Mackinlay, S.K. Card, and G.G. Robertson, "A Semantic Analysis of the Design Space of Input Devices," *Human-Computer Interaction* **5** pp. 145-190 (1990).