POINTING ON A COMPUTER DISPLAY

by

Evan D. Graham
B.Math., Waterloo, 1984
M.A.Sc., Waterloo, 1986

THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
in the School of Kinesiology

SIMON FRASER UNIVERSITY
June, 1996

APPROVAL

NAME: Evan D. Graham

DEGREE: Doctor of Philosophy

TITLE OF THESIS: POINTING ON A COMPUTER DISPLAY

EXAMINING COMMITTEE:

Chair: Dr. Daniel J. Weeks

Dr. Christine MacKenzie
Senior Supervisor
Professor, School of Kinesiology

Dr. Thomas Calvert
Professor & Director
Centre for Systems Science

Dr. John Dickinson
Professor, School of Kinesiology

Dr. Kellogg S. Booth
Professor, Computer Science U.B.C.

Dr. John Dill
Internal Examiner
Professor, School of Engineering

Dr. Colin Ware
External Examiner
Professor, Computer Science
University of New Brunswick

Date Approved: 3 June 1996

Abstract

Four experiments were performed to study planning and control of pointing movements in human-computer interaction (HCI). We examined how theories and models derived from motor control research on natural movements apply to HCI, in order to understand why mediated pointing is difficult, and to suggest strategies for improving interfaces. An OPTOTRAK sensor and computer system mapped 3-D hand movements to a 2-D display with less than 25 ms lag. Subjects viewed targets on a graphics display and pointed to locations on the table top. The index finger controlled the position and orientation of a corresponding pointer on the display to reach targets of different widths (W) at a variety of distances (A). For each experiment, both finger (hand space), and pointer (display space) movement trajectories were analysed to examine kinematic features related to movement planning and control.

Experiment 1 demonstrated that planning and control are based in hand space, regardless of the size of A and W on the display. Experiment 2 revealed that asymmetric effects of A and W on kinematics and movement time are a feature of the pointing task, and not a function of the spatial relationship of hand space and display space. Experiment 3 showed evidence for fast corrective processes when target position was perturbed at movement c ·set, even when gaze was directed away from the target location for the hand. Experiment 4 revealed that even when gain is perturbed during a pointing movement, adaptation is achieved through adjusting kinematic parameters of the same underlying plan.

Results indicate that modelling and analysis of performance in mediated pointing must be based on hand actions required for a task, regardless of the size of the display. Pointing performance will be improved by reducing the scale of hand movements; thus future effort should be focused on improving the sensing resolution of input devices for manual interaction. Emerging from this research is a refined two-part model for predicting movement time, providing a more detailed view than Fitts' law. In addition, the model's parameter space predicts performance changes when hand movements are scaled, and characterises tasks and devices in terms of difficulties in the two movement phases. The experiments provide a model of performance when pointing with an unencumbered hand, a benchmark against which pointing devices can be compared. The measurement and analysis techniques, and the identification of separable direction, amplitude, and width effects may prove useful for predicting target location while a movement is ongoing, and can be applied to the design of an intelligent interface, employing precomputation to respond to object selection.

Table of Contents

Abstract iii
List of Figures vii
List of Tables x
List of Equations xi

Introduction 1
  Human-Computer Interaction Studies of Pointing 4
  Theories and Models of Human Pointing 5
    Movement Amplitude and Target Width Effects 6
    Fast Visually-Based Corrections 8
  HCI Setup versus Interacting with the Physical World 9
    Spatial Relationship of Hand and Display Space 10
    Control-Display Gain 10
    Non-Linear Mapping 11
  Overview of Experiments 12

Experiment 1: Control-Display Gain 15
  Method 17
    Subjects 17
    Apparatus 18
    Procedure 20
    Data Analysis 24
  Results 27
    Display Space Analysis 28
    Hand Space Analysis 37
    Modelling Movement Times with Fitts' Law 44
  Discussion 48
    Theory 48
    Application 49

Experiment 2: Physical Versus Virtual Pointing 52
  Method 52
    Subjects 52
    Apparatus 53
    Procedure 53
    Data Analysis 55
  Results 55
  Discussion 64

Experiment 3: Response to Perturbation of Target Direction 68
Method 71
Subjects 71
Apparatus 71
Procedure 73
Data Analysis 76
Results 77
Target Direction 78
Response to Perturbation 82
Discussion 96
Experiment 4: Non-Linear Control-Display Mappings 102
Method 103
Subjects 103
Apparatus 103
Procedure 106
Data Analysis 106
Results 107
Control-Display Mapping Effects 107
Amplitude Effects 110
Width Effects 114
Fitting Movement Times with the Two-Part Model 114
Discussion 118
General Discussion 120
Overview 120
Asymmetry of Amplitude and Width Effects 122
Movement Phases 126
Visually-Based Corrections 127
Two-part Model 134
Predicting Target Location 140
Conclusion 142
Bibliography 143

List of Figures

Figure — Page

Experiment 1:

1. Layout of the experimental setup 19
2. Placement of infrared markers (IREDS) on the hand and arm 21
3. Movement amplitude and target width combinations tested in each gain condition, and subsets of amplitudes and widths used in display space and hand space analysis 23
4. Pointing movement kinematics 26
5. The interaction between gain and movement amplitude for movement time on the display 30
6. Timing of the initial acceleration peak for the cursor showed a gain by amplitude interaction 31
7. The interaction between gain and movement amplitude on cursor peak velocity 32
8. Proportion of time in deceleration for the cursor showed a gain by amplitude interaction 34
9. The timing of the initial velocity peak for the hand demonstrated a small gain by amplitude interaction 39
10. The proportion of time after peak deceleration for the hand showed an amplitude by gain interaction 40
11. Graphical overview showing main effects from display space analysis and hand space analysis 43
12. Movement time data for all amplitude-width combinations in the gain = 1 condition 45
13. Movement time data from Figure 11 replotted to illustrate the relative magnitudes of amplitude and width effects 47

Experiment 2:

14. Layout of the experimental setup 54
15. Mean movement times as a function of target width for virtual and physical displays, showing a display by target width interaction 56

16. Proportion of time spent in deceleration to targets of different widths for virtual and physical display, showing a display by target width interaction 59

17. Illustration of the two-part models from Equations 4.1 and 4.2 63

18. Parameter space for the model of Equation 3.3, showing the data points for virtual and physical pointing 66

Experiment 3:

19. Computer displays used in the experiment, the standard HCI setup and superimposed display 72

20. Placement of the infrared (IRED) markers on the index finger. Positions of the IREDs over the work space are mapped to screen coordinates 74

21. Positions and size of targets displayed in the experiment 75

22. In non-perturbed trials the target was presented in a fixed position for the duration of the trial. In perturbed trials the target was initially presented in the centre position, then shifted to the left or right 81

23. The time of the initial acceleration peak showed a significant interaction between display and direction for non-perturbed trials 84

24. The time of the initial velocity peak showed a significant interaction between display and direction for non-perturbed trials 85

25. The magnitude of the initial velocity peak showed a significant interaction between display and direction for non-perturbed trials 86

26. Movement times revealed an interaction of perturbation and trial block 87

27. The interaction of display and trial block for nonperturbed and perturbed trials to the centre target 88

28. Summary of directional measures for nonperturbed and perturbed trials 94

29. Movement paths in the x-y plane for representative trials in the HCI setup 95

Experiment 4:

30. Gain as a function of hand speed in the plane of the table for the three control-display mappings used in the experiment 105

31. Interactions between control-display mapping and movement amplitude as evidenced by the timing of the initial velocity peaks for the cursor and the hand 111

32. The relationship of the magnitude and timing of the first velocity peak for different control-display mappings 112

33. Parameter-space description for the two-part model of Equation 3.3 relating changes in movement time to changes in amplitude and width 117

General Discussion:

34. Predicted movement times from the model for physical pointing in Experiment 2. In this case, movement times will actually decrease as ID increases 124

35. Base trajectory representation of velocity profiles introduced in Experiment 4 128

36. Experiment 1: schematic view of amplitude and width effects on hand velocity profiles 129

37. Experiment 2: schematic view of amplitude and width effects on hand velocity profiles from both virtual and physical pointing 130

38. Experiment 4: schematic view of amplitude and width effects on hand velocity profiles in constant gain, discrete gain change, and continuous gain change conditions 131

39. Points in the parameter space for models from Experiments 1, 2, 4 and MacKenzie et al. (1987) 136

40. Points in the parameter space of the two-part model for individual subjects in Experiment 1 138

List of Tables

Table | Page

Experiment 1:

1. Display space analysis: Summary of main effects of gain on two-dimensional cursor kinematics 29

2. Display space analysis: Summary of main effects of movement amplitude on two-dimensional cursor kinematics 35

3. Display space analysis: Summary of main effects of target width on two-dimensional cursor kinematics 36

4. Hand space analysis: Summary of main effects of gain on three-dimensional hand kinematics 38

5. Hand space analysis: Summary of main effects of movement amplitude on three-dimensional hand kinematics 38

6. Hand space analysis: Summary of main effects of target width on three-dimensional hand kinematics 42

Experiment 2:

7. Summary of the main effects of display condition. 57

8. Summary of the main effects of movement amplitude 60

9. Summary of the main effects of target width 61

Experiment 3:

10. Summary of effects of trial block on nonperturbed trials to left, centre, and right targets 79

11. Summary of effects of display on nonperturbed trials to the left, centre, and right targets 80

12. Summary of effects of direction on nonperturbed trials to the left, centre, and right targets 83

13. Summary of main effects of trial block on nonperturbed and perturbed trials to the centre target 89

14. Summary of main effects of display on nonperturbed and perturbed trials to the centre target 91

15. Summary of main effects of perturbations: nonperturbed trials to the centre target, perturbed trials to the right, and perturbed trials to the left 92

Experiment 4:

16. Main effects for control-display mapping on cursor and hand kinematics. 108

17. Main effects for movement amplitude (A) on cursor and hand kinematics. 113

18. Main effects for target width (W) on cursor and hand kinematics. 115

List of Equations

Equation	Page
1	6
2	6
3	17
4	44
5	46
6	46
7	48
8	58
9	58,122
10	62
11	62
12	104
13	116
14	116
15	116
16	116
17	116
18	116
19	123
20	134

Introduction

As with all new technologies, the full advantage of modern information systems is felt only by those who can devote time and effort towards mastery. Some of the difficulties of acquiring these abilities can be reduced by careful study and redesign of the interface through which humans interact with these machines. In the last decade, the field of human-computer interaction (HCI) has emerged as an important area for multidisciplinary study, with the fields of computer science, information systems, psychology, and human factors forming the foundation (Dillon, 1995). At present, researchers and designers are concerned with multiple simultaneous goals (Wickens, 1992; Schneiderman, 1992; Carroll, 1991; Laurel, 1990): to increase the speed, storage, communications bandwidth of hardware; to develop systems which effectively support work and other activities; to design interfaces which match and complement human capabilities; and to increase utility and usability of computers so that a larger population may enjoy the advantages of information technology. HCI can be approached from the perspective of computer science, focusing on the technology. People, however, are very different than computer systems, and while it is relatively easy to redesign hardware and software, basic human capabilities are not so easy to change. By drawing on other disciplines in addition to computer science, HCI can also be approached from the perspective of human performance, emphasising the personal experience of the user. This approach, often referred to as "user-centred" (Norman & Draper, 1986), originally inspired the author and led to to the research which is the topic of this thesis.

The traditional form of human-computer interaction is a conversational metaphor: the computer is an agent which responds to requests (typed, spoken, or selected) to perform work. This approach is being extended to the next generation of systems through ideas such as "personal assistants" (Reichman, 1986), or a computer that simulates social interaction, using expressions and body language to participate in group activities with other humans (Takeuchi & Naito, 1995). For many computer-aided tasks, systems which embody the concept of an agent remain the focus of our attention, rather than being a tool through which we perform the work. In addition to task expertise, users also need to develop the skills to direct the computer to perform the task. Traditionally, this additional level of difficulty has been addressed by training, trading off effort saved in automating the task with effort expended in learning to operate the interface.

Computer "pointing", using an input device to indicate a location on a graphics monitor, is an elemental gesture in many forms of human-computer interaction. With the increasing popularity of mouse-based interfaces, users spend less time typing on a keyboard, and more time communicating with computer systems using pointing movements. Augmented reality and virtual reality systems offer additional potential in this direction by supporting more complex interactions with three-dimensional displays (Welner, 1993). Other developments, such as "ubiquitous computing" (Weiser, 1993), envision distributing computational power amongst a large number of artifacts, small objects with which people interact in a natural manner as part of their everyday activities.

All of these developments suggest that computer interfaces will increasingly rely on human prehension to support interaction. The term "direct manipulation" was coined by Schneiderman (1982) to describe interfaces based on a strong metaphor for physical interaction with objects in the world. Direct manipulation[1] is designed to give users the illusion that they operate on a system (representing the task) without the impression of an intervening agent, by using simple actions. Examples of direct manipulation interfaces (Sutherland, 1966; Brennan, 1990; Schneiderman, 1992) take advantage of some human spatial, cognitive, visuomotor, and prehensile abilities in order to reduce information processing loads (Kay, 1990), and have shown high user acceptance. The advantage of such a strong metaphor over the traditional conversational approach is that users can transfer their knowledge of how to interact with the physical world, and use their natural prehensile abilities to operate the interface. A disadvantage is that poor design violates the users' assumptions about the the effects of their actions on the representation of a task, leading to errors and difficulties.

A system which represents a user's movements, such as pointing, in a direct way is

---

[1] Direct manipulation should not be confused with windows-icons-menus-pointing, or WIMP (Dix, 1993), which usually indicates a superficial technique applied to existing command-based programs to make them easier to use, while still retaining the same basic command structure.

"visually coupled", and requires a high level of performance to sustain the illusion. The performance requirements for spatial and temporal accuracy must be investigated and understood in terms of human perception and behaviour (Kalawsky, 1993). Current technology has been developed without much regard for these requirements, and its limitations give rise to difficulties ranging from mild annoyance to nauseating simulator sickness (Kennedy et al., 1987). Hand movements with a pointing device which senses forces, motion, or position of a handle or puck to control the movement of a cursor on the display screen have been in general use for several decades. A number of studies (for a summary see S. MacKenzie, 1992) have suggested that theories and models of human performance from motor control research on natural pointing movements may be brought to bear on the study of more abstract interactions with a two-or three-dimensional graphical interface mediated by a pointing device.

The author has been involved with pointing devices and interactive graphics software as part of his work in educational computing for the last decade, and experience with colleagues and thousands of university students suggests that it is difficult to control an interface and the simple gestures which underlie the interaction, even though manipulation of objects in the physical world is performed with ease using the same actions. Comparison of the results from a significant body of both human factors and motor control research supports this observation: movement times are generally longer, and error rates higher for HCI studies than for motor control studies where a subject points with a finger or stylus. Why should this apparently simple task be more difficult in the HCI context, where objects and actions represented on a graphics display are mediated by a pointing device and computer system? Some preliminary investigations (Graham, MacKenzie, Calvert & Crawford, 1993) suggested that such mediated pointing is not simply a slower version of the natural task, and that the fine structure of positioning movements is markedly different for pointing on a computer display than for natural pointing at objects in the physical world.

The purpose of this thesis is to investigate how the limitations and constraints of two- and three-dimensional virtual environments affect human movement planning and control. Even with future developments of immersive displays and sensing hardware, these constraints will persist. Visual information in a virtual environment is necessarily an abstraction: first, there are practical limitations on the complexity of a world model (computational), and a visual representation (display resolution) which can be simulated in real-time; second, there are limitations on the spatial accuracy and lag in sensing positions of objects and body segments, in particular the hands and arms, in order to be represented as part of the virtual environment. Recent advances in methodology have enabled researchers to address questions about movement planning and control by recording and analysing the trajectories of moving limbs, using an opto-electronic technique which leaves the subject relatively unencumbered by mechanical linkages. For the experiments described in this thesis, an optical tracking system was used to simulate input devices for human-computer interaction which sense position or motion, but with much finer temporal and spatial resolution than current commercially available systems.

To examine the question of why mediated pointing is difficult, detailed analyses of movement trajectories were performed in order to reveal systematic effects of the constraints on human performance, and to suggest strategies for design of human-computer interfaces which more effectively support manual interaction. The research concentrated on how, and when to apply models and descriptions of behaviour from a large body of research on human and perceptual-motor abilities to mediated pointing on a computer display. The author hopes to be able to extend these results and methodologies to understand more complex goal-directed movements such as reaching, grasping, and manipulating in virtual environments.

Human-Computer Interaction Studies of Pointing

Empirical studies of human performance using input devices to position a cursor on a display screen spans at least two decades. One approach has been to compare two or more pointing devices or tasks (Albert, 1982; Card et al., 1978; Ewing et al., 1986; Goodwin, 1975; Haller et al., 1984; Jellinek & Card, 1990; Karat et al., 1984; S. MacKenzie et al., 1991; Sears et al., 1991; Walker, 1990; Zhai, 1992) to determine which is the best combination. The tasks usually involved selecting a target, dragging a cursor over an area, tracking a target, or tracing a path on the screen. The devices studied range from step keys to a light pen, touch screen, touch pad, graphics pad, joystick, and mouse. One general conclusion from this research is that a mouse is better than most other devices for selecting a target. Another approach has attempted to quantify human performance with novel prototypes (Chung, 1992; Gobel et al., 1995; Hill et al., 1992; Hirota et al., 1993; Kabbash et al., 1995; Lippman, 1981; Maggioni, 1993; Murakami, 1994; Rosenberg, 1993; Stammers & Bird, 1980) or applications of existing input devices (Flach et al., 1987; Arthur et al., 1993) to create a novel form of interaction.

The bulk of this earlier research has concentrated on simple behavioural measures: reaction time, movement time, and spatial accuracy. In addition, the measures are usually derived from the motion of a cursor, not from the human movements elicited to perform the task under study. Although this approach is useful to evaluate comparative levels of performance under different conditions, it does not provide much insight into the planning and control of human movements in a HCI context. Because the research has concentrated on evaluating existing devices and systems, there has been little attempt to systematically manipulate design features to infer their effects on human behaviour. Some exceptions are recent studies by Walker et al. (1993), who looked at the microstructure of positioning movements, and Ware & Balakrishnan (1994) and Bryson (1995) who systematically manipulated lag and display update rate.

Theories and Models of Human Pointing

The action of pointing to locations with a stylus or finger has been studied since Woodworth (1899) first proposed that human pointing movements can been understood in terms of two movement phases: an initial planned impulse which covers most of the distance, followed by a second phase of deceleration to the target under current control. Important kinematic features of the movement phases can be represented schematically by a bell-shaped velocity profile. The initial movement phase is characterised by the timing and magnitude of the velocity peak which arises from the planned impulse. Woodworth also demonstrated that when vision was used to home in on the target, subjects were more accurate, but approached the end point of the movement more slowly. Following on these ideas, Craik (1947) reintroduced the question of how vision contributes to movement control in aiming, pointing and tracking tasks. He proposed that corrections to the movement trajectory occurred at discrete intervals as the error was sampled through vision. Since then, the time required to correct errors in movement trajectories on the basis of visual information, and the nature of the mechanisms of visuomotor control have been the subject of continued investigation.

Movement Amplitude and Target Width Effects

According to Fitts (1954), the average time for discrete or repetitive pointing movements results from a tradeoff between speed required for distance covering, and controlling accuracy of the final position. Formulated on the basis of information theory, Fitts' law describes movement time (MT), on average:

$$MT = a + b\,ID \qquad (1)$$

$$ID = \log_2(2A/W) \qquad (2)$$

as a linear function of the index of difficulty (ID). ID is determined by the ratio of distance to the target (A) and target width (W). The coefficient b (seconds per bit) in Fitts' law is interpreted as the multiplicative inverse of the index of performance (bits per second), an information theoretic measure of the bandwidth of the channel for movement control. Support for this model has been demonstrated by a number of studies on discrete and repetitive tapping and pointing with physical targets, using an implement such as a stylus (e.g. Fitts & Peterson, 1964; Crossman & Goodeve, 1983; Welford, 1968). Fitts' law has also been suggested as a model for movement time when a cursor on a computer display is controlled by a two-dimensional pointing device such as a mouse or trackball; a number of studies (Boritz, Booth & Cowan, 1991; Jellinek & Card, 1990; S. MacKenzie 1989, 1992; MacKenzie & Buxton 1992; MacKenzie, Sellen & Buxton, 1991; MacKenzie & Ware, 1993; Trankle & Deutschmann, 1991; Walker, Meyer, & Smelcer, 1993) show results to support the use of the model in HCI.

The information theoretic approach of Fitts resulted in several models of how vision and kinesthesis combine to control movement trajectories. Crossman and Goodeve (1983) proposed that after the initial impulse, visually determined error between limb and target position resulted in a series of deterministic iterative corrections to guide the effector to the goal position. Schmidt et al. (1979) proposed an impulse variability model, in which error in the initial phase of the movement was in proportion to the magnitude of the impulse, resulting in the speed-accuracy tradeoffs implied by Fitts' law. Carlton (1980) observed that vision during the initial phase of a pointing movement contributed little to accuracy, and supported the view that a single correction was made at the time of peak velocity. Meyer et al. (1980) incorporated both iterative corrections and impulse variability in their stochastic optimised submovements (SOS) model, where a series of overlapping impulses combine to form a relatively smooth movement trajectory. Although these models all predict speed-accuracy tradeoffs described by Fitts, they each imply a different fine structure for movement trajectories. Impulse variability predicts smooth velocity profiles. The iterative corrections model predicts multiple peaks in the velocity profile, the average number related to the log of the distance to the goal position. The SOS model also predicts multiple peaks, but with trajectories which may appear relatively smooth depending on how the submovements overlap.

In natural pointing, Soechting (1984) and Langolf et al. (1976) examined pointing movement velocity profiles, and showed that the shape of the profiles was skewed for smaller, more difficult targets. Mackenzie et al. (1987) demonstrated that this effect was systematic for three-dimensional pointing movements, suggesting that movement amplitude (A) and target size (W) gave rise to separable features of the movement trajectory. The effect of increasing A (moving over greater distances) was to increase the overall size of the velocity profile, in terms of peak speed and time. The overall shape of the velocity profile was changed by decreasing W (using smaller targets), and became skewed so that a greater proportion of time was spent in the deceleration phase. The skew of the velocity profile is termed the "precision effect" (MacKenzie et al., 1987), and has been shown to be a feature of both pointing and grasping movements where error in the final position must be reduced due to task constraints. MacKenzie et al. (1987) also suggested that these separable effects must combine so as to cancel their effects for similar IDs when A and W are scaled in proportion. Velocity profiles for similar IDs but different distances will have different shapes and sizes, but will take the same total time, because the increase in overall size due to increasing A is cancelled by the shortening of the deceleration phase due to increasing

W. Their results support a view that speed along the path of the movement can be represented by a base trajectory profile, with overall size scaled by A, and shape determined by W.

Fast Visually-based Corrections

For humans, visually guided reaching is rapid and precise when directed at targets in the work space, and the hand usually moves from one stationary point to another tracing out a relatively straight path (Flash & Hogan, 1985). These fast pointing movements, with a finger or stylus, generally exhibit smooth velocity profiles which are somewhat skewed due to the precision effect. In contrast, fine movement features of velocity profiles using a mouse to control a cursor on a computer display reveal that a target is often reached in one or more submovements spaced at a period of several hundred milliseconds (Walker, Meyer & Smelcer, 1993; Graham, MacKenzie, Calvert, & Crawford, 1993). Iterative corrections of this type appear to include a simple reaction time with a latency of about 200 ms, and can be understood to include the time to identify an error in the position of a cursor in relation to a visual target, plus the time required to organise a corrective motor response. For goal-directed movements in the work space, peak speed is usually reached in under 300 ms; thus corrective responses based on visual comparison to generate an error signal will be observed in the deceleration phase of a movement.

More recent studies, in which subjects aimed or pointed without vision of the hand and arm, have used detailed analyses of movement kinematics to suggest involvement of a more direct mechanism which continuously corrects the trajectory of a moving limb (Goodale, Pelisson, Prablanc, 1986; Pelisson, Prablanc, Goodale, Jeannerod, 1986; Zelaznik, Hawkins & Kisselburgh, 1983). This fast sensorimotor control loop relates visual target position directly to the felt position of the hand, resulting in very short latency (somewhat less than 100 ms) corrections to the movement trajectory. This has been shown to occur when the position of a target is perturbed at the time the movement is initiated. In some cases, smooth corrections to the direction of movement are observed. In other cases, reorganisation of the movement is indicated by changes in the timing of characteristic peaks in acceleration and velocity that occur early in the movement.

HCI Setup versus Interacting with the Physical World

There are a number of differences between the standard configuration of input device and visual display that are typically used for interaction with a graphical interface, compared to natural interaction with the physical world. At present, a typical workstation configuration requires the user's gaze to be fixed on a CRT display screen in front of their body, while the right hand moves a pointing device on the table top in the right half of the work space. We will refer to this as the "standard HCI setup". For pointing to locations on the screen, this setup represents a simple virtual environment where hand actions control a visual representation on a display. More technically sophisticated environments are termed "augmented" or "virtual" (Wellner, 1993). In augmented reality, optical techniques superimpose a computer display on the user's view of physical objects in the work space. Virtual reality provides the user with an immersive display, completely replacing his or her view of the world with a computer-generated image.

In terms of human perceptual-motor processes, there are fundamental differences between virtual environments and natural interaction with physical objects in the work space. Some of these differences can be understood in terms of the relationship of the work space for the hand (hand space), and the region in which the representation of actions and objects are displayed (display space). In HCI, compared to natural interaction with physical objects in the workspace:

a) display space need not be in the same location as hand space;
b) the size of the representation of actions and objects in display space may be arbitrarily varied, even though identical hand actions are used for interaction; and
c) the mapping between motions of the hand to motions of its displayed representation is arbitrary.

In addition to the limited computational model, and constraints on display resolution and sensing mentioned above, any one of these additional factors could account for the difficulties with mediated pointing due to the disruption or modification of natural perceptual-motor processes when pointing on a computer display. The issues of spatial relationship between hand and display space, the relative sizes of hand and display space, and the mapping of hand movements to movement on the display are discussed in the following sections.

Spatial Relationship of Hand and Display Space

Normally, when pointing at physical targets viewed directly in the work space, hand space and display space are exactly superimposed. In this situation, the act of fixing gaze on a target (head and eye movements) provides information in a body-centred coordinate system about target location; this information contributes to accurate positioning of the hand in pointing and grasping (Carnahan & Marteniuk, 1991; Gentilucci et al., 1994). In the standard HCI setup, hand space (where the pointing device is operated) is on a horizontal surface in front of and to the right (or left) of a user, whereas display space is usually directly in front of the body in a near vertical plane. Gaze is directed at a location on the display, and provides no direct information about the goal location for the hand to control a cursor in reaching the target on the display. Assuming that hand and cursor movements are represented at the same scale, predicting how hand movements on the table will affect cursor position requires that the user learn or adapt to a particular spatial transformation (translation and rotation) from hand space to display space. There is some evidence that humans can adapt to simple displacements of the visual array (e.g. Elliot & Roy, 1981), and can retain and recall a particular spatial transformation appropriate for the context in computer pointing (Cunningham & Welch, 1994). But if this adaptation requires additional information processing, it might disrupt the fast corrective processes used to control natural movements, and account for some of the increased difficulty in mediated pointing.

Control-Display Gain

If hand space and display space are superimposed, the relative scale of object sizes and distances is necessarily identical for the eye and hand. If hand and display space are not superimposed, their relative scale can be altered. When the input and output is one-dimensional (as in much of the human factors research on controls and displays), this relationship is quantified as a control-display ratio, or control-display gain (the inverse of control-display ratio). Control-display gain is the preferred description, as it is more directly related to the user's experience (Buck, 1980).

In two or three dimensions, the same concept can be used to describe the relationship between the scale of hand movements and their representation on a display. With a linear position-to-position mapping, for example, a control-display gain of 2 means that each 1 mm of hand movement results in 2 mm of cursor movement in a corresponding direction. For movement planning, applying this particular input-output mapping in mediated pointing results in a visual presentation of amplitude (A) twice as big as the actual A required for the hand to reach the goal position. Similarly, the accuracy constraint for the end point (target width, W) is also visually represented as larger than the actual constraint for the hand. One view of coordination and control is that the goal of the movement plan is to control the trajectory of the cursor on the display. In this case, the separable effects of A and W would remain invariant in the cursor kinematics regardless of control-display gain. An opposing view is that visual information is translated into body or hand-centred coordinates, in which case the separable effects would be related to the size of A and W for the hand, regardless of control-display gain. A somewhat unexplored view is that, with control-display gains other than unity, there is some confusion in movement planning and control, which could account for additional difficulty in mediated pointing. For example, if planning were based in hand space and deceleration to the target based on the display, then A for the hand and W on the display would give rise to characteristic features of the velocity profile.

Non-Linear Input Mapping

As discussed above, in the standard HCI setup, not only are hand and display space in different locations, but the relative sizes of objects and distances is arbitrary. In addition, hand and arm movements may also bear little direct relation to movements of the cursor on the display. Interaction of the hand with an input device can be sensed in a variety of ways (position, motion, force), and use a variety of strategies to convert the sensed value to control the behaviour of a cursor, changing its position, velocity, or acceleration over time (Baecker, Grudin, Buxton, & Greenberg, 1995, pp. 469-479; Card, MacKinlay, & Robertson, 1991). Even when the display is constrained to one or two dimensions, the hand still moves in a three dimensional working space.

For simulating prehension in a direct manipulation interface, a simple linear position-to-position mapping seems most appropriate. Such a mapping has the property that a given set of movements in any order from a fixed starting point will always leave the cursor in the same end position, and moving faster or slower over the same path produces the same end result.

Although users of computer pointing devices perceive that commercial systems use a position-to-position mapping, this is not the case. The mouse (or other pointing device) provided with most workstations uses a higher control-display gain for fast movements, and a lower gain for fine positioning movements, resulting in a non-linear mapping (e.g. Rose, 1985; Scheifler & Krikorian, 1988). Empirical studies have failed to prove any advantage to such a mapping other than reducing the space required in which to operate the mouse (Jellinek & Card, 1990; Flach et al., 1987). The fine features of the movements in operating these devices have not been examined from a motor control perspective. In an experimental context, changing the gain during the course of a positioning movement amounts to a perturbation, and may provide a method for determining which features of movement kinematics are visually based, and which are hand-based. By studying how subjects adapt to non-linear mappings, features of the cursor kinematics which remain invariant when gain is perturbed provide evidence for visual control. Features of hand movements which remain invariant provide evidence for planning and control based on kinesthesis.

Overview of Experiments

In order to make a more direct connection between previous research on pointing at physical targets in the work space and pointing on a computer display, we chose a natural movement for the task to be studied. In all experiments, a subject rested their hand on the table top, with the index finger extended in a start position. To reach a target, the subject lifted the hand and pointed with the index finger, touching the target location. This pointing movement was kept constant across all experiments, although direction, amplitude, and target width varied. Visual information was manipulated by changing the location of targets, by representing targets and hand movements at different scales, and by perturbing the target position or the dynamic mapping of the hand movement. The finger was displayed as a red pointer on the computer screen. Subjects did not interact with an input device, but used the free hand to control the position and orientation of the pointer on the display. This system was designed to support experiments in which it is important to avoid artifacts introduced by the size, configuration or weight of a pointing device; as well as the lags and lack of spatial and temporal accuracy inherent in many sensing technologies. The spatial resolution was at least .2 mm for the hand, and the image of the pointer was updated at 60 Hz. so that its position lagged the position of the index finger by approximately 25 ms (Graham, 1996, in preparation).

The system simultaneously recorded two-dimensional data describing the position of the pointer on the display, and three-dimensional data describing the position of markers placed on the hand. Data for the hand and pointer were subjected to similar analysis techniques to determine characteristic features of their kinematics. This approach allowed us to compare results for the hand directly with previous studies on pointing to physical targets. Also, the results for the pointer could be compared with previous research on using input devices to position a cursor on a computer screen. In addition, by comparing the results for hand and pointer, we were able to make inferences about how visual information from the display and kinesthesis from the arm and hand combine in order to plan and control pointing movements.

Each of the following four experiments addresses a particular aspect of the general question of how mediated pointing differs from direct interaction with physical targets in the work space, and why it is more difficult than the natural version of the task. Experiment 1 examined how the scale of display space, relative to the scale of hand space, affects movement planning and control. Experiment 2 contrasted performance with two-dimensional, versus three-dimensional visual information, and also provided a baseline for pointing at physical targets in the workspace for comparison with the mediated version of the task. In Experiment 3, target locations were perturbed during the movement to show how processes responsible for fast corrections to movement trajectories are affected by the relationship of hand and display space. Experiment 4 examined how movement planning and control adapt when gain is changed during a pointing movement.

In the general discussion that follows, results of the four experiments are integrated and the implications are considered in two ways. First we consider the theoretical implications in terms of studying motor control in virtual environments, and second we consider the application of results in terms of general principles for interface design, and empirical techniques for analysis of human-computer interaction in pointing on a computer display. In closing, we consider some issues raised by the results of the four experiments, and propose an approach for further research.

Experiment 1: Control-Display Gain[1]

In human computer interaction (HCI) it is common to view a graphic display in one location while operating a hand-held device in another location to control the position of a cursor on the display. With a typical input device that maps positions to positions (e.g., a graphics tablet) the effects of hand movements on cursor position can be understood in terms of a three-dimensional linear spatial transformation involving translation, rotation, and scaling which relates points on the operating surface for the device to the surface of the display. In this context, one important issue is the relative sizes of the hand space and the display space. In the one-dimensional case, this relationship is characterised by the control-display gain, the ratio between the size of a movement and the size of its effect on the display. A higher gain means that a given hand or device movement will have greater effect on the position of the cursor on the display.

Human factors references such as McCormick (1976) report that there will be an ideal gain setting for any particular task and system, so as to scale the amount of control movement to a comfortable range. The effect of changing control-display gain on movement times has been investigated using different two-dimensional computer pointing devices. Users of a particular system often have a preferred setting (Card, Moran & Newell, 1983) even though there is little evidence that gain has a significant effect on pointing performance (Jellinek & Card, 1990), at least in terms of time to select a target on the display. Trankle and Deutschmann (1991) reported that changing control-display gain for a mouse had no effect on positioning times, but that display size itself was a decisive factor. This result seems confusing, since changing display size with a given pointing device should be equivalent to changing control-display gain.

Any linear mapping of positions in hand space to positions on the display which scales all dimensions equally exhibits a constant control-display gain for all directions of movement. With all dimensions scaled equally, changing the gain does not change the ratio of movement amplitude (A) to target width (W) for a particular A and W combination. As a result, the index of difficulty (ID) is preserved regardless of the choice of control-display

[1] A short paper describing this experiment entitled "Pointing on a Computer Display" was published in the CHI '95 Conference Companion, Human Factors in Computing Systems, Denver, Colorado, May 1995.

gain, and Fitts' law would predict no difference in average movement time for pointing with a given index of difficulty. Most HCI studies have measured A and W on the display (Card, English & Burr, 1978; Epps, 1986; Gillan, Holden & Adam, 1990; S. MacKenzie, 1991; Walker, Meyer & Smelcer, 1993; Boritz, Booth & Cowan, 1991), where the effect of gain is to change the scale of hand movements for the same display representation. Buck (1980) took the opposite view, analysing control device kinematics in a positioning task, where the effect of gain was to change the size of the display representation for the same hand movements. For the task of turning a dial to control the position of a pointer in one dimension, he concluded that the width of the target as represented on the input device, i.e. as an accuracy constraint for hand movements, determined positioning time. None of these studies simultaneously manipulated both A and W to measure performance over a wide range of IDs.

The separable effects of A and W on velocity profiles first suggested by Langolf et al. (1976), and Soechting (1984), and systematically investigated by MacKenzie et al. (1987) were that:

a) the shape of the profile was skewed so that a greater proportion of time was spent in the deceleration phase for smaller targets, and;

b) the size of the movement profile was scaled up for greater distances.

MacKenzie et al. concluded that these separable effects must combine so as to cancel their effects for similar IDs, when movement distance and target width are scaled in proportion. In other words, velocity profiles for similar IDs but different distances will have different shapes and scales, with constant movement time.

In an HCI context using various gain settings, there are really two movement distances to consider: distance of the hand movement (A_h), and distance of the cursor movement on the display (A_d). In addition, there are also two target widths: visual size on the display (W_d), and width as an accuracy constraint for hand movement (W_h). It is not clear which A and W will give rise to the separable effects on velocity profiles when gain is different than one. Since the initial impulse is planned for distance-covering, and the distance is to be covered by movement of the hand, we would predict that movement amplitude in hand space (A_h) is the planning variable to determine the overall size of the velocity profile. Since the deceleration phase is presumably controlled by visual comparison of the relative positions of the cursor and target, we would predict that it is target width on the display (W_d) which would determine the characteristics of the second movement phase. Under these assumptions, an effective index of difficulty would be:

$$ID = \log_2(2A_h/W_d) \qquad (3)$$

Consider the effect of increasing control-display gain when the size of the display is held constant. In this case, the higher gain reduces A_h, but since the size of the display is held constant W_d remains unchanged. By reducing A_h without affecting W_d, the predicted index of difficulty is lowered. Thus, in terms of Equation 3, increasing gain lowers the effective index of difficulty for a positioning task without changing visually determined constraints. Since the performance increments and decrements for higher gains reported in HCI and human factors studies are not dramatic, it is more likely that the separable A and W effects combine in some other manner to affect movement trajectories when control-display gain is varied.

Experiment 1 was designed to address the question of exactly how A_d, A_h, W_d, and W_h affect planning and control of hand movements in computer pointing. In this study, A_d, W_d, and gain were varied systematically by powers of two, allowing two simultaneous analyses of movement trajectories: display space analysis, based on two-dimensional cursor kinematics as done in most previous HCI studies; and hand space analysis, based on three-dimensional kinematics of hand motion, as done in most previous motor control studies.

Method

Subjects

Six volunteers, ranging in age from 19 to 41 years, were paid $10 for participating in a 90 minute session. All had experience with the use of computer systems and pointing devices, and preferred to use their right hand in this context. Subjects had normal or corrected-to-normal vision. As in all experiments in this thesis, subjects provided informed consent for their participation.

Apparatus

The experiment layout is shown in Figure 1 (see caption). The subject was comfortably seated at a table, with seat height adjusted so that the forearm was approximately parallel to the table top when moving the hand over the workspace. The computer display was placed at a viewing distance of 65 cm, the top of the display surface tilted away from the subjects at an angle of 15 degrees from vertical. The working space for the hand was centred at 30 cm to the right of the body midline, with the starting position for the index finger 12 cm from the edge of the table. The display was viewed at an angle of from 15 to 20 degrees below horizontal, depending on the proportions of a subject's upper body.

Movements were recorded with an OPTOTRAK motion analysis system, which measured the three-dimensional position of infrared markers (IREDs) placed on the subject's hand and arm. The markers were placed as shown in Figure 2:

1) on the central aspect of the nail of the second digit (index finger), as distal as possible but not interfering with contact of the finger tip on the table surface;
2) near the medial proximal aspect of the metacarpal-phalangeal joint (MCP);
3) on the skin above the medial aspect of the styloid process of the ulna;
4) on the forearm about five cm distal to the head of the radius, so as to be in view of the OPTOTRAK camera;
5) near the medial head of the deltoid, on the distal aspect.

Data from the OPTOTRAK were sampled and recorded at 60 Hz by a Silicon Graphics Indigo Extreme computer workstation. The workstation also controlled a graphics display presented on an RGB monitor, using 1280 by 1024 pixels over an illuminated area of 350 by 280 mm. Before each experiment session, a reference coordinate frame was recorded by placing three markers on the table top corresponding to the origin, maximum x position, and maximum y position on the computer display. The dimensions of the reference frame, which covered the workspace for the hand, were determined by the control-display gain for the experiment session as explained below. The reference frame was used for subsequent data analysis, as well as to generate in real time an

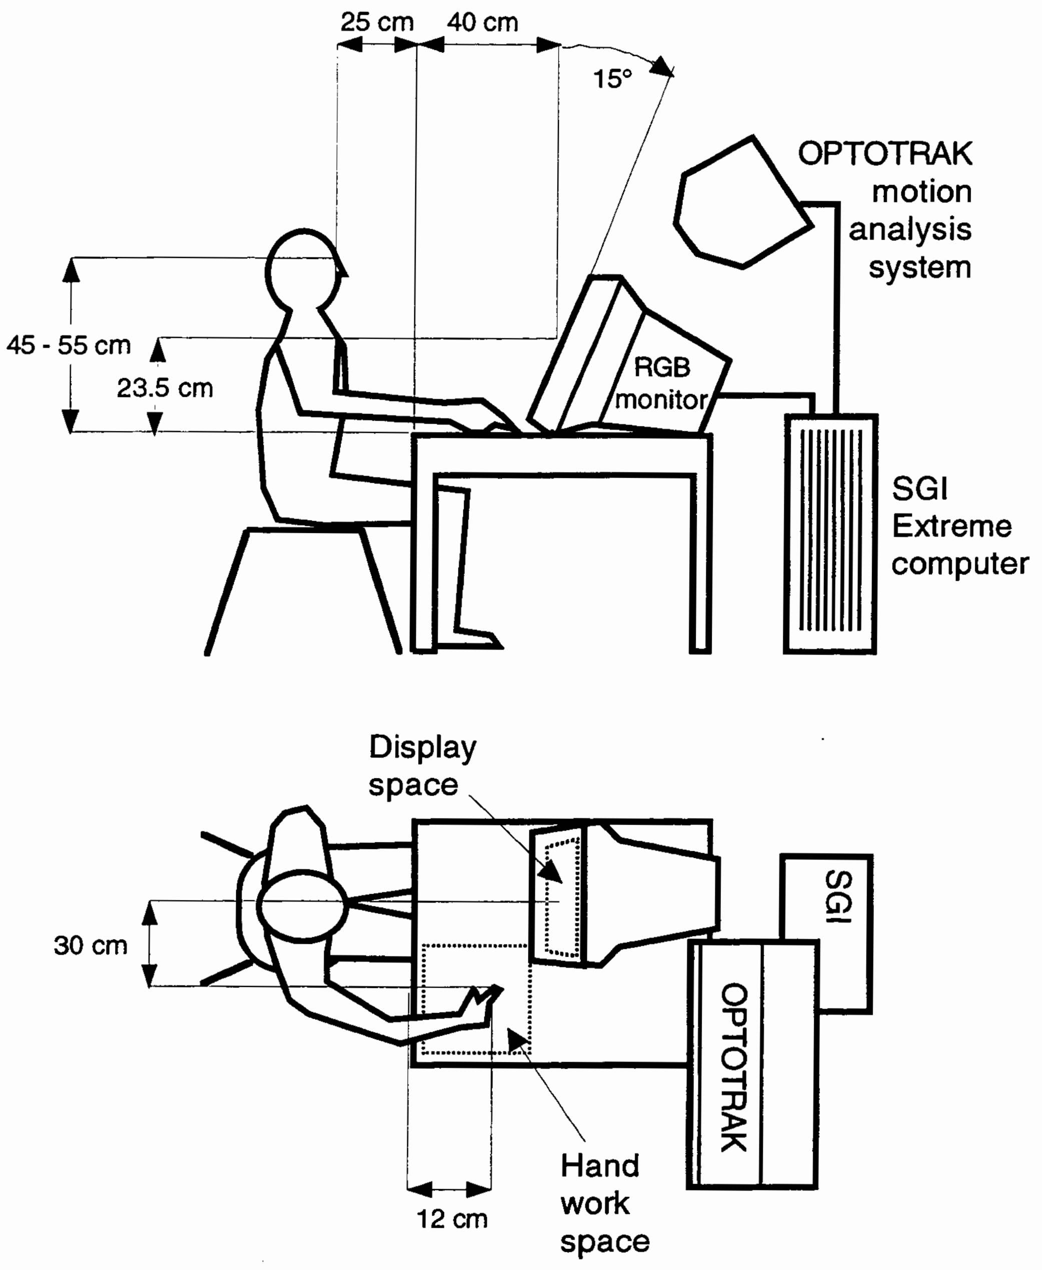

Figure 1. Layout of the experimental setup. Hand movements on the table in front of and to the right of the subject are tracked by the OPTOTRAK motion analysis system. A pointer, controlled by the index finger, and experimenter-defined targets appear on the display in front of the subject. The spatial arrangement of the workspace and display is intended to imitate a standard HCI setup.

orthogonal (orthographic parallel plan view of the workspace) projection of the three-dimensional position of the hand markers above the two-dimensional table surface.

The computer displayed white circular targets on a black background, as well as a 3 mm diameter white circle to define the start position for each trial. A subject moved the hand above the table top to control the position and orientation of a red pointer (35 by 9 mm) on the display. As illustrated in Figure 2, the projected position of the first IRED marker controlled the position of the tip of the pointer which was oriented along a line projected on the table surface through the positions for the first and second markers. Projected coordinates of the virtual finger were also recorded in synchronisation with data from the IRED markers on the hand and arm.

The position of the red pointer on the display was updated using the latest position data from the OPTOTRAK at 60 Hz. The total system lag from position-sensing with the OPTOTRAK to the appearance of the red pointer with its position updated was determined to be approximately 25 ms by two methods. First, timing calculations from the data acquisition and double-buffered graphics display program predict this value. Second, we have done preliminary tests to quantify subjects' perception of the lag by having them move an IRED marker to produce a 10 mm spatial misregistration between the marker and its image on the display (Graham, 1996, in preparation).

Procedure

The computer system continuously updated the display during the entire experiment session. As a result, the subject experienced no interruptions nor any need to readjust to the difficulty of controlling the position of the pointer between trials. Three-dimensional position data for the hand and pointer were recorded from the start to end of each trial, and stored in separate computer files for subsequent analysis. To initiate a trial, the subject viewed the display, held the tip of the pointer on the small start mark, and indicated they were ready for the next trial by saying "ready". The task of holding the pointer still on the start mark was sufficiently difficult to require the subject's gaze to be directed to the start position. After a short delay (from one to two seconds), during which the experimenter determined that the pointer was in the correct position and that the subject's gaze remained

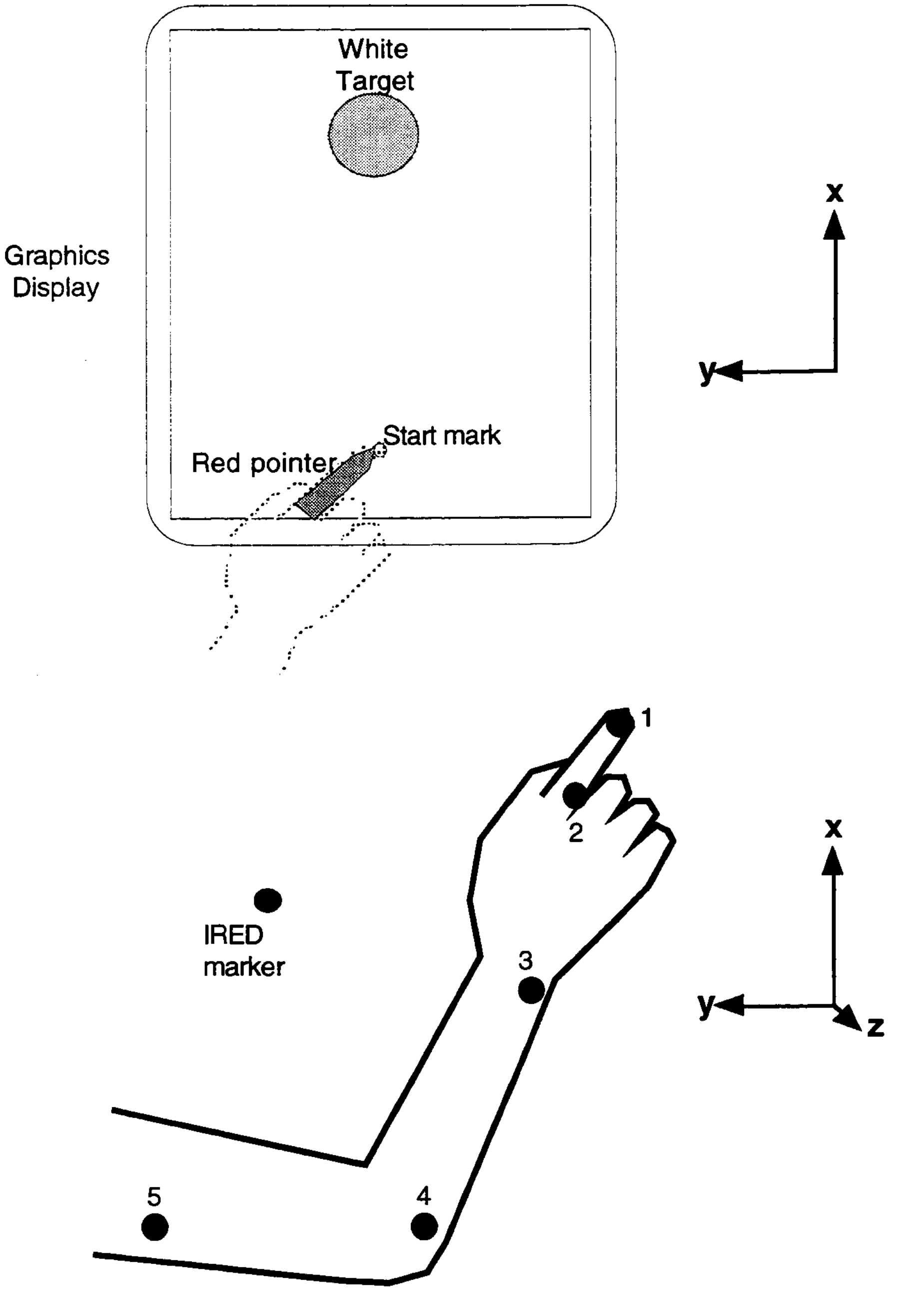

Figure 2. Below, the placement of infrared markers (IREDs) on the hand and arm. The positions of IREDs in the work space are monitored by the OPTOTRAK system in three dimensions. Above, the positions of markers on the hand and index finger are mapped to two-dimensional screen coordinates in real time, and control the position and orientation of a red pointer on the display.

on the start mark, a target appeared. Still viewing the display, the subject lifted the hand and pointed with the index finger to a position on the table top so that the tip of the red pointer stopped anywhere inside the white circular target. The subject was instructed not to rush to initiate the movement, but to move as quickly as possible while still landing in the target, and to make a small correction if necessary. After the tip of the pointer was held still (speed less than 4 mm/s for 600 ms) on the target, data collection for that trial was stopped, the target was erased, and the start mark was again displayed.

In the low control-display gain condition, the dimensions of the working space for the hand were 350 by 280 mm on the table top, which matched exactly the size of the representation on the display (control-display gain of 1). In the medium gain condition, the the dimensions of the working space for the hand were 175 by 140 mm, so that each mm of movement in hand space resulted in two mm on the display (control-display gain of 2). In the high gain condition, the dimensions of working space for the hand were reduced to 87.5 by 70 mm, mapping each mm of hand movement to four mm on the display (control-display gain of 4).

Subjects performed three blocks of trials, one in each control-display gain condition. A block started with a practice session to allow the subject to achieve stable performance for the current gain setting. During practice, all subjects were eventually able to reach the criterion of four targets in a row quickly, without corrections. The order of presentation of gains was counterbalanced across the six subjects.

For each gain, combinations of five target widths (W = 3, 6, 12, 24, and 48 mm on the display) and five movement distances (A = 18.75, 37.5, 75, 150, and 300 mm on the display) were randomly presented within a block until the subject had performed twelve trials in each of the A and W combinations of interest. Sufficient A and W combinations were selected in each gain condition to provide at least nine (three A by three W) of identical size and index of difficulty (ID) for both the display space and hand space analyses, as depicted in Figure 3 and explained below. The start mark, and all target positions were centred in hand and display space, so that movement direction for all trials was away from the body in the sagittal plane.

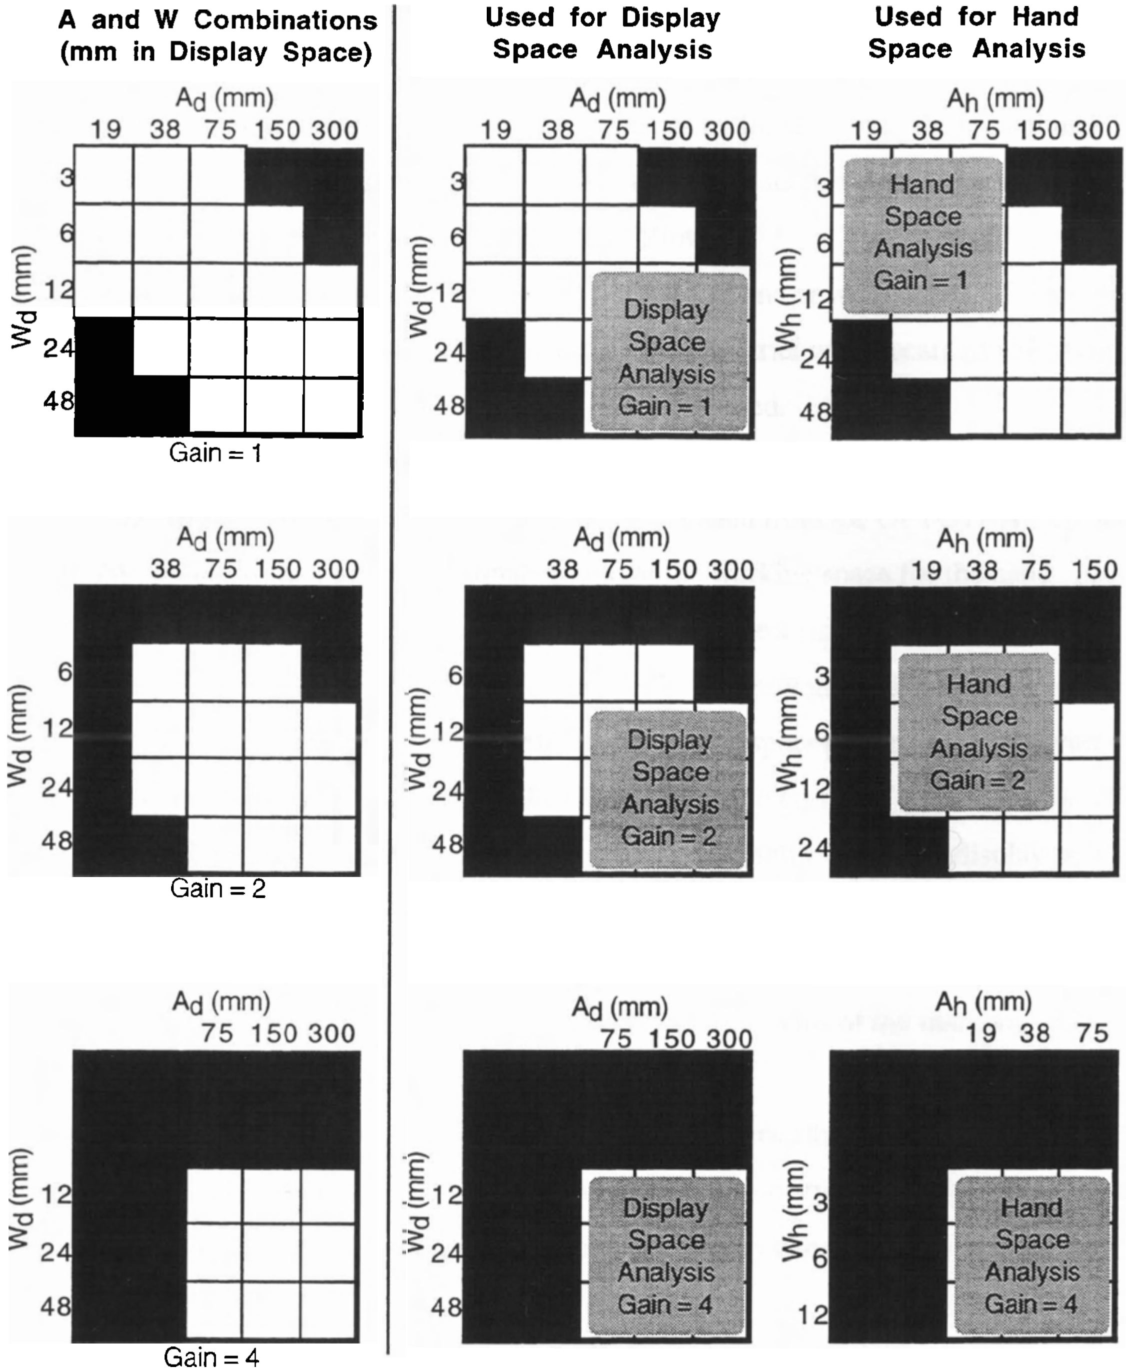

Figure 3. On the left, the movement amplitude and target width combinations (white squares) in display space tested in each gain condition. On the right, the subsets (grey squares) of three amplitudes and three widths, for a total of nine combinations, used for the display space analysis, and the hand space analysis. Note that at gain = 2 and gain = 4, A and W are smaller in hand space than in display space.

Data Analysis

Files of OPTOTRAK position data for each trial were processed by various computer programs to produce dependent measures, as outlined below. Both the three-dimensional position of the first marker (hand position), and its two-dimensional projection on the display (cursor position) were analysed as follows:

1) Missing data points were interpolated over gaps of one or two frames. If there were any gaps greater than two frames (33 ms), the trial was discarded. Missing data points were usually due to IREDs being obscured.
2) Using the reference frame recorded at the start of a block of trials, three-dimensional position data were translated and rotated from the OPTOTRAK camera coordinate system to the coordinate frame of the working space for the hand. From the subject's point of view, the origin was at the closest right corner of the work space, with the x axis pointing away parallel to the body midline, the y axis to the left, and z axis upward. This corresponded to the display projection used during the experiment, with the origin at the lower right hand corner of the screen, the x axis pointing upward, the y axis pointing to the left, and z out of the display surface towards the subject.
3) Data were filtered using an 8 hertz low-pass second-order bidirectional Butterworth digital filter to remove digital sampling artifacts, vibrations of the markers, and tremor from the hand movements.
4) Three- and two-dimensional position data were numerically differentiated using a three-point algorithm to provide an estimate of the velocity profile along the x, y, and z axes, and resultant for the hand; and the x and y axes, and resultant for the pointer.
5) A computer program determined the start and end of movement for each trial from the three-dimensional velocity profile. The start of movement was defined to be the first point at which the resultant velocity was greater than 4 mm/s and continued to increase to a criterion value of 100 mm/s. To determine the end of movement, a 400 ms window was advanced in time over the velocity profile, looking for the first occurrence of a mean velocity over the time window less than 12 mm/s. The end

point was then defined to be the first sample (within the window) where the resultant velocity was less than 12 mm/s. The start and end were then confirmed by inspecting a graph of the velocity profile. The criterion values for the above algorithm were chosen to produce similar values for start and end of the movement as would be determined by visually examining a graph of the velocity profile. In case of disagreement between the experimenter and the computer program, that trial was rejected for analysis. At this point, trials were also rejected if the algorithm failed to find a start and end, the trial included a false start, or if the subject failed to make a proper movement in the direction of the target.

6) Acceptable velocity profiles were analysed by a computer program on a trial by trial basis to determine the timing and magnitude of the first peak in resultant velocity after the start of movement.
7) The resultant of each velocity profile was further numerically differentiated to provide an estimate of acceleration along the movement path. The timing and magnitude of the first positive peak after the start of movement, and the first negative peak (after peak velocity) were determined by a computer program.

Dependent measures outlined below were calculated, using individual subject means over the last ten good trials in each gain by amplitude by width combination. The reader who is unfamiliar with analysis techniques to characterise features of movement trajectories is directed to Figure 4, which illustrates how dependent measures are calculated for a representative trial. Measures calculated from three-dimensional hand trajectories are coded by standard acronyms (e.g. MT, TPV). For corresponding measures calculated from two-dimensional cursor trajectories on the display, we have prefixed the standard acronym with the letter C:

MT Movement time calculated as the time in milliseconds from the start to the end points as determined above from the three-dimensional velocity profile;

TPA Time in milliseconds from the start of the movement to the first peak in the acceleration profile for the hand;

PA The magnitude of acceleration along the three dimensional path of hand movement at TPA in mm/s/s;

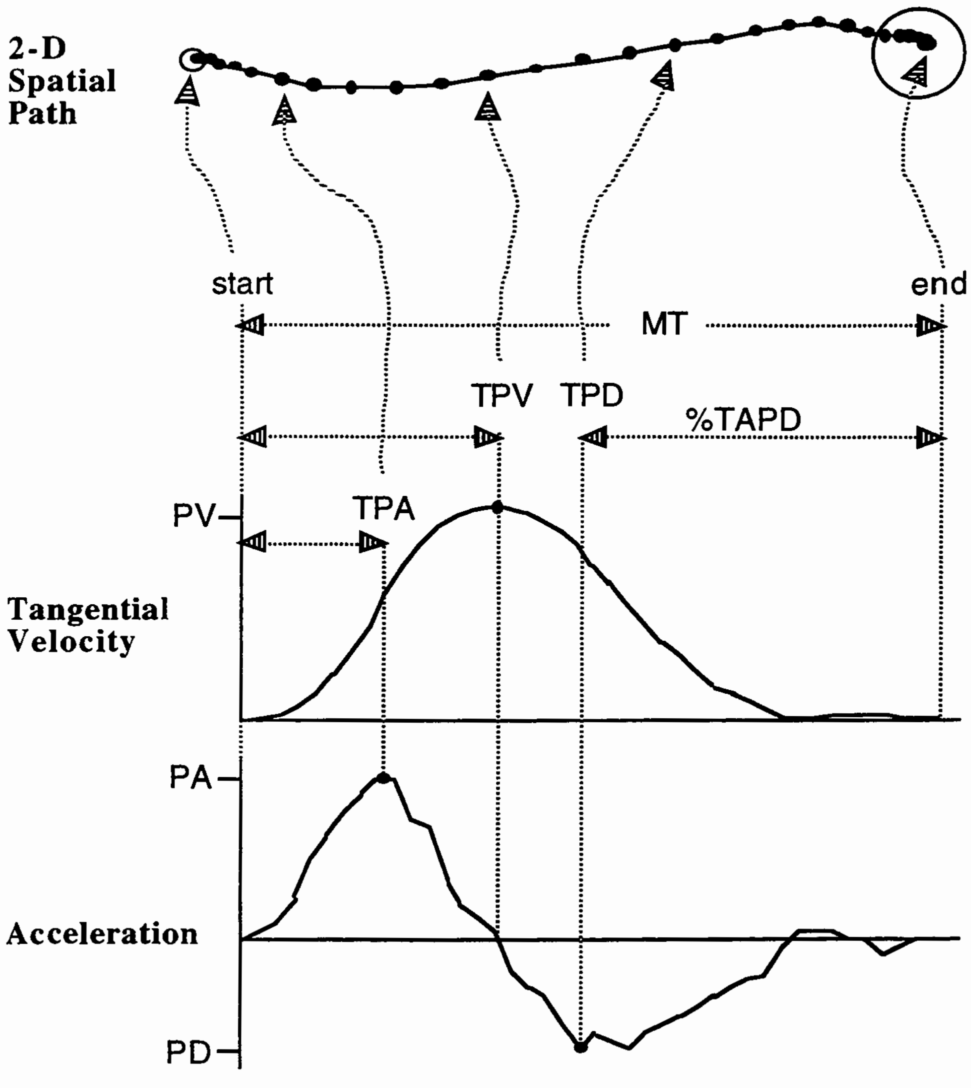

Figure 4. Pointing movement kinematics. The two-dimensional spatial path for a representative trial is shown at the top; each dot represents one sample of data taken at 60 Hz. The corresponding velocity and acceleration profiles are shown below. Annotations indicate the measurements of timing and magnitudes of peaks used to derive dependent measures for statistical analysis in the experiment.

CTPA Time in milliseconds from the start of the movement to the first peak in the acceleration profile for the cursor;

CPA The magnitude of acceleration along the two-dimensional path of cursor movement at CTPA in mm/s/s;

TPV Time in milliseconds from the start of movement to the first peak in the velocity profile for the hand;

PV The magnitude of velocity along the three-dimensional path of hand movement at TPV in mm/s;

CTPV Time in milliseconds from the start of the movement to the first peak in the velocity profile for the cursor;

CPV The magnitude of velocity along the two-dimensional path of cursor movement at CTPV in mm/s;

%TAPD Per cent time after peak deceleration, the proportion of movement time after the first negative peak in the acceleration profile for the hand, occurring after TPV.

C%TAPD Cursor per cent time after peak deceleration, the proportion of movement time after the first negative peak in the acceleration profile, occurring after CTPV;

Wex Effective target width in the principal direction of movement, derived from the standard deviation of hand position (S. MacKenzie, 1992) on the x-axis at the end point of the movement;

Wey Effective target width orthogonal to the principal direction of movement, derived from standard deviation of hand position on the y-axis at the end point of the movement.

Results

Separate analyses were performed on dependent measures derived from two-dimensional cursor kinematics (display space analysis), and on measures derived from three-dimensional hand kinematics (hand space analysis). Figure 3 illustrates the combinations of movement amplitude (A_d) and target width (W_d) for which data were collected in each gain condition, in terms of dimensions on the display. For the display space analysis, a subset of conditions was chosen to include three identical target widths and three identical movement amplitudes in each gain condition. In this analysis, the effect of gain was to reduce both movement amplitude (A_h) and target width (W_h) for the hand, while leaving them unchanged on the display. The hand space analysis used a different subset of conditions where A_h and W_h were matched for different gains. In this analysis, the effect of gain was to increase the size of A_d and W_d, while leaving amplitude and target width unchanged for the hand. Both analyses used ANOVA with repeated measures on a Subjects x Gain x Amplitude x Width (6 x 3 x 3 x 3) design. The main effects for gain, amplitude, and target width are summarised in Tables 1 through 6. F and p values are also given in the tables for each effect; these are not included in the text.

Display Space Analysis

The main effects for gain on cursor kinematics are summarised in Table 1. Movement time decreased for higher gains; the effect was significant and large (98 ms). A significant interaction ($F_{4,20} = 7.56$, $p < .0007$) between gain and movement amplitude on the display is illustrated in Figure 5. The increase in MT with amplitude predicted by Fitts' law was most pronounced in the lowest gain condition, and the size of the effect was reduced as gain was increased.

The magnitude of the initial acceleration peak (CPA) was also affected significantly by gain. The peak was almost twice as large (6155 versus 3411 mm/s/s) for the highest gain compared to the lowest gain condition. There was no significant effect of gain on the timing of the initial acceleration peak (CTPA), but a significant gain by amplitude interaction ($F_{4,20} = 3.28$, $p < .032$) is illustrated in Figure 6. At the lowest gain, the timing of the peak is systematically delayed as movement amplitude increases. At the highest gain, timing is relatively constant for all amplitudes.

Gain also significantly affected the magnitude and timing of the initial velocity peak. The magnitude of the peak (CPV) increased (485 to 571 mm/s) from the lowest to the highest gain, and CTPV occurred earlier (233 versus 292 ms) for the highest gain compared to the lowest gain condition. The interaction between gain and movement amplitude ($F4,20 = 6.36$, $p < .0018$) revealed a small effect illustrated in Figure 7. The

TABLE 1

Display space analysis: Summary of main effects of gain on two-dimensional cursor kinematics

	Control-Display Gain 1	2	4	$F_{2,10}$	p <
MT (ms)	858	791	760	6.73	.0141
CPA (mm/s/s)	3411	4145	6155	20.30	.0003
CTPA (ms)	91.7	86.5	80.7	2.77	.1102
CPV (mm/s)	485	512	571	19.00	.0004
CTPV (ms)	292	274	233	25.87	.0001
C%TAPD	53.1	56.3	65.0	17.79	.0005
Wex (mm)	15.7	18.5	18.9	5.53	.0241
Wey (mm)	10.2	12.7	14.8	8.04	.0083

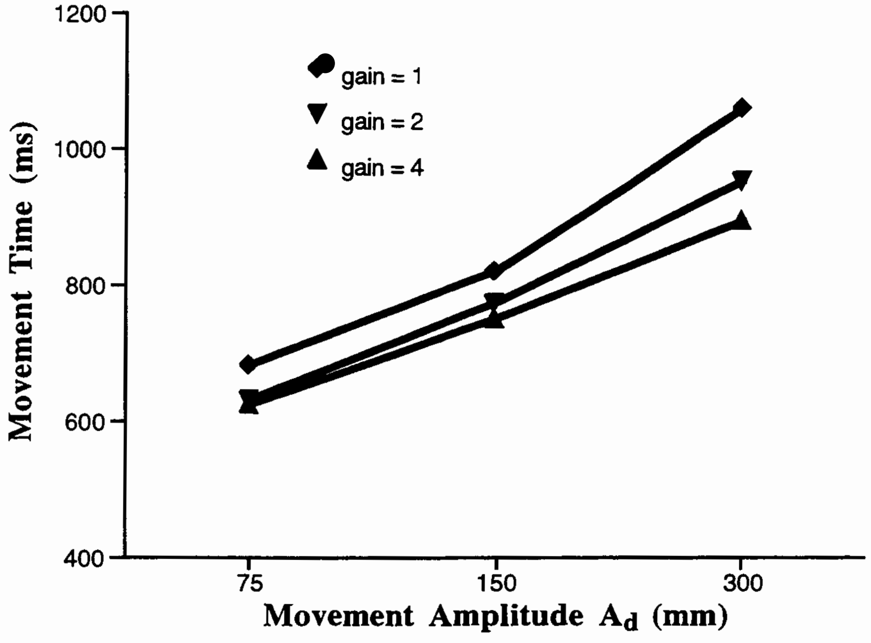

Figure 5. The interaction between gain and movement amplitude for movement time on the display. The increase in MT for larger amplitudes is most pronounced in the lowest gain condition, while the size of the effect is reduced for higher gains.

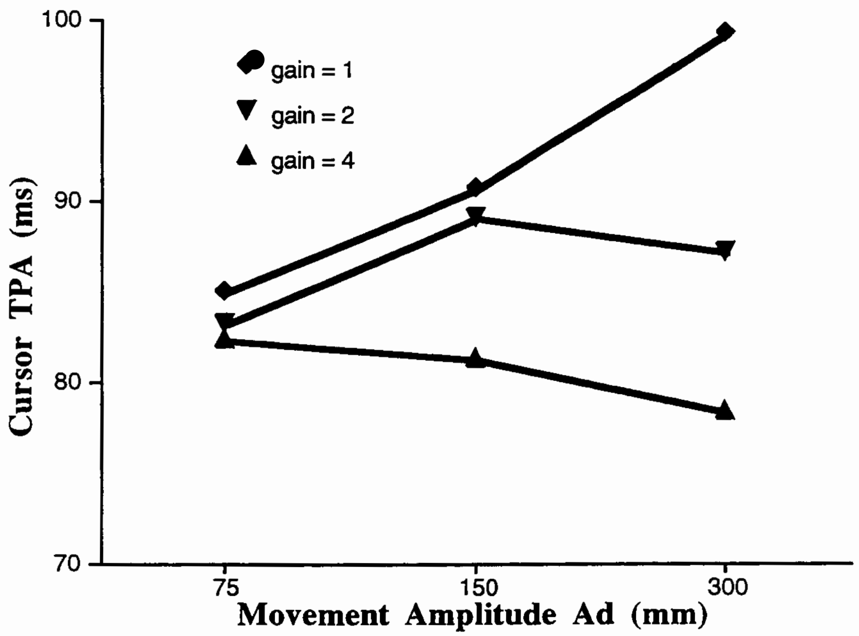

Figure 6. Timing of the initial acceleration peak for the cursor showed a gain by amplitude interaction. At the lowest gain, the timing of the peak is systematically delayed as movement amplitude increases. At the higher gains, timing is relatively constant for all amplitudes.

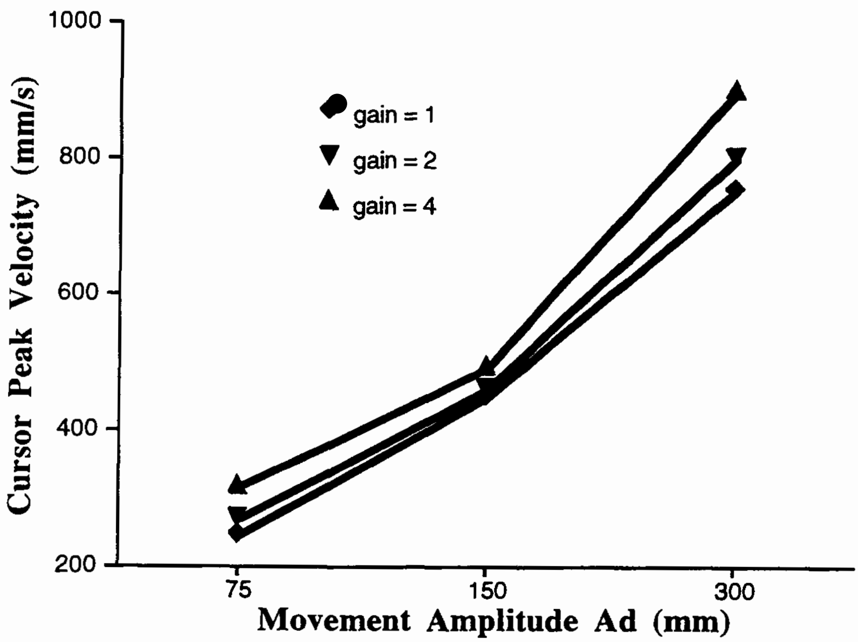

Figure 7. The interaction between gain and movement amplitude on cursor peak velocity. The magnitude of the peak increased more with increased amplitude for the lowest gain than in the highest gain condition.

magnitude of the velocity peak increased more with increased amplitude for the lowest gain than in the highest gain condition.

Gain also affected the proportion of time decelerating to the target (C%TAPD). Figure 8 illustrates a significant gain by amplitude interaction ($F_{4,20} = 303.25$, $p < .0005$). At the lowest gain, the proportion of time in deceleration increased for larger distances. At the medium gain, no differences were evident, but for the highest gain, C%TAPD decreased for larger distances.

Significant effects of gain on effective target width in both the x (Wex) and y (Wey) directions indicate the same trend. As gain increased, the variability of the end point also increased. Overall, there was more variability, and thus a larger effective target width, in the x direction along the principal axis of movement compared to the y direction.

The main effects for movement amplitude are summarised in Table 2. As predicted by Fitts' law, movement time increased for larger distances, from 650 ms for the smallest, to 973 ms for the largest amplitude. The effects of amplitude were mainly confined to the initial phase of the movement. The magnitude of the initial velocity (CPA) and acceleration (CPV) peaks increased systematically for larger, compared to smaller amplitudes. In addition, the initial velocity peak occurred later for larger distances, with CTPV varying from 223 ms for the smallest, to 309 ms for the largest amplitude.

Amplitude showed no significant effect on the proportion of time spent in deceleration to the target. Effective target width in the direction of movement (Wex) showed a slight increase for longer distances, but this effect was not significant. The effective width orthogonal to the direction of movement (Wey) showed a significant increase (11 to 14.3 mm) from smallest to largest movement amplitude.

Table 3 summarises the main effects of target width. Movement time increased for smaller targets, as predicted by Fitts' law, from 725 ms for the largest, to 884 ms for the smallest target. There were no significant effects of width on the timing or magnitude of the initial acceleration peak. Width showed a significant, but small effect on the timing and magnitude of the initial velocity peak. From the smallest to the largest target, the magnitude increased slightly (CPV from 513 to 537 mm/s), and the peak occurred later (CTPV from 262 to 272 ms).

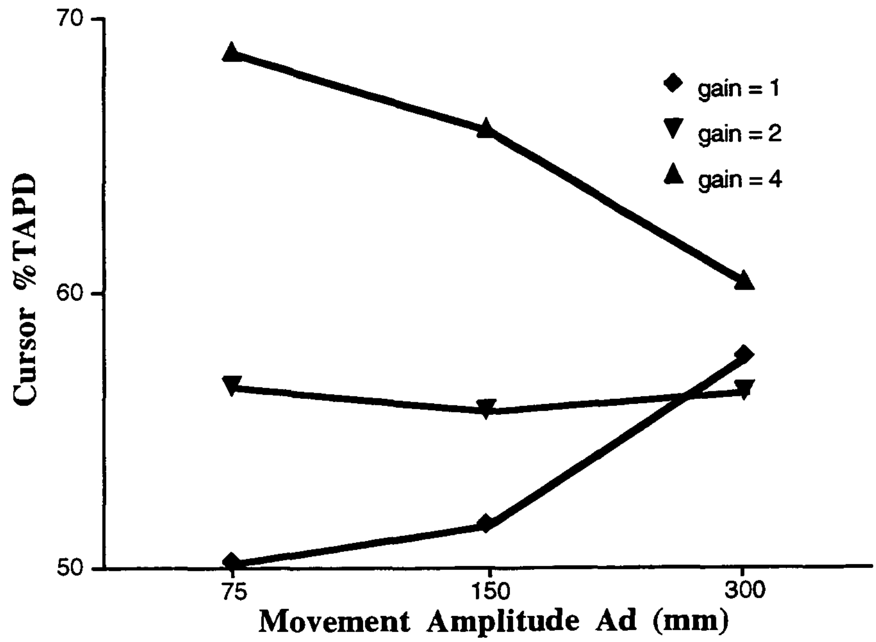

Figure 8. Proportion of time in deceleration for the cursor showed a gain by amplitude interaction. At the lowest gain, the proportion increased for larger distances. At the medium gain, no differences were evident, but for the highest gain, the proportion decreased for larger distances.

TABLE 2

Display space analysis: Summary of main effects of movement amplitude on two-dimensional cursor kinematics

	A_d (mm)				
	75	150	300	$F_{2,10}$	p <
MT (ms)	650	785	973	206.85	.0001
CPA (mm/s/s)	3586	4591	5533	47.07	.0001
CTPA (ms)	83.5	87.1	88.3	1.87	.2043
CPV (mm/s)	278	469	821	329.63	.0001
CTPV (ms)	223	266	309	55.05	.0001
C%TAPD	58.5	57.8	58.2	.06	.9427
Wex (mm)	16.5	18.0	18.5	1.67	.2364
Wey (mm)	11.0	12.4	14.3	7.93	.0087

TABLE 3

Display space analysis: Summary of main effects of target width on two-dimensional cursor kinematics

	W_d (mm)				
	12	24	48	$F_{2,10}$	p <
MT (ms)	884	800	725	40.05	.0001
CPA (mm/s/s)	4602	4543	4565	.21	.8120
CTPA (ms)	87.1	86.6	85.3	.52	.6107
CPV (mm/s)	513	518	537	13.84	.0013
CTPV (ms)	262	265	272	4.88	.0332
C%TAPD	62.8	58.7	52.9	22.48	.0002
Wex (mm)	9.27	16.0	27.8	51.39	.0001
Wey (mm)	8.07	12.39	17.2	18.16	.0005

The largest effects of width were concentrated in the deceleration phase. Proportion of time after peak deceleration (C%TAPD) increased systematically from the largest to the smallest target (52.9 to 62.8 per cent), demonstrating a precision effect. Both effective target width in the direction of movement (Wex), and orthogonal to the direction of movement (Wey) increased significantly as the actual target width increased. Both Wex and Wey exhibited a range effect, where the change in effective target widths was somewhat less than the change in dimensions of the actual targets.

Hand Space Analysis

The main effects of gain on three-dimensional hand kinematics are summarised in Table 4. The only significant main effect was on the magnitude of the initial velocity peak which decreased slightly from 187 to 161 mm/s from the lowest to the highest gain condition. This was also evident in a gain by amplitude interaction ($F_{4,20} = 7.17$, $p < .0009$) illustrated in Figure 9. The magnitude of the amplitude effect, where TPV increased for larger compared to smaller distances, was slightly greater at higher, compared to lower gains. Proportion of time in the deceleration phase (%TAPD) also showed a gain by amplitude interaction ($F_{4,20} = 4.04$, $p < .015$) as illustrated in Figure 10. The %TAPD remained relatively constant for all distances in the lowest gain conditions, but was noticeably greater for the two shortest distances in the highest gain condition. Overall, the magnitude of these effects were very small, and other main effects for gain did not approach statistical significance.

Table 5 summarises the main effects for amplitude, in terms of distances covered by the hand to reach targets, on three-dimensional hand kinematics. As expected, movement time increased systematically from the smallest to the largest amplitude. Amplitude effects were most evident in the initial phase of the movement. The magnitude of the initial acceleration peak (PA) increased as distance to the target increased, whereas amplitude showed no significant effect on the timing of the peak (TPA). The timing and magnitude of the initial velocity peak showed systematic effects due to amplitude. The magnitude of the peak (PV) increased from 120 to 251 mm/s, and the timing varied from 154 to 230 ms from the smallest to largest movement amplitude.

TABLE 4

Hand space analysis: Summary of main effects of gain on three-dimensional hand kinematics

	Control-Display Gain 1	2	4	$F_{2,10}$	p <
MT (ms)	724	724	760	1.44	.2818
PA (mm/s/s)	2611	2425	2234	.99	.4508
TPA (ms)	87.9	84.2	77.0	1.73	.2263
PV (mm/s)	187	179	161	5.67	.0226
TPV (ms)	188	190	185	.16	.8568
%TAPD	67.9	66.6	71.4	3.07	.0915
Wex (mm)	5.14	5.33	4.73	1.74	.2248
Wey (mm)	4.03	4.03	3.70	.97	.4124

TABLE 5

Hand space analysis: Summary of main effects of movement amplitude on three-dimensional hand kinematics

	A_h (mm) 18.75	37.5	75	$F_{2,10}$	p <
MT (ms)	599	723	866	153.63	.0001
PA (mm/s/s)	2109	2340	2821	48.29	.0001
TPA (ms)	83.5	84.3	81.3	.56	.5907
PV (mm/s)	120	156	251	659.96	.0001
TPV (ms)	154	179	230	47.01	.0001
%TAPD	68.7	70.2	67.1	1.87	.2037
Wex (mm)	4.90	5.03	5.26	1.07	.3804
Wey (mm)	3.52	3.79	4.44	24.83	.0001

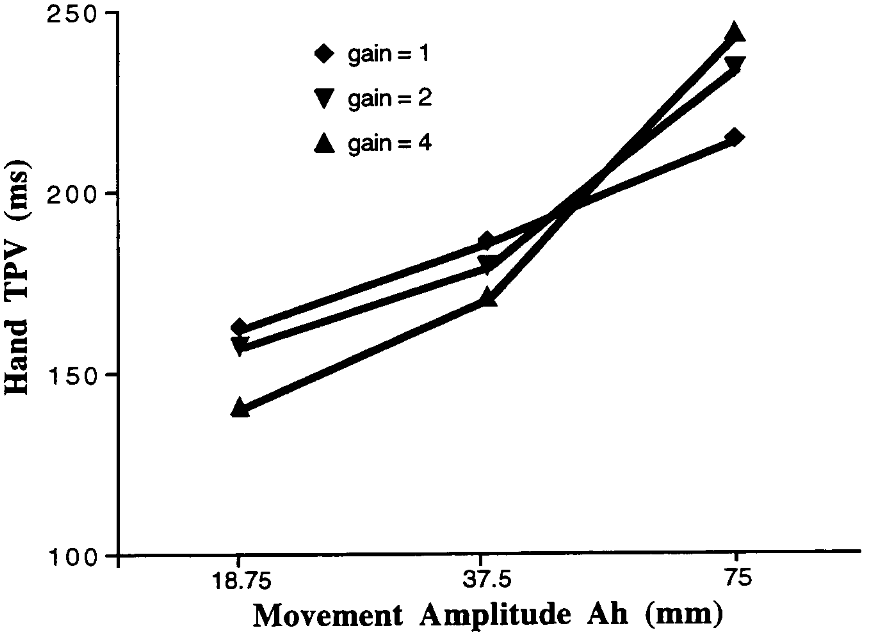

Figure 9. The timing of the initial velocity peak for the hand demonstrated a small gain by amplitude interaction. The peak occurred later for larger amplitudes, and this effect was exaggerated for high, compared to low gains.

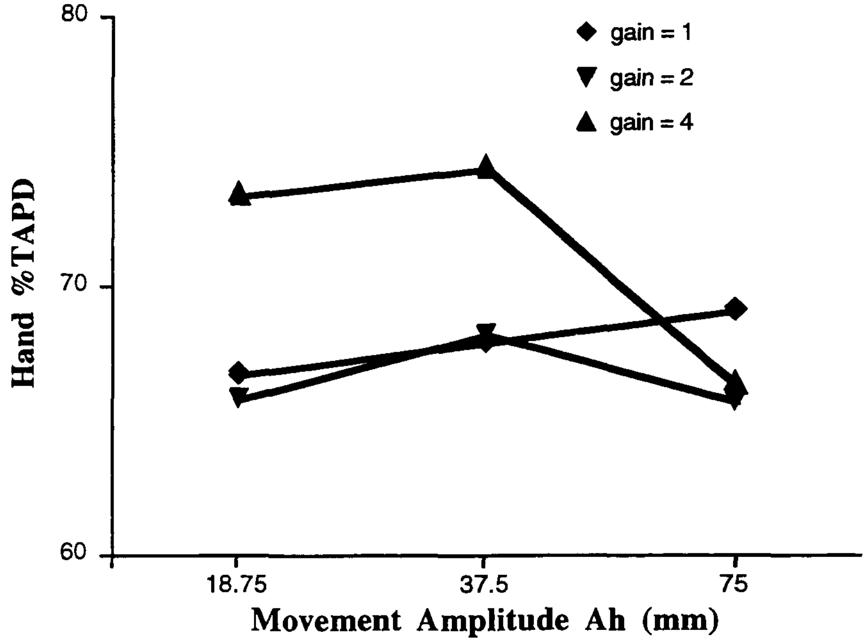

Figure 10. The proportion of time after peak deceleration for the hand showed an amplitude by gain interaction. At the lowest gain %TAPD increased slightly for larger amplitudes. At the highest gain, %TAPD actually decreased for larger, compared to smaller amplitudes.

There was no significant effect of amplitude on proportion of time spent in deceleration (%TAPD). Effective target width in the direction of movement (Wex) was not affected by amplitude, although there was a slight increase in effective width orthogonal to the direction of movement (Wey) for larger amplitudes.

Table 6 displays the main effects of target width on three-dimensional hand kinematics. Movement time decreased for larger compared to smaller targets, as predicted by Fitts' law. There were slight effects on the magnitudes of the initial acceleration peak (PA) and velocity peak (PV). Both peaks increased for larger, compared to smaller distances, but both these effects were very small. The timing of the peaks was not significantly affected by target width.

Width effects were concentrated in the deceleration phase. The proportion of time in deceleration (%TAPD) increased from 58 to 71.6 per cent from the largest to the smallest target, demonstrating a large precision effect. In addition, effective target widths both in the direction of movement (Wex) and orthogonal to the direction of movement (Wey) increased as actual target size increased. As observed above in the display space analysis, effective target width exhibited a slight range effect in that variations in We were somewhat less than variations in actual target width across experimental conditions.

The main effects from both display space and hand space analyses are summarised graphically in Figure 11. In comparing the two sets of graphs, an obvious difference is the number of significant and large effects for gain in display space, compared to one slight effect for gain in hand space. The lack of gain effects in hand space suggests that for a given hand movement, A and W on the display can be varied with very little effect on movement kinematics. Conversely, for a given cursor movement, varying distances and target widths for the hand changes not only the features of the hand movement but also the movement kinematics of the pointer on the display. Overall, the hand space analysis demonstrates that movement kinematics vary systematically only as a function of A_h and W_h, and provides a clearer picture of movement planning and control. In display space the picture is more complicated; in addition to A_d and W_d, control-display gain is also required to explain systematic variations in kinematic features of the pointer movement.

TABLE 6

Hand space analysis: Summary of main effects of target width on three-dimensional hand kinematics

	W_h (mm)				
	3	6	12	$F_{2,10}$	p <
MT (ms)	811	732	645	56.94	.0001
PA (mm/s/s)	2363	2422	2485	4.69	.0365
TPA (ms)	83.3	83.9	81.9	.38	.6948
PV (mm/s)	173	174	179	9.27	.0053
TPV (ms)	189	187	187	.22	.8085
%TAPD	71.6	68.8	58.0	21.92	.0002
Wex (mm)	2.94	4.52	7.73	68.25	.0001
Wey (mm)	2.67	3.65	5.43	22.26	.0002

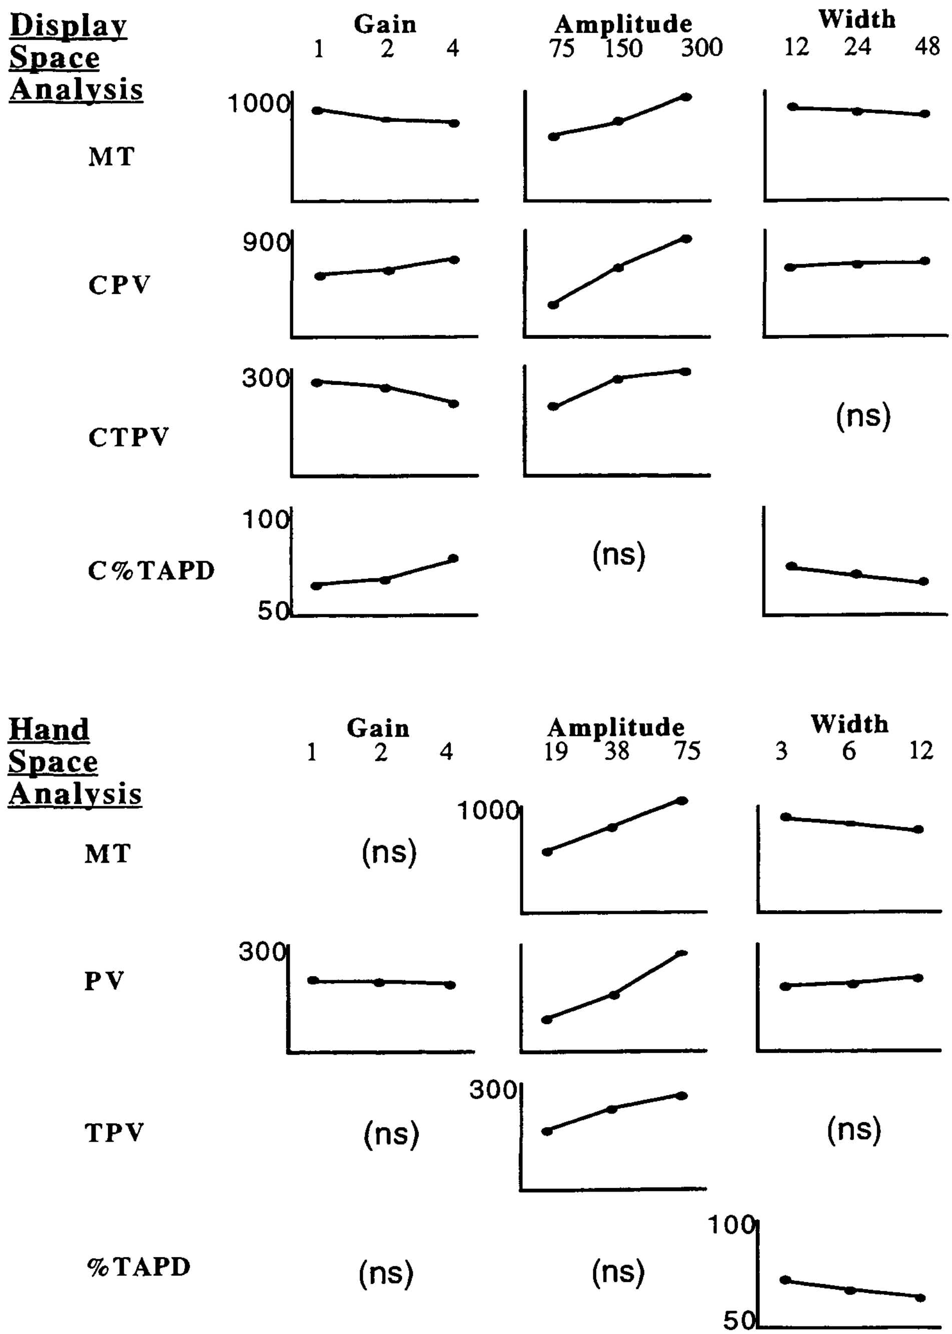

Figure 11. Graphical overview showing main effects from display space analysis (upper, from Tables 1, 2 and 3) compared to main effects from hand space analysis (lower, from Tables 4, 5 and 6) on a subset of dependent measures. Individual graphs indicate effects from the ANOVA with $p > .05$, while non-significant effects are labelled (ns). There are fewer significant effects for gain in hand space, compared to display space. The magnitude of gain effects in display space is considerable, while the magnitude of the gain effect in hand space (PV) is very small.

Modelling Movement Times with Fitts' Law

Considering the effects on movement time from the hand and display analyses presented above, one feature appears counterintuitive. In display space, MT decreased from 858 to 760 ms from the lowest to the highest gain. The magnitude of this effect was almost as large as the effect for target width, where MT decreased from 884 to 725 ms for the smallest, compared to largest target. In hand space, there was no significant effect of gain on movement time, as would be predicted by Fitts' law. Why should movement times in display space vary so widely in conditions where the index of difficulty was held constant? In addition, it is not apparent how this effect could appear in display space, but not in hand space, given that amplitude and width effects on all kinematics in both analyses showed similar patterns. The obvious difference between the analyses was the large number of gain effects in display space, compared to one significant, small gain effect in hand space.

Data from the low gain condition (gain = 1) were modelled in terms of Fitts' Law. This condition contained the largest number of A and W combinations, and since there were no significant effects of gain in the hand space analysis of timing of kinematic measures, this subset of the data was felt to represent a clear picture of the essential features of the task in terms of speed-accuracy tradeoffs. Fitting means across subjects of MT versus the index of difficulty for each A and W combination, from Equation 1:

$$MT\ (ms) = 351 + 118\ ID \qquad (R^2 = .75) \qquad (4)$$

which accounts for only 75% of the variance. The data points and regression line are shown in Figure 12, where it can be noted that there is a considerable spread of values for different A and W combinations within given ID. Individual subjects' data were also fit to Equation 1, yielding correlations ranging from .46 to .82.

The rather poor correlations and range of values within ID led to another approach due to Welford (1968), who attempted to characterise movement time data with a similar spread. His task involved repetitive pointing with a pencil, and demonstrated a characteristic which required a more detailed model than Fitts' law: the rate of change of movement time for changes in target width was markedly different than for changes in

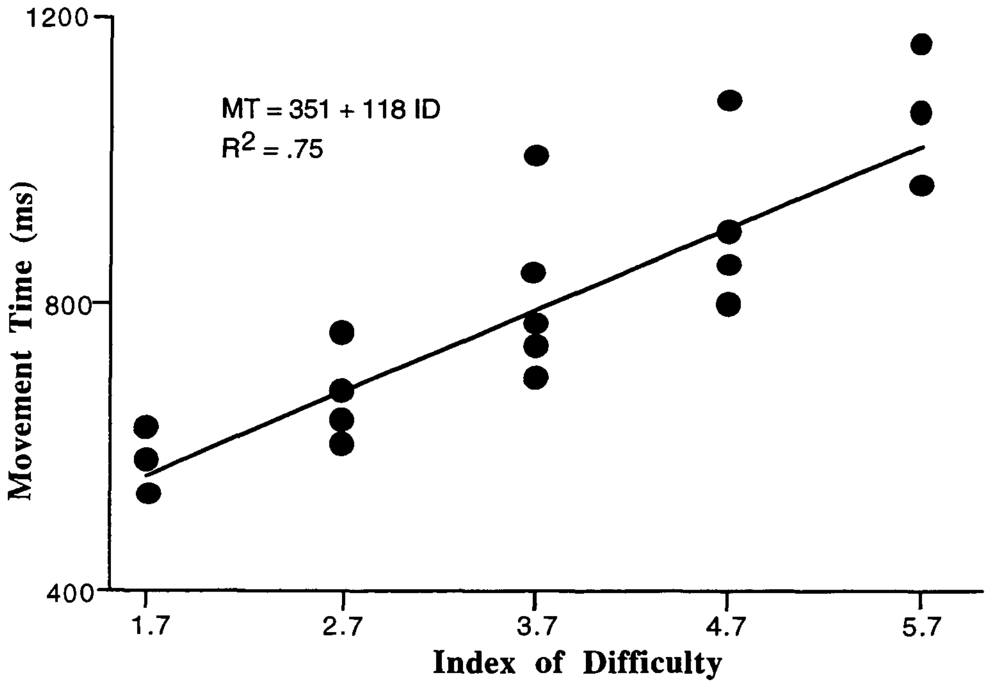

Figure 12. Movement time data for all amplitude-width combinations in the gain = 1 condition. The vertical spread of the points at each index of difficulty indicates that movement time is affected by proportional scaling of A and W for the hand. Within each index of difficulty, movement time increases as A and W are increased by factors of two.

movement amplitude. He derived a model which separates the contributions of distance-covering (movement amplitude) and accuracy (target width) to total movement time. This two-part model can be simplified to a form which allows a fit using multiple linear regression:

$$MT = a + b_1 \log_2 A - b_2 \log_2 W \qquad (5)$$

In the special case that b_1 and b_2 are equal, the Welford two-part model is equivalent to Fitts' law. Means from the present experiment across subjects for all combinations of A and W in the low gain condition were fitted to the two-part model, with the result:

$$MT\ (ms) = 129 + 153 \log_2 A - 83 \log_2 W \qquad (R^2 = .96) \qquad (6)$$

accounting for 96 per cent of the variance of movement times. The effect of the separate coefficients for the A and W terms can be understood by comparing the graphs in Figure 13. In the upper graph, actual data points for equal W are considered in series. In this case, the increase in ID on the abscissa is due to increasing movement amplitude, and MT increases at approximately 153 ms/bit as a function of ID. In lower graph, the series are actual data points for equal A. In this case, the increase in ID on the abscissa is due to decreasing target width, and MT increases at approximately 83 ms/bit as a function of ID.

In the present experiment, targets were relatively small compared to their distances from the start mark in each trial. The end point data were checked to verify that the average distance subjects' moved in each A and W combination closely matched the actual distance to the target. For target width, there was no similar constraint forcing subjects to make use of the entire target area. A possible explanation for the asymmetry of amplitude and width effects in the data is that subjects were not taking advantage of the relaxed accuracy constraint of larger targets to reduce their movement times. In order to examine this possibility we fit the same movement time data to the two-part model, but using effective, instead of actual target width. Effective target width (We) was estimated from the end-point variability in the direction of movement (Wex), since this measure captured more of the positioning errors than variability orthogonal to the direction of movement (Wey).

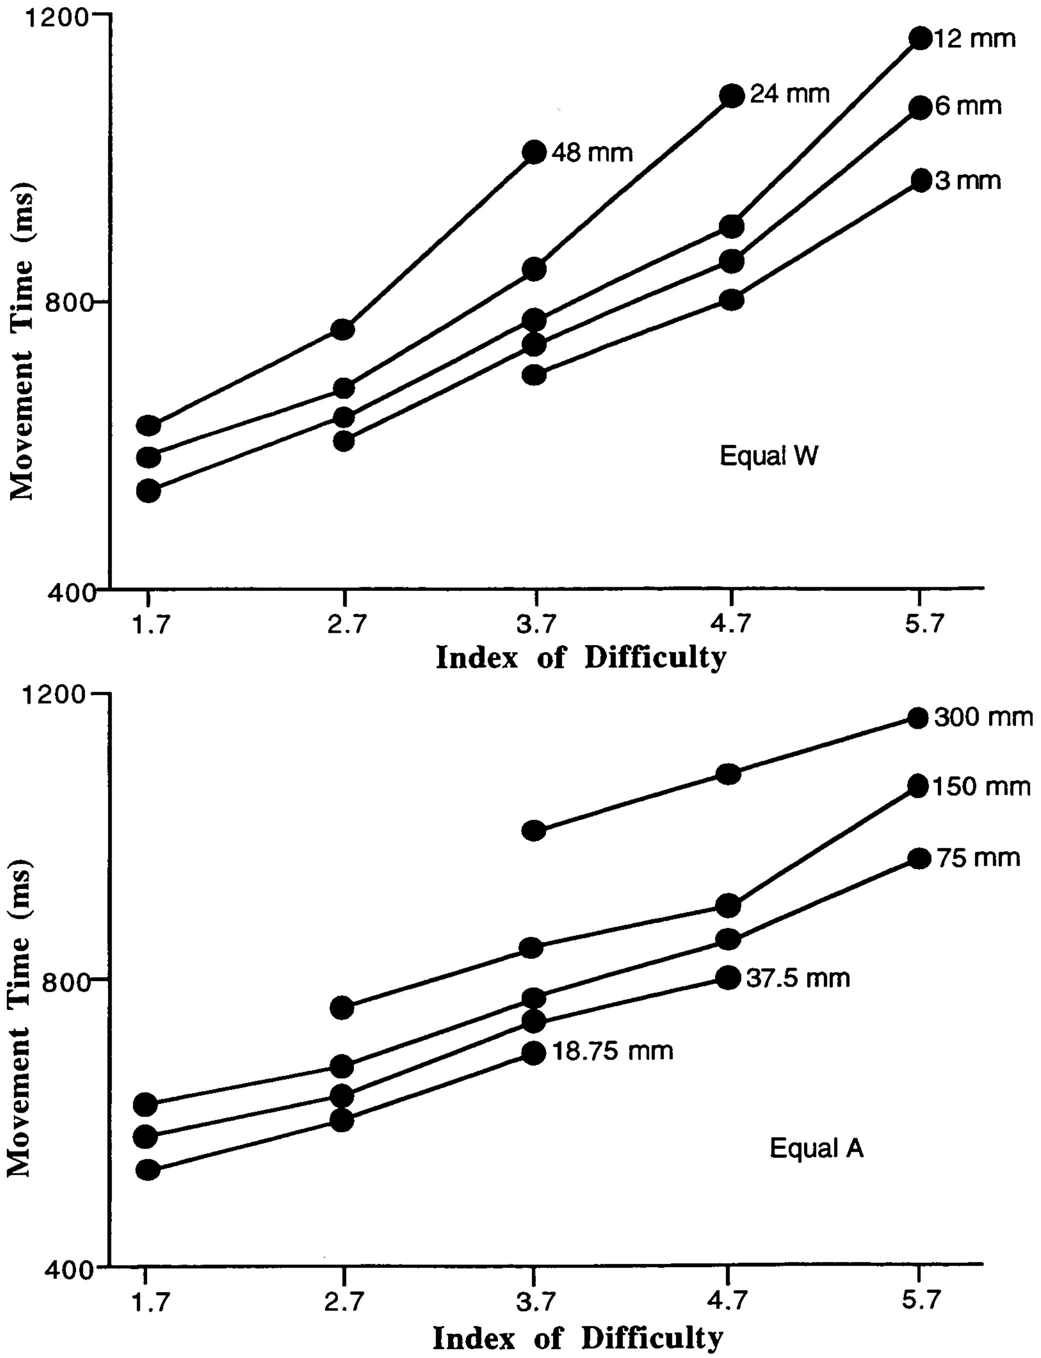

Figure 13. Movement time data from Figure 12 replotted to illustrate the relative magnitudes of amplitude and width effects. Above, the slope of lines connecting points associated with the same target width (W_h) indicates how movement time (MT) increases as a function of amplitude (A_h). Below, lines between points associated with equal A_h, show how MT increases with decreasing in target size (W_h).

Using effective target width:

$$MT\ (ms) = 351 + 151 \log_2 A - 102 \log_2 We \quad (R^2 = .96) \qquad (7)$$

the two-part model accounts for 96 per cent of the variance in movement times. The essential feature of the data, that is the fact that MT is more sensitive to change in A than in W, is still preserved. Thus, the model of Equation 6 effectively describes human performance in the task.

Discussion

Theory

As in other studies (Welford, 1968) involving large variations in distance, target width has less effect than distance on movement time. If target width were held constant, and ID varied by changing movement amplitude, the present experiment would have produced data which would be well characterised by Fitts' law. Similarly, if targets of various sizes were presented at a constant distance, the data would be also well characterised by Fitts' law, albeit with a different index of performance than the previous example. In fact, Fitts' model holds for small subsets of data from the present experiment, and from aggregated data even when movement time is markedly more or less sensitive to changes in the distance to be covered than to changes in the accuracy constraint provided by the target. The two-part model simply provides a more complete description of the data than Fitts' law, and accounts for the major difference between the pattern of results for the display space analysis versus the hand space analysis as we will explain below.

First, the lack of gain effects in hand space supports the idea that movement distance for the hand is the key variable for movement planning. Timing of the initial velocity peak is shown to depend on distance for the hand (A_h), regardless of distance on the display (A_d). In addition, the proportion of time decelerating to the target is a function of target size as an accuracy constraint for the hand (W_h), regardless of the size of the visually presented target (W_d) on the display. The results from the hand space analysis clearly show separable A_h and W_h effects on the velocity profile, A_h accounting for its overall size (PV and TPV), and W_h determining its shape or skew (%TAPD). What the two-part model reveals is that the magnitude of the amplitude effect on the size of the velocity profile is much greater than the magnitude of the width effect. This asymmetry results in an increase in movement time for the same index of difficulty when both target width and amplitude are increased in proportion. Similarly, if both amplitude and width are decreased in proportion, movement time will decrease. This is evident in the hand space analysis, where the effect of increased gain is to reduce the scale of hand movements for a given index of difficulty, reducing A_h and W_h in proportion. Thus, the effect of increasing gain is to reduce cursor movement times on the display. The same argument also provides a reasonable explanation of the gain effects evident in the display space analysis, for which there are no corresponding gain effects in hand space.

One interpretation of the two-part model, due to Welford, is that the A term (b_1) is related to distance-covering, and hence movement planning, and the W term (b_2) arises from difficulty with visual control when homing in on the target. Similar to the index of performance derived from Fitts' law, a larger b_2 coefficient indicates greater sensitivity to changes in target size, and can be interpreted as a reduced bandwidth for the corresponding channel for movement control.

In describing data from the present experiment, the larger coefficient for the amplitude term (b_1) in Equation 6 might be explained by a dissociation, compared to pointing at physical targets in the work space, between natural mechanisms for movement planning and visual control. That is, in the HCI context, movement planning requires a spatial transformation to predict the effects of hand movements on the cursor, in terms of distance and direction. This additional information processing load may result in a less accurate movement plan, making speed-accuracy tradeoffs more difficult, with the result that movement time increases more dramatically for larger movement amplitudes.

Application

Overall, the results suggest that movement planning and control are based in hand space, and that the size of the visual representation of the virtual environment on the display has little effect on movement kinematics. The obtained asymmetrical amplitude and width effects on movement time suggest an explanation for the confusion about the effect of control-display gain on performance in previous HCI and human factors studies. As in the present study, gain effects result from a combination of two factors: asymmetry of changes in kinematic features of hand movements due to scaling A and W, and the analysis of performance by measuring cursor (or control) movements rather than hand movements. We refer to the systematic changes in kinematic measures produced by this combination of factors as "scaling effects", effects produced by changing the size of A and W for the hand in proportion so that ID remains constant.

The scaling effects on hand movement profiles reported in this experiment suggest that, for pointing devices like a mouse or tablet, analysis and modelling can be done more effectively by considering the hand actions required for a task, regardless of the scale of the display. At present, most graphical interfaces are programmed using pixel-based coordinates to describe distances between, and sizes of objects on the display. In addition, programming models support device independence (e.g. Foley, Van Dam, Feiner & Hughes, 1991, section 4.6) in order to simplify software development of interfaces based on standard interaction techniques. As a result, the interface designer often has little knowledge about the actual hand movements required of a user to operate an interface. The present study demonstrates that for mediated pointing, a predictive model of movement time for this form of human-computer interaction must be human-centred, and in fact hand-centred, in order to effectively capture the details of user performance under different conditions.

One result from the present study with potential application is the observation that timing and magnitude of the initial velocity peak in a pointing movement is strongly correlated to the distance to be covered by the hand. This was also suggested by Walker et al. (1993) for a task with one degree of freedom, and agrees with results from MacKenzie et al. (1987). The complementary measure chosen here to capture the effect of target width is per cent time after peak deceleration (%TAPD). This measure is strongly related to per cent time after peak velocity (%TAPV), used to indicate the skew of velocity profiles or precision effect. Previous studies show that %TAPV is affected mainly by target width, but also by movement amplitude, and these effects have also been confirmed for data from the present experiment. We found that %TAPD exhibits a truly separable effect for target width, regardless of movement distance.

An intelligent interface could use advance knowledge about object selection to improve the system response through precomputation. This may be useful for two reasons. First, users are typically looking at the target they are about to select, and display resolution can be concentrated on this spot, with less resolution in the peripheral field. Second, if the response of an object on the interface requires complex computation, this calculation can be started in advance of selection in order to improve the response time. The systematic effects observed here suggest that gesture recognition of a pointing movement while the movement is ongoing could be used to detect the magnitude and timing of the initial velocity peak in real time. This information gives a prediction of distance to the target at about the time of peak velocity, usually somewhat less than half of total movement time. By the time of the first deceleration peak it may be possible to estimate the time remaining before the target is reached, based on %TAPD; this time could be used for precomputations in order to respond more quickly to the action of selecting the target object.

Experiment 2: Physical Versus Virtual Pointing

The asymmetrical amplitude and width effects on movement kinematics in Experiment 1 could be a characteristic of mediated pointing, compared to pointing at physical targets in the work space. If this were the case, the asymmetry would indicate an important difference in movement planning and control due to the constraints of a virtual environment. As discussed in the previous experiment, the increased sensitivity of movement time to changes in distance may be due to difficulties in adapting to the constraints of the HCI setup, and forming a sensorimotor transformation to map the direction and extent of hand movements to corresponding cursor movements on the display. Another possibility is that the asymmetry is due to another fundamental difference between the HCI setup and pointing at physical targets: the two-dimensional display presents no visual information about height of the hand above the table top, even though this dimension is important for control of the hand to achieve contact with the target location.

In order to further address the question of the differences between HCI and pointing at physical targets, the previous experiment was repeated using only the unity gain condition, and the same graphical display of pointer and targets as in Experiment 1. In the present experiment, the need for spatial transformations between hand space and display space was eliminated by superimposing the display image directly on the work space for the hand. As explained below, this allowed comparison of virtual pointing with two-dimensional visual information provided entirely by the computer, to physical pointing, where the hand and targets were viewed directly.

Method

Subjects

Six volunteers from the local community were paid to participate in the experiment. All had normal or corrected vision, had experience with computer pointing devices, and preferred to use the right hand in this context.

Apparatus

The subject was comfortably seated in front of a table, viewing the work surface through a half-silvered mirror as shown in Figure 14. As in Experiment 1, an OPTOTRAK motion analysis system recorded the three-dimensional position of infrared markers (IREDs) placed on the index finger, hand and arm, and stored the results in data files for further analysis. Coordinates from the OPTOTRAK were also used in real time by a Silicon Graphics Indigo Extreme workstation to update a 19" RGB monitor. Subjects performed two versions of the pointing task:

1) Virtual pointing, where the half-silvered mirror was blocked so that only the graphics display image was visible. Targets were represented as white circles on a black background, while the two IRED markers on the index finger were used to superimpose an image of a red pointer on the actual position of the subjects' finger over the work surface.
2) Physical pointing, where the graphics display was turned off and the mirror was unblocked, so that subjects could see through the mirror to the workspace below. The work space was illuminated by a task light (fluorescent, 15 watts). White circular targets were painted on a black background to correspond to the images generated by the graphics display for the virtual task. The subject wore a fluorescent red pointer, having the same dimensions as the pointer on the display in the virtual condition, on the index finger.

Procedure

For each trial, a subject held the tip of the red pointer on a start mark. On instruction from the experimenter, the subject lifted the hand and quickly pointed to a spot on the table surface so that the tip of the red pointer ended anywhere inside the target circle. Movement speed was emphasised, with small corrections allowed if the target was missed. Four target sizes (W = 6, 12, 24, 48 mm) were combined with four movement distances (A = 37, 75, 150, 300 mm) for a total of sixteen A and W combinations.

Data for each subject were collected during a single session lasting about and one half hours. After a practice session with one of the display configurations, subjects

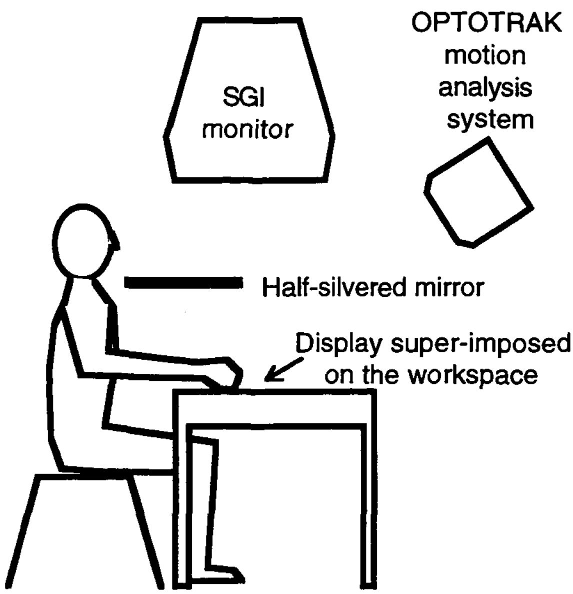

Figure 14. Layout of the experimental setup. A subject made pointing movements under two conditions: a) virtual pointing - viewing a graphics image of targets and a pointer superimposed on the workspace; and b) physical pointing - viewing a small pointer attached to the index finger and targets painted on the table surface. Kinematic data for hand position was captured by an OPTOTRAK motion analysis system.

performed twelve movements in each A and W combination, with a short break between each set of twelve trials while the target was changed. The order of presentation of A and W combinations was randomised, and the order of presentation of the virtual or physical task was counterbalanced across subjects.

Data Analysis

As in the previous experiment, three-dimensional position data for each trial were rotated and translated to a convenient coordinate system with: the x-axis representing the principal direction of movement, away from the subject; the y-axis pointing to the left of the subject, and the z-axis pointing up. Further analysis was accomplished in the same manner as in Experiment 1, except that processing was done only on hand movements, and cursor kinematics were not included. Three-dimensional data from the index finger IRED were used to calculate the following dependent measures, using the same techniques and definitions as in Experiment 1: movement time (MT), timing (TPV) and magnitude (PV) of the first velocity peak, timing (TPA) and magnitude (PA) of the initial acceleration peak, per cent of movement time after the first deceleration peak (%TAPD), and variability of the end point in the principal direction of movement, or effective target width (We).

The within subject means of the last ten good trials in each A and W combination were analysed using ANOVA with repeated measures (BMDP 8v and 2v) on a subjects (6) by display (2) by movement amplitude (4) by target width (4) design. We also performed separate multiple regressions (BMDP 2r) on means for MT as a function of A, W, and We for the two-part model of Equation 5 for both virtual, and physical pointing.

Results

Overall differences in 3-D kinematic measures for the hand due to the display are shown in Table 7. Movement times were about 80 ms slower with the virtual display than for physical pointing. Differences in MT were evident in a display by target width interaction ($F_{3,15} = 22.9$, $p < .01$). Illustrated in Figure 15, the increase in movement times as target size decreased was greater for the virtual display compared to the physical display condition.

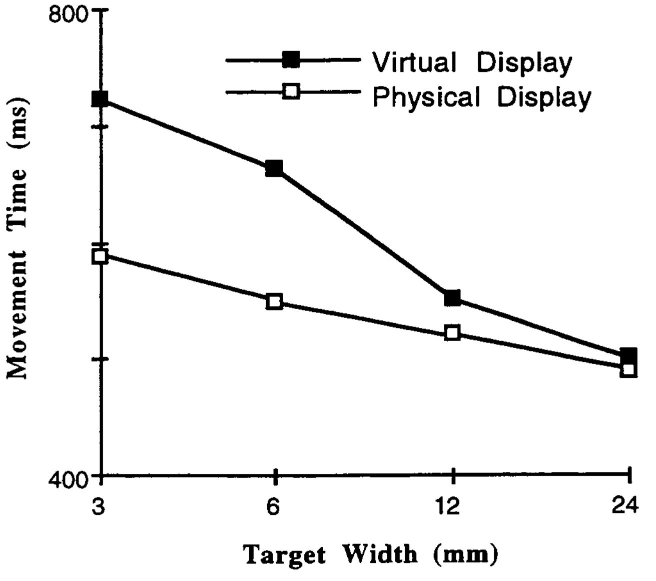

Figure 15. Mean movement times as a function of target width for virtual and physical displays, showing a display by target width interaction. Target width effects on MT are more pronounced for the virtual than the physical display.

TABLE 7.

Summary of the main effects of display condition.

Display Condition	Virtual	Physical	F1,5	p <
MT (ms)	610	538	25.35	.0040
PA (mm/s/s)	7688	8514	.75	.4248
TPA (ms)	100	97.8	1.48	.2781
PV (mm/s)	721	751	.34	.5864
TPV (ms)	192	186	.61	.4713
%TAPD	56.1	53.7	2.41	.1812
We (mm)	8.56	7.72	8.09	.0360

There were no significant differences between the physical and virtual display condition for the initial movement phase, up to the time of peak velocity. The slightly longer proportion of movement time (%TAPD) for the virtual display (56.1 versus 53.7 per cent) was not statistically significant, but revealed itself in a display by width interaction ($F_{3,15}$=4.97, $p < .02$) as shown in Figure 16. For smaller targets, the %TAPD increased more with the virtual display than with the physical display. In contrast, the %TAPD was similar across displays for the largest targets. This interaction parallels the interaction for MT, and demonstrates that the longer movement times for virtual compared to physical pointing are mainly due to extra time in deceleration to smaller targets when the virtual display was used.

The main effects for movement amplitude are shown in Table 8. As predicted by Fitts' law, movement time increased for longer distances. Differences in the velocity profile due to movement amplitude were evident up to the time of the first velocity peak. In particular, the magnitudes of the initial acceleration and velocity peaks (PA and PV) increased with longer distances, and the timing of these peaks was also delayed for longer distances in a systematic fashion. All these effects were highly significant. Although there were small differences in %TAPD for different movement amplitudes, this effect did not approach statistical significance.

Main effects for target width are shown in Table 9. As expected, movement time decreased for larger targets. Target size showed little effect on the first movement phase, although there was a slight increase in the magnitude of the first acceleration peak for larger targets. The proportion of movement time spent in deceleration increased systematically from 49.1 to 61.7 per cent as target size decreased from the largest to smallest.

We captured the effects of A and W on movement time using multiple linear regression to fit the model of Equation 5 to means of MT for each A and W combination, using units of mm and ms. For virtual pointing:

$$MT\ (ms) = 22 + 123 \log_2 A - 79 \log_2 W \qquad (R^2 = .99) \qquad (8)$$

and for physical pointing:

$$MT\ (ms) = -24 + 98 \log_2 A - 32 \log_2 W \qquad (R^2 = .96) \qquad (9)$$

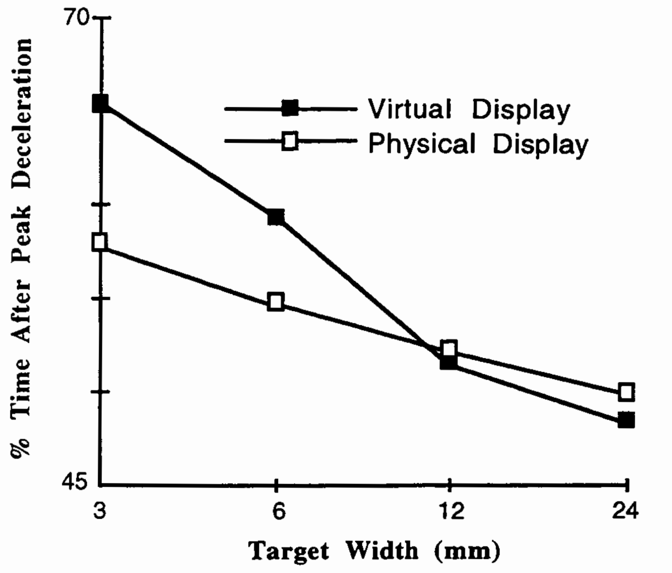

Figure 16. Mean proportion of time spent in deceleration to targets of different widths for virtual and physical display, showing a display by target width interaction. Target width effects are more pronounced for the virtual than the physical display.

TABLE 8.

Summary of the main effects of movement amplitude.

Amplitude (mm)	37.5	75	150	300	F3,15	p <
MT (ms)	420	510	613	753	229.89	.0001
PA (mm/s/s)	5148	6114	9187	11957	17.5	.0001
TPA (ms)	89.1	99.3	104	104	10.86	.0005
PV (mm/s)	324	471	837	1312	85.66	.0001
TPV (ms)	143	177	201	236	99.33	.0001
%TAPD	57.1	53.6	54.3	54.6	1.09	.3855
We (mm)	7.08	8.17	8.69	8.62	4.7	.0167

TABLE 9.

Summary of the main effects of target width.

Width (mm)	3	6	12	24	F3,15	p <
MT (ms)	657	607	536	496	44.30	.0001
PA (mm/s/s)	7487	8041	8672	8806	5.59	.0089
TPA (ms)	97.1	99.5	98.7	101	1.28	.3159
PV (mm/s)	719	747	712	766	.30	.8264
TPV (ms)	190	190	190	187	.76	.5355
%TAPD	61.7	57.0	51.9	49.1	22.14	.0001
We (mm)	4.33	5.92	9.09	13.2	42.38	.0001

Note that in both conditions MT is more sensitive to changes in A (virtual 123 ms/bit, physical 98 ms/bit) than in W (virtual 79 ms/bit, physical 32 ms/bit), and this asymmetry is even more marked for the physical display than for virtual pointing. Contrasting virtual and physical pointing, there is a greater difference for the W term (79 versus 32 ms/bit) than for changes in A (123 versus 98 ms/bit). One explanation for this phenomenon is that subjects were not taking advantage of the relaxed accuracy constraint to improve their speed in reaching larger targets, so that effective target width was not varying as much as the actual width presented during trials. Modelling MT with effective target width (We) rather than actual target width (W) revealed the same pattern of coefficients as Equations 8 and 9. For virtual pointing (units are mm and ms):

$$MT\ (ms) = 30 + 132 \log_2 A - 89 \log_2 We \quad (R^2 = .93) \qquad (10)$$

and for physical pointing:

$$MT\ (ms) = -28 + 101 \log_2 A - 41 \log_2 We \quad (R^2 = .95) \qquad (11)$$

The similar pattern of coefficients when modelling with We indicates two things: first, that the subjects were in fact taking advantage of larger targets to increase the movement speed, and second, that Equations 8 and 9 effectively describe subjects' performance in the task.

To visualise the two-part models of Equations 8 and 9, in Figure 17 we have plotted actual and predicted movement times against ID for virtual and physical pointing. In these figures:

a) Movement times are plotted as black squares. Differences for various A and W combinations giving rise to the same ID are evident by the vertical spread of the data points.

b) The dashed lines represent predicted movement times when target width is held constant, and ID is changed by varying A.

c) The solid lines represent predicted movement times when amplitude is held constant, and ID is changed by varying W.

Looking at Figure 17, two features are evident. First, changing A (dashed lines) gives rise to a steeper slope than changing W (solid lines) for both display conditions. Second, the

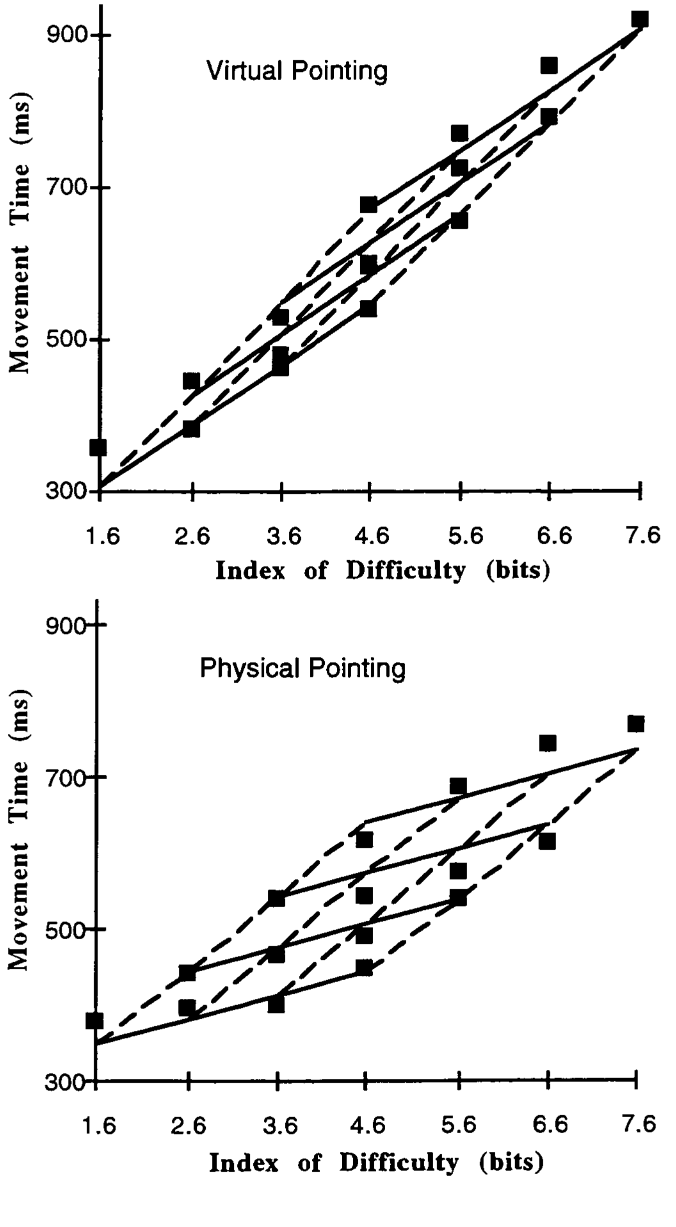

Figure 17. Illustration of the two-part models from Equations 8 (above) and 9 (below). Squares are actual movement times as a function of ID for physical pointing. Solid lines predict changes in movement time as a function of target width (W). Dashed lines predict changes in movement times as a function of movement amplitude (A). The shape of the parallelogram reveals that the asymmetry in sensitivity to changes in A and W is more marked for physical (below) than for virtual pointing (above).

difference in slopes for A (dashed lines) and W (solid lines) is more marked for physical pointing than for virtual pointing.

Discussion

The results suggest that movement planning is similar for both virtual and physical pointing, evidenced by the lack of differences in kinematic features of the initial phase of the movement. The difference between the virtual and physical display is apparent only in the second movement phase, where visual control of deceleration to the smaller targets in the virtual task took more time than in the physical task. This interpretation is also supported by contrasting the models of Equations 8 and 9: The coefficients (b_1) for the A term are of similar magnitude, but there is a marked difference in the coefficients for the W term (b_2), making MT much more sensitive to the increased accuracy constraint of a smaller target in virtual pointing.

What differences between virtual and physical display conditions would account for this pattern of results? A key feature of the virtual display in this experiment is that it is flat — it presents no visual information to show the height of the finger above the table top, even though this dimension is used as part of the deceleration strategy to contact the table surface at the target location. This lack of the third dimension could account for the difficulty with visual control during final positioning in virtual pointing. Second, the virtual display system operates with 25 ms lag, whereas with physical pointing there is presumably zero lag. This difference, though small, may increase the difficulty of controlling deceleration to smaller targets. The effects of lag in a positioning task have been modelled in terms of Fitts' law (Ware and Balakrishnan, 1994), and have been shown to increase movement time for higher indices of difficulty which would result from smaller target widths. Only further study using a more realistic 3-D virtual environment, or by systematically varying lag can clarify this issue.

A surprising result of this experiment was that, even with physical pointing, changes in the size and shape of the velocity profile did not cancel when A and W were increased in proportion, as would be predicted by Fitts' Law. The same asymmetry was also observed in Experiment 1. The spread of movement times for similar IDs in physical pointing was even greater than either virtual pointing, or our previous results using the standard HCI configuration in Experiment 1. Thus, the fact that movement times increase as A and W are scaled up in size seems to be a feature of the pointing task per se, not a result of limitations in motor control processes associated with using a computer display in the HCI setup.

Although the two-part model from Welford (Equation 5) is difficult to interpret in terms of an information-theoretic account of movement planning and control processes, it does capture the large and significant differences in movement time due to scaling. We suggest that the two-part models captured in Equations 8 and 9 promise some utility from an engineering perspective. First, as a predictive model of performance, it takes in to account the size or scale of movements appropriate for a particular interface. Second, as a diagnostic, it can serve to identify the type of constraint leading to speed-accuracy tradeoffs in a variety of positioning tasks.

In a graphical interface design, there is often some leeway in the choice of size for a representation. For example, a set of buttons can be sized and spaced in various ways on the screen. There are obvious limitations on the lower limit of size, such as visual acuity and pointing device and display resolution, as well as on the upper limit, such as the footprint for working with a pointing device, but within these bounds the designer has to make an informed choice. One way to understand how the results of the present experiment can be brought to bear on this issue is illustrated in Figure 18, which represents a parameter space for the amplitude and width coefficients of the two-part model. In this figure we have plotted the data points for virtual and physical pointing from Equations 8 and 9. In addition, we have added the point for Experiment 1, showing where Equation 6, representing the standard HCI setup, falls in the parameter space.

If the data from our studies were well characterised by Fitts' law, they would fall on the line b1 = b2. In our case, all points fall below this line by different amounts. In this region, below the Fitts' law line, movement time increases as the movement distance (and target width) are scaled up in size. Tasks associated with points in this region of the parameter space (b1 > b2) will be performed faster when hand movements are scaled down

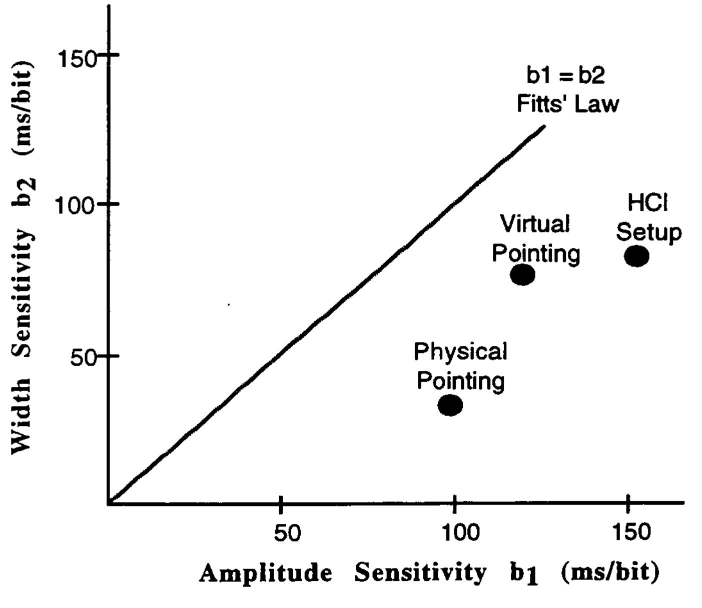

Figure 18. Parameter space for the model of Equation 5, showing the data points for virtual and physical pointing, and also for pointing using the HCI setup in Experiment 1.

in size — smaller is better. A glance at the points on the graph of Figure 18 tells us that our tasks will be performed faster if the movement distance and target size are reduced to the smallest practical value. In contrast, if we were to find a two-part model for a task which fell above the Fitts' law line (b1 < b2), this task would benefit from hand movements being scaled up — bigger would be better.

The parameter space of Figure 18 may also be useful to classify pointing devices and display systems. Approximate values for the two-part model coefficients for a particular configuration can be readily identified by testing a small number of representative users with a large and small target, and a large and small movement amplitude, a total of four A and W combinations. We suggest that data points that fall higher on the graph (b2 is large) indicate a task or device which involves more difficulty in final positioning on the target. A data point which falls more to the right (b1 is large) indicates difficulty in planning and control of the distance covering phase of the movement with that system.

The results of this experiment invite further study in at least two ways. First, the experiment should be repeated using a 3-D stereographic display superimposed on the workspace, in order to see if the differences between the virtual and physical task we reported are due to lack of visual information about the height of the hand above the workspace in our 2-D virtual environment. Second, it would be useful to revisit data reported both in the HCI and motor control literature on discrete and repetitive pointing, in order to model and plot additional points in the parameter space shown in Figure 18. With some additional points to support our results, the parameter space approach could be developed into a useful technique for characterising and quantifying human performance for different combinations of positioning task and input device.

Experiment 3: Response to Perturbation of Target Direction

In visually guided pointing, one mechanism by which endpoint accuracy is achieved can be attributed to visual reafference of the moving limb, whether viewed directly in the work space or inferred by viewing its representation as a cursor on a computer display, and its comparison to a goal position. The minimum time necessary to use vision in this way to correct for initial errors in rapid aimed movements, when the hand is viewed directly, has been estimated at approximately 200 ms (Keele & Posner, 1968; Elliott & Allard, 1985). When a cursor is being controlled on a graphics display, this minimum time may be as great as 500 ms (Smith, Smith, Stanley, & Harley, 1956) depending on the spatial relationship between the directions of hand movements and directions of corresponding effects on the display.

More recent studies, in which subjects aimed or pointed without vision of the hand and arm, have used detailed analyses of movement kinematics to suggest involvement of another, more direct mechanism which continuously corrects the trajectory of a moving limb (Goodale, Pelisson, & Prablanc, 1986; Pelisson, Prablanc, Goodale, & Jeannerod, 1986; Zelaznik, Hawkins & Kisselburgh, 1983). This fast sensorimotor control loop relates visual target position directly to the felt position of the hand, resulting in very short latency (<100 ms) corrections to the movement trajectory. This has been shown to occur when the position of a target is perturbed at the time of peak velocity of saccade, or near the time that movement is initiated (Komilis, Pelisson, & Prablanc,1993). In some cases, smooth corrections to the direction of movement are observed. In other cases, reorganisation of the movement is indicated by changes in the timing of characteristic peaks in acceleration and velocity that occur early in the movement.

Jeannerod (1991) proposed that this fast form of sensorimotor control is mediated by neural maps that encode visuo-spatial information in the form of proprioceptive information about corresponding limb position and are subsequently used for motor control by the central nervous system. There is strong evidence that the act of looking at a target location provides accurate information for positioning the limb. First, eye, head and hand movements form a synergy for reaching and pointing (Carnahan & Marteniuk, 1991).

Second, subjects are still reasonably accurate in pointing to a location without vision of the hand and arm, and even when the target is not visible as long as they keep their gaze fixed on the goal position (Prablanc et al., 1986, Pelisson et al., 1986). Third, even when the target position is well learned through proprioception without any involvement of vision (Hocherman, 1993), positioning movements are relatively slow, inaccurate, and demonstrate iterative rather than continuous corrections to the end point.

There seems little doubt that using a pointing device such as a mouse or graphics tablet to position a cursor on a graphics display is more difficult than pointing with the hand directly to a target in the work space. In motor control studies, accurate movements are often accomplished in as little as 200 to 300 ms, whereas in HCI studies there are many examples of 1 to 2 second movement times for a task with a fairly low index of difficulty (see S. MacKenzie, 1992 for a review). As discussed in the introduction to this thesis, one major difference between interacting with objects in the physical world, and interacting with a computer display using an input device can be understood in terms of the spatial relationship of hand and display spaces. To briefly review this issue, here we define hand space as the work space in which the arm and hand make goal-directed movements, and display space as a region (not necessarily in the work space) where the visual representation of objects and the moving limb is viewed. Normally, when pointing at physical targets viewed directly in the work space, hand space and display space are exactly superimposed. In a typical HCI setup, hand space (where the pointing device is operated) is on a horizontal surface in front of and to the right of a user, whereas display space is usually directly in front of their body in a near vertical plane. Predicting how hand movements will affect cursor position in a HCI setup, with a mouse or graphics tablet, requires that the user learn or adapt to a particular spatial transformation (translation, rotation, and scaling) from hand space to display space.

Can difficulties with pointing accurately and quickly on a computer display be understood in terms of the interaction of the sensorimotor mechanisms for movement control, and the difficulty of processing spatial transformations to relate movements in hand space to their effects in display space? In a typical HCI setup, the felt position of the hand in the workspace no longer corresponds directly to the seen position of a target on the display, suggesting that continuous fast correction to errors in the movement trajectory might not be possible. The mechanism proposed by Jeannerod is based on a coupling between gaze and control of limb movements to the location in space where gaze is fixed. In an HCI context, the goal position for limb movements in the work space bears no direct relation to gaze, which is fixed to a target on the display. At the very least, it seems that some additional time would be required to process the spatial transformations necessary to control goal-directed hand movements under these conditions.

In order to examine this question, the present experiment was designed to contrast performance when subjects made fast pointing movements while viewing a computer display under two conditions: first, when display and hand space were superimposed; and second, when hand space and display space were spatially separate, in the standard HCI setup for interacting with a computer display using a mouse or graphics tablet. With both displays, movements were recorded and kinematic features analysed when subjects pointed with the index finger to stable targets, and also to targets which were initially presented in one position and then perturbed to a new position at movement onset. This approach is similar to the double step paradigm used in studies by Komilis, Pelisson, & Prablanc (1993) and others.

The hypothesis was that the response to perturbations would be different depending on the display condition. With the superimposed display there would be fast, continuous corrections early in the movement trajectory to compensate for the perturbation, indicated either by a change in the timing or magnitude of the initial velocity and acceleration peaks (as in Paulignan, Jeannerod, MacKenzie, & Marteniuk, 1991b), or by a change in the direction of movement in the initial phase. With the HCI setup, because of the spatial transformations required, there would be no early continuous corrections. The response to perturbations would take the form of discrete corrections resulting in a lengthening of the deceleration phase of the movement.

Method

Subjects

Eight subjects with normal or corrected vision, ranging in age from 19 to 41 years, were paid for their participation in the experiment. All had experience using computer systems with pointing devices, and preferred to use their right hand in this context. Subjects were naive with regard to the purpose of the experiment.

Apparatus

During the experiment, the subject was seated comfortably in front of a table, the surface of which provided a workspace for the hand. Two display configurations were used, as shown in Figure 19:

(A) In the standard human-computer interaction (HCI) setup from Experiment 1, the computer monitor was placed directly in front of the subject at a comfortable viewing distance. The work space for the hand was offset to the right, as is commonly used with pointing devices such as a mouse or graphics tablet.

(B) In the superimposed setup from Experiment 2, the workspace was viewed through a half-silvered mirror so that the image of the monitor, mounted in a frame above the work space, appeared on the table surface. In this case, the work space for the hand, and also the display space lay centred directly in front of the subject. During the experiment, a black cover on the underside of the mirror prevented the subject from seeing their hand and arm directly. The cover was removed for the initial calibration step, described below.

In both display conditions the subject viewed the monitor image, and could not see the hand or arm while pointing to targets using the tip of the index finger. As in the previous experiments, hand movements were monitored by an OPTOTRAK motion analysis system which resolved the three-dimensional position of infrared emitting diode markers (IREDs) placed on the subject's hand and index finger. Figure 20 illustrates the computer display, showing white circular targets appearing on a black background, and a red pointer whose position and orientation were controlled by the IREDs on the subject's index finger.

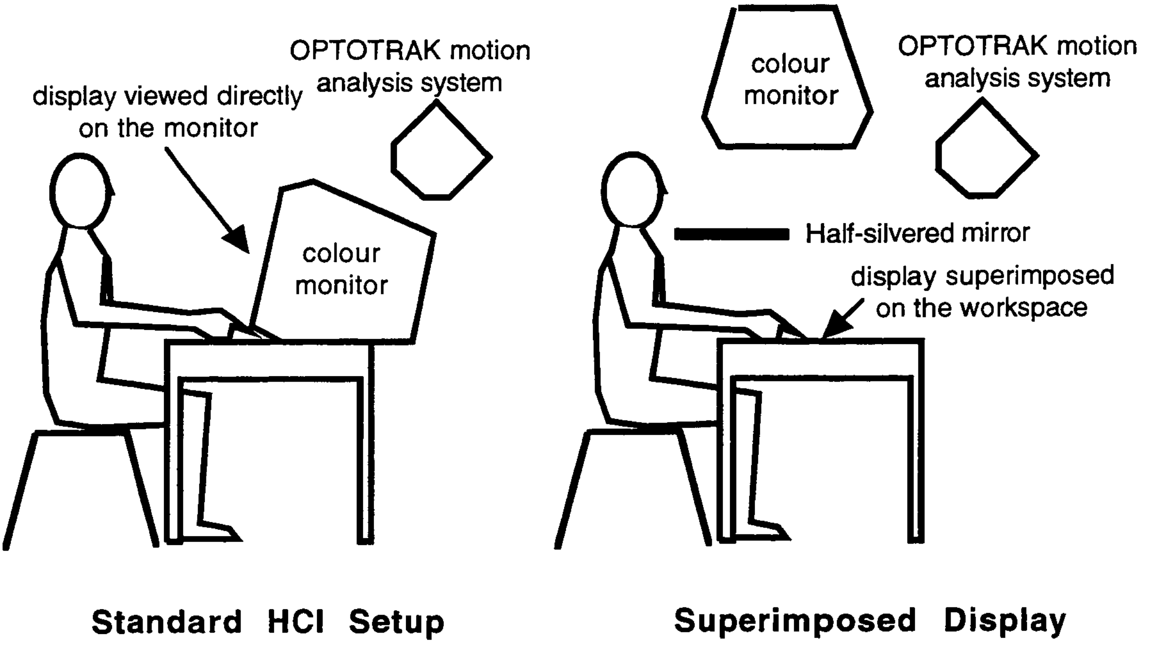

Figure 19. Computer displays used in the experiment. In the standard HCI setup, used in Experiment 1, a pointer controlled by the index finger, and experimenter-defined targets appear on a monitor directly in front of the user. Hand movements on the table in front of and to the right of the user are tracked by the OPTOTRAK motion analysis system. In the superimposed display (shown on the right), the user looks down at the workspace through a mirror to view the same image of targets and a pointer. The image of the pointer appears to match exactly the position and orientation of the index finger.

Procedure

Data for each subject were collected during two separate one hour sessions on different days. At the start of a session in which the subject used either the HCI or superimposed display, the relationship between the display and work space for the hand was calibrated. For the HCI setup, a work space with the same dimensions as the monitor screen (350 by 280 mm) was defined on the table surface, and measured by the OPTOTRAK. These data were used to map the position of the index finger to the pointer on the monitor screen during the remainder of the current session. For the superimposed display, IREDs viewed through a half-silvered mirror were aligned with reference marks from the corners of the projected display to define a work space for the hand (350 by 280 mm). These data enabled the computer to position the red pointer, viewed through the mirror, exactly over the position of the subject's index finger (Figure 20).

A trial was initiated when the subject held the tip of the red pointer on a start mark for a sufficient period of time. This period of time, the delay after which a trial was started, varied (567, 767, 967, or 1167 ms) from trial to trial. Holding the tip of the red pointer still in the very small start mark (3 mm diameter) required the subject to focus, and fix gaze on the start position. The difficulty in holding still, and the unpredictable delay was designed to prevent subjects from starting the pointing movement before the target was presented.

Once the trial began, the computer presented a white circular target 24 mm in diameter in one of five positions (varying at angles of 10°), at a distance of 220 mm from the start position as shown in Figure 21. A subject was instructed to wait for a target, and then to lift the hand and point to a position on the table so as to position the tip of the red pointer anywhere inside the target. It was emphasised that this was not a test for reaction time, but they were to make their movement as fast as possible while still landing in the target. Subjects were instructed to make a small correction to reach the target if they missed.

In nonperturbed trials, target position remained constant. In some trials (eight per cent), the target was shifted to another position as soon as the hand speed exceeded 20 mm/s. The end of a movement was detected by the computer when the tip of the red

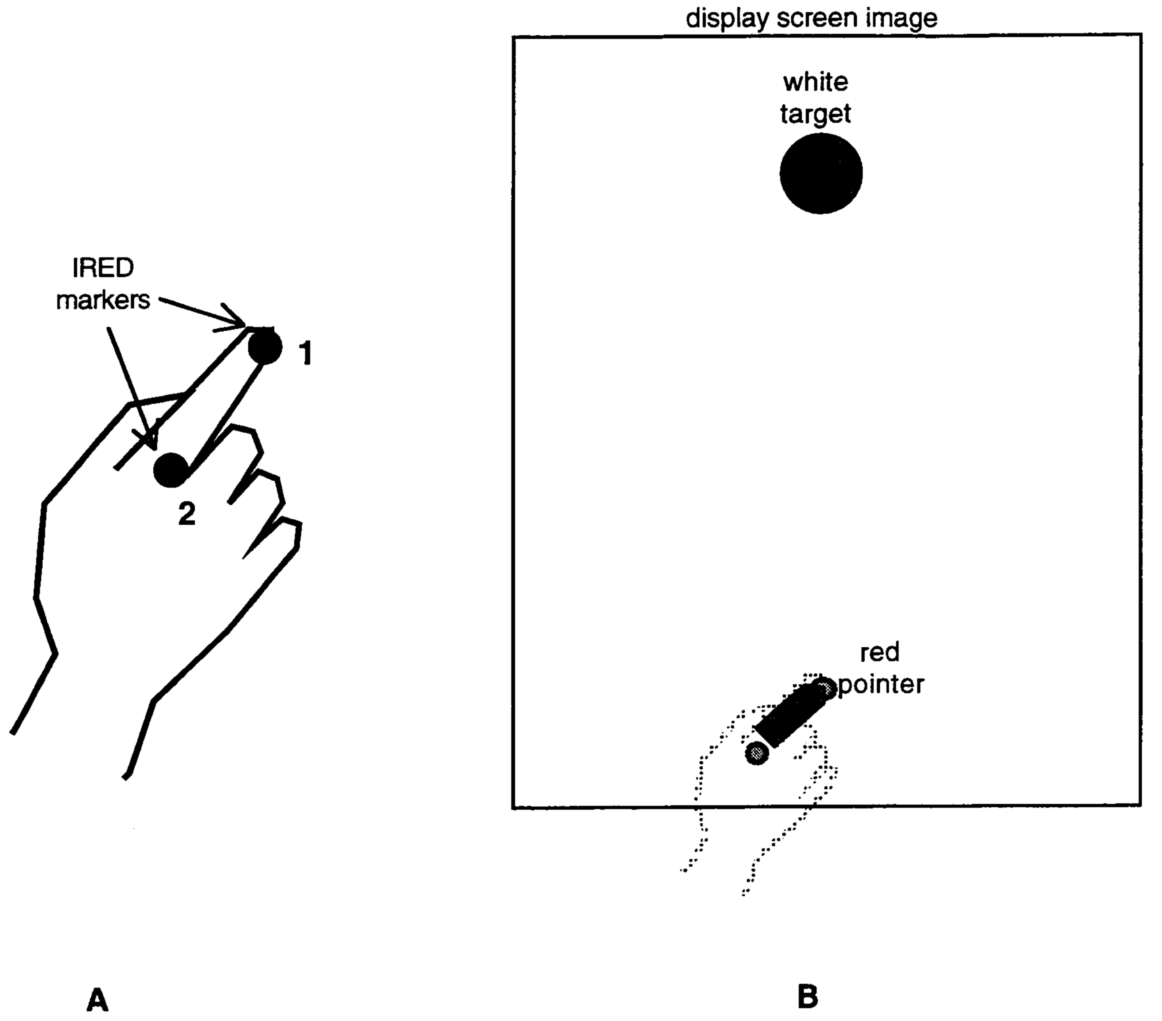

Figure 20. (A) Placement of the infrared (IRED) markers on the index finger. The positions of the IREDs are monitored by the OPTOTRAK motion analysis system. (B) The positions of the IREDs over the work space are mapped to screen coordinates, and are used in real time to control the red pointer and targets on the computer monitor. With the superimposed display, the red pointer appears over the actual position of the subject's index finger on the table top.

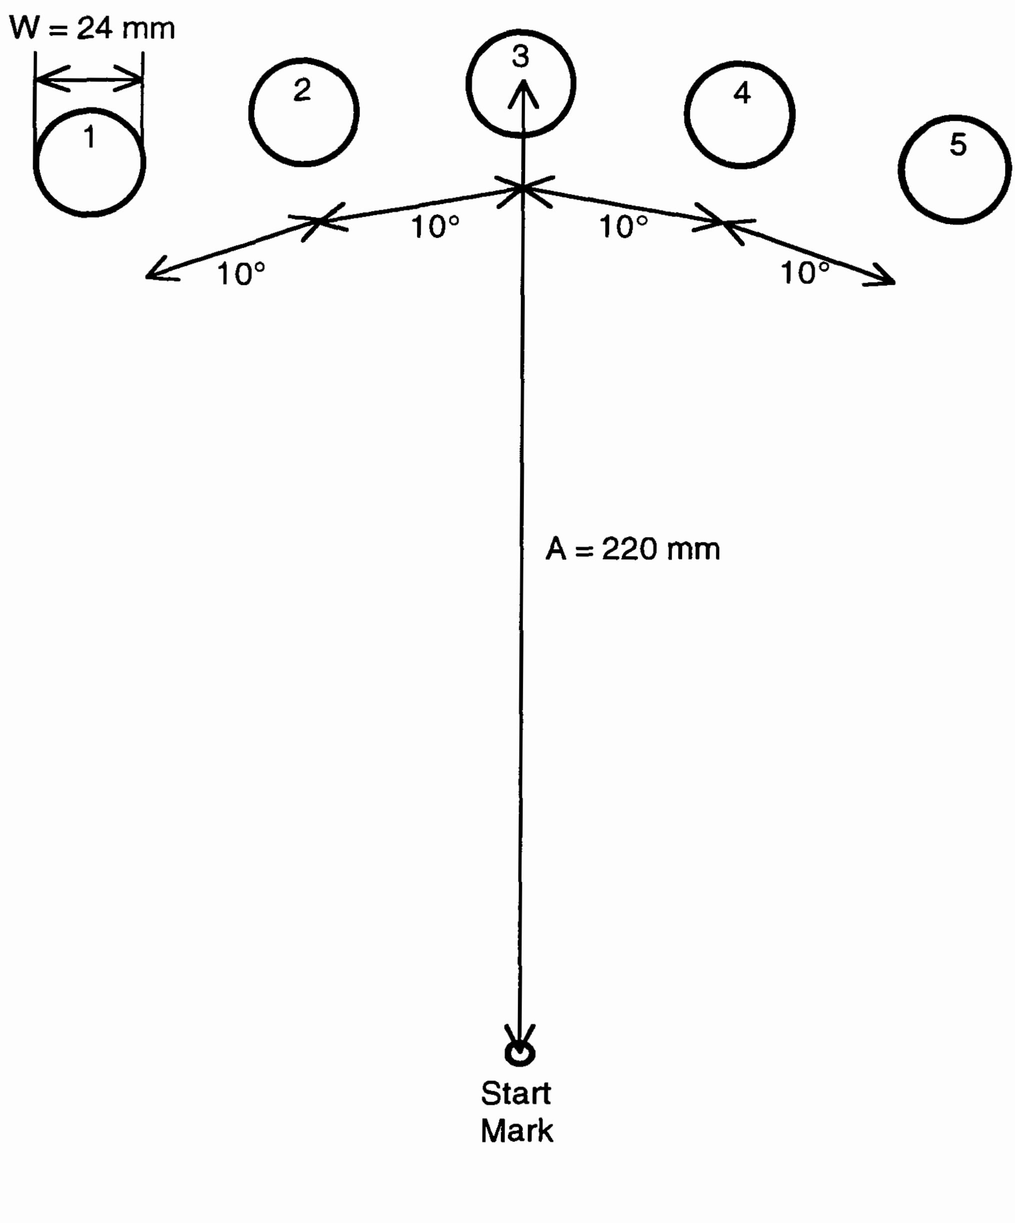

Figure 21. Positions and size of targets displayed in the experiment. The start mark, and position 3 fall along the subject's body midline in both the HCI and superimposed display conditions. Although there were an equal number of trials to all five targets in the experiment and subsequent analysis, for clarity we have presented results considering only three of the targets: left (2), centre (3), and right (4).

pointer remained on or in the target circle for 440 ms, during which time the hand speed also remained below 20 mm/s. If a subject failed to reach the target in the maximum allotted time of 5 seconds, the trial was ended and discarded. At the end of each trial, the target was erased from the screen, the start mark displayed, and the computer waited for the subject to initiate the next trial.

After the display was calibrated for a session, the subject performed three blocks of trials separated by a short break. The first block of 50 practice trials (all nonperturbed) contained ten presentations of each of the five target locations in completely random order. Two subsequent blocks (218 randomised trials per block) were used for analysis, and contained 40 trials to each of the five target positions as well as 18 trials during which the initial target position was perturbed. In 6 of the perturbed trials, the initial centre target position was shifted 10° to the left, and in another 6 it was shifted 10° to the right at movement onset (see Figure 22). The remaining 6 perturbed trials to initial positions other than the centre were included in order to prevent a subject from realising that only the centre target was likely to be shifted, and performing more cautiously in this condition.

Four of the eight subjects performed the session using the HCI display first; the remaining four used the superimposed display in their first session. In all blocks, the selection of starting delay for trials, the order of target positions, and occurrence of perturbed trials was completely randomised.

Data analysis

The same data analysis as in previous experiments was used to produce the following dependent measures: movement time (MT), timing (TPV) and magnitude (PV) of the first velocity peak, timing (TPA) and magnitude (PA) of the initial acceleration peak, and per cent of movement time after the first deceleration peak (%TAPD). In addition, two-dimensional directional measures were calculated by projecting the spatial path on the horizontal plane of the workspace (x and y axes) to determine the tangent to the path of the movement at the time of the first acceleration peak (DIRPA), and at the time of the first velocity peak (DIRPV). Directions were in units of degrees from the body midline (x-axis), with movements to the left resulting in a positive value, and movements to the right a negative value.

For each block of trials, within-subject means for all dependent measures were calculated for each target position (using the 40 trials to each position), and for perturbed trials starting toward the centre target but perturbed to the left (6 trials), and perturbed trials to the right (6 trials). The remaining 6 perturbed trials were not analysed. We performed two separate analyses for each dependent measure in order to determine the effects of the perturbation on movement kinematics, and to distinguish differences which normally occur when pointing to different directions in the work space from perturbation effects. First, we examined the effects of movement direction for nonperturbed trials only (direction analysis) by performing an ANOVA with repeated measures on a subjects (8) by block (first, last) by direction (1 to 5) design using within-subject means. Second, we contrasted nonperturbed trials towards the centre target with perturbed trials (perturbation analysis) by performing a separate ANOVA with repeated measures on a subjects (8) by block (first, last) by direction (nonperturbed, perturbed left, perturbed right) design.

Results

In the following sections, we first present the direction analysis of nonperturbed trials in order to highlight systematic differences in kinematics which arise naturally as a function of movement direction. Although the ANOVA compared trials to targets in all five directions, for clarity we present here results for the three middle target positions (left, centre, and right) which are directly comparable to the final target positions in the perturbation analysis. Second, we present the perturbation analysis, highlighting the differences between nonperturbed and perturbed trials to the centre target position, and demonstrating that responses to the perturbations differ from directional effects.

The analyses are discussed in terms of factors, rather than by examining each dependent measure separately, since we realise that measures such as movement time (MT) and time to peak velocity (TPV) can be highly correlated. Effects due to trial block are examined to see if subjects adapted their behaviour over the course of the experiment, either due to the presence of perturbed trials, or because they had not achieved stable performance as a result of the practice session. Effects of the display configuration are considered separately in order to highlight overall differences in performance between the HCI setup and the superimposed display.

Target Direction

The main effects for trial block on nonperturbed trials to the left, centre, and right targets are summarised in Table 10. None of the differences due to block were large or statistically significant, suggesting that subjects' performance had stabilised during the practice session, and they exhibited a constant level of performance throughout the two blocks of trials used for analysis.

Table 11 shows the main effects for display on nonperturbed trials. Consistent with results from Experiments 1 and 2, movements using the superimposed display were faster than with the HCI setup by approximately 92 ms. The increase can be attributed to a faster and shorter initial phase of the movement. Both the initial acceleration peak, PA and velocity peak, PV exhibited larger values for the superimposed display than the HCI setup. There was no main effect on the time to peak acceleration (TPA), but Figure 22 illustrates the significant ($F_{4,28} = 14.7$, $p < .001$) interaction between display and direction, where TPA occurred sooner for movements to the right with the superimposed display than the HCI setup, but was similar for the leftmost target in both the HCI and superimposed display conditions.

The initial velocity peak occurred somewhat sooner (145 compared to 172 ms) for the superimposed display than with the HCI setup. There were also significant interactions, shown in Figures 24 and 25, between display and direction on both the magnitude ($F_{4,28} = 12.1$, $p < .001$) and timing ($F_{4,28} = 16.91$, $p < .001$) of the initial velocity peak. PV increased systematically from left to right for the superimposed display, but remained relatively constant across directions for the HCI setup. In addition, TPV decreased systematically from left to right for the superimposed display, while there was little difference in the timing of the velocity peak for the HCI setup. The proportion of movement time decelerating to the target was slightly less for the superimposed display (51.9 per cent) compared to the HCI setup (53.7 per cent), although this effect was not

TABLE 10.

Summary of effects of trial block on nonperturbed trials to left, centre, and right targets

Trial Block	First	Last	$F_{1,7}$	p <
MT (ms)	540	530	2.19	.1826
PA (mm/s/s)	12880	12607	.74	.4183
TPA (ms)	88.0	91.7	2.12	.1866
DIRPA (degrees)	-.63°	+.50°	2.87	.1342
PV (mm/s/s)	1130	1114	.98	.3563
TPV (ms)	158	160	.31	.5972
DIRPV (degrees)	+.11°	+.17°	.07	.8026
%TAPD	53.3	51.7	2.38	.1670

TABLE 11.

Summary of effects of display on nonperturbed trials to the left, centre, and right targets

Display	HCI	Superimposed	$F_{1,7}$	p <
MT (ms)	581	489	63.75	.0001
PA (mm/s/s)	11139	14349	14.73	.0064
TPA (ms)	90.8	88.1	1.00	.3499
DIRPA (degrees)	+1.3°	-1.4°	17.96	.0038
PV (mm/s/s)	1045	1198	9.95	.0160
TPV (ms)	172	145	13.43	.0080
DIRPV (degrees)	+1.1°	-.81°	4.68	.0672
%TAPD	53.7	51.9	1.46	.2658

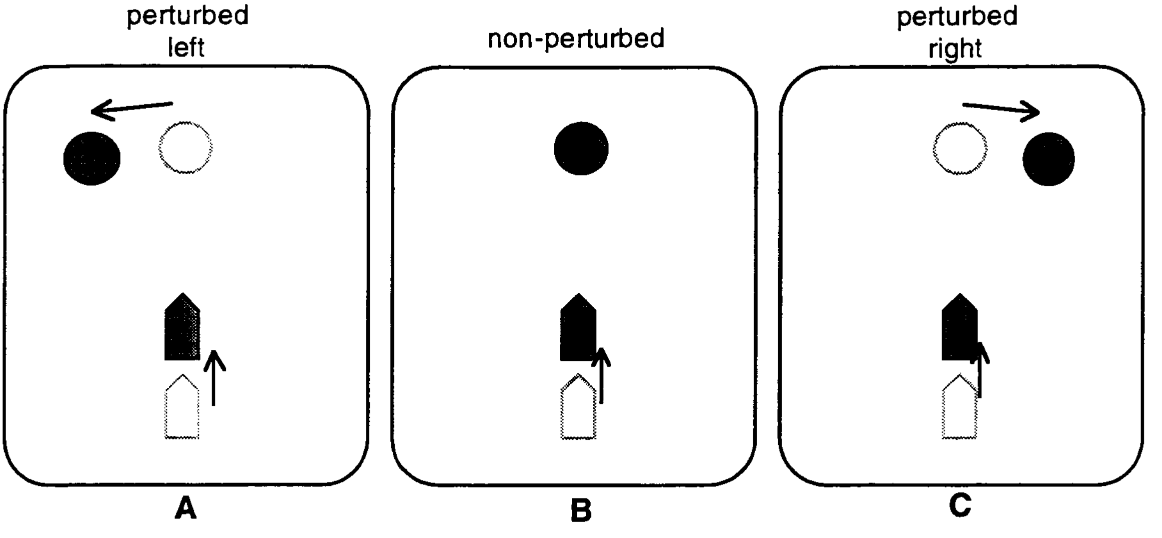

Figure 22. In non-perturbed trials (B), the target was presented in a fixed position for the duration of the trial. Pointing movements to the centre position shown here were used for comparison with perturbed trials. In perturbed trials the target was initially presented in the centre position, then shifted to the left (A) or right (C) as soon as the subject's finger moved from the start position.

statistically significant.

There were a large number of highly significant and systematic effects due to the target direction, summarised in Table 12. Differences in the directional measures at the time of peak acceleration, DIRPA and peak velocity, DIRPV showed a consistent pattern, as expected. The directions were slightly exaggerated with a range of 24 ° for DIRPA and a range of 22° for DIRPV, although the range of actual directions from the start mark between the leftmost and rightmost targets was 20°.

Movement time decreased systematically, going from left to right, by a few milliseconds. These small differences were mainly due to systematic variations in the initial phase of the movement, as both the initial acceleration and velocity peaks occurred sooner for targets to the right than to the left. In addition, the magnitudes of both peaks, PA and PV increased systematically from left to right. The general pattern of faster movements to the right than the left is mainly due to the superimposed display. This is evident in the interactions described in Figures 23, 24, and 25, where the peaks occur sooner and velocities are greater for the superimposed display, but show little variation due to direction with the HCI setup.

Response to Perturbation

The main effects of trial block (first and last) on kinematic measures for nonperturbed trials to the centre, and perturbed trials to the left and right targets are summarised in Table 13. The only significant effect was a slight decrease in movement time of 24 ms from first to last trial block. This decrease can be understood in terms of two interactions. The first interaction between block and direction ($F_{2,14} = 3.97$, $p < .05$) is displayed in Figure 26, where MT decreases significantly only for trials in which the target was perturbed to the right, but stays relatively constant for trials perturbed to the left or nonperturbed trials. The second interaction, between block and display ($F_{1,7} = 5.28$, $p < .056$) is shown in Figure 27. Here, the proportion of time spent decelerating to the target (%TAPD) was initially higher in the first block for the HCI setup, but decreased to a level similar to the superimposed display by the last block of trials.

TABLE 12.

Summary of effects of direction on nonperturbed trials to the left, centre, and right targets

Target	Left	Centre	Right	$F_{4,28}$	p <
MT (ms)	538	533	531	6.12	.0011
PA (mm/s/s)	12212	12700	13049	12.34	.0001
TPA (ms)	92.5	90.8	86.9	11.75	.0001
DIRPA (degrees)	+11°	+.55°	-13°	407.93	.0001
PV (mm/s/s)	1099	1118	1134	15.66	.0001
TPV (ms)	162	161	157	9.71	.0001
DIRPV (degrees)	+11°	+.11°	-11°	3784.98	.0001
%TAPD	51.8	52.2	53.1	1.77	.1629

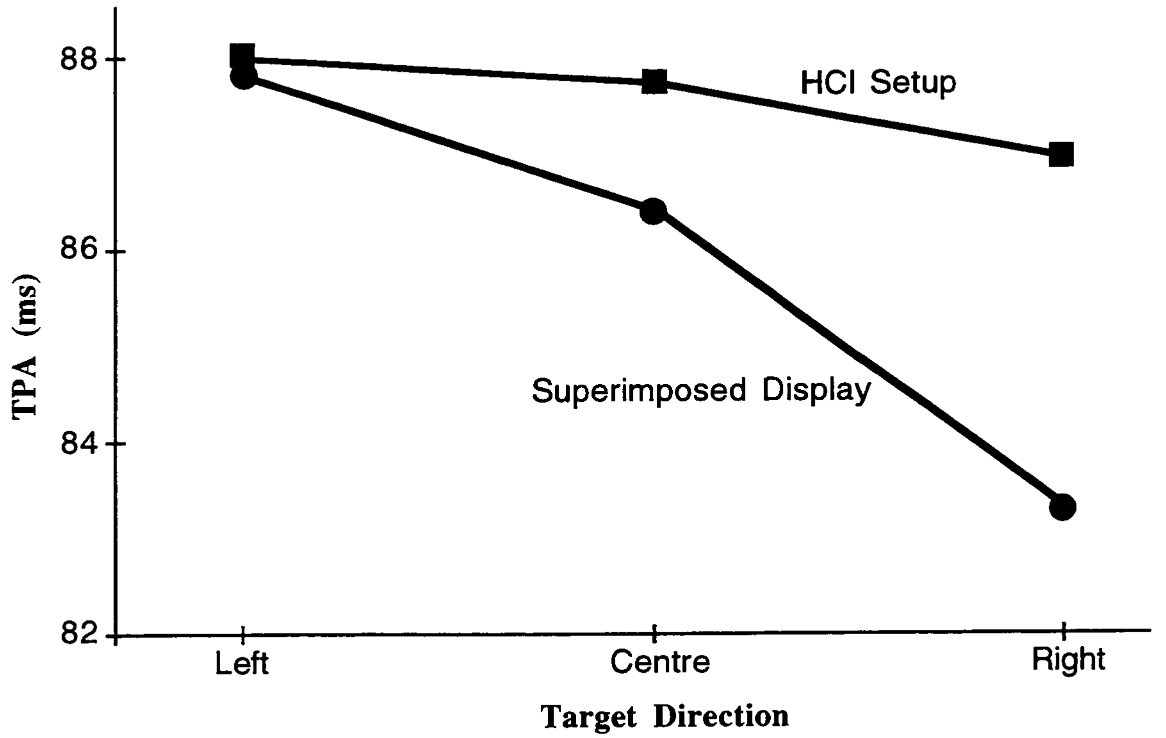

Figure 23. The time of the initial acceleration peak showed a significant interaction between display and direction for non-perturbed trials. The timing of the peak was similar for with the two displays trials to the left target, but for trials to the right the peak occurred sooner with the superimposed display than with the HCI setup.

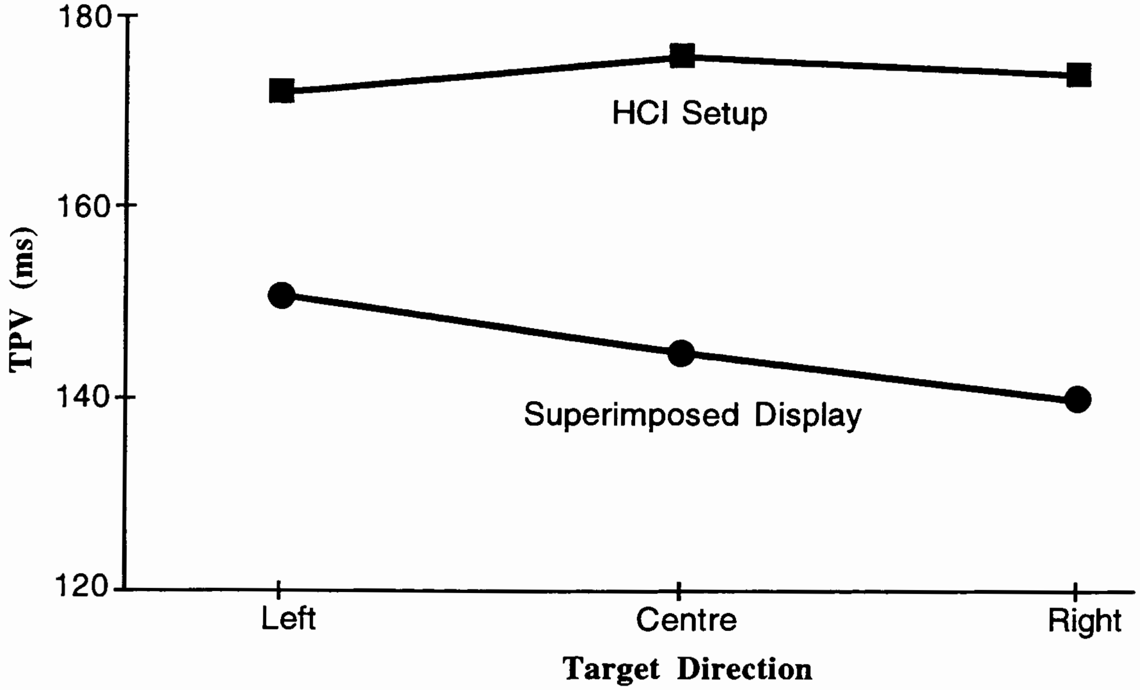

Figure 24. The time of the initial velocity peak showed a significant interaction between display and direction for non-perturbed trials. The timing of the peak decreases systematically from left to right targets with the superimposed display. In contrast, the timing of the peak was similar in all directions with the HCI setup.

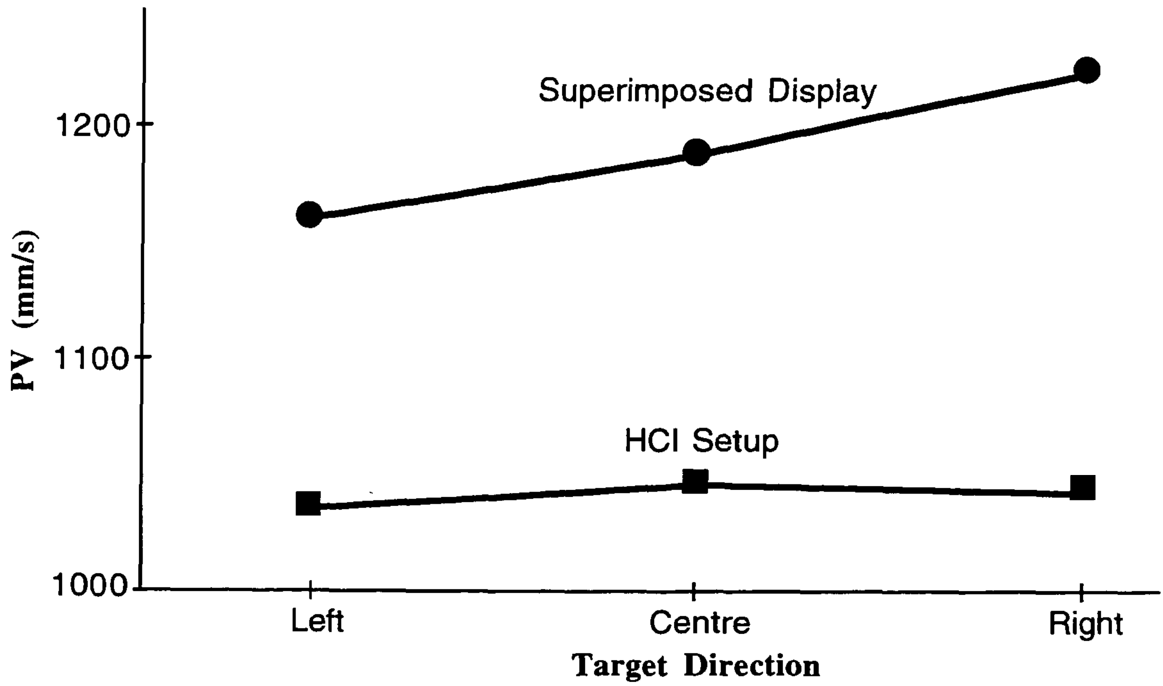

Figure 25. The magnitude of the initial velocity peak showed a significant interaction between display and direction for non-perturbed trials. The peak increased systematically from left to right targets with the superimposed display. In contrast, the magnitude of the peak was similar in all directions with the HCI setup.

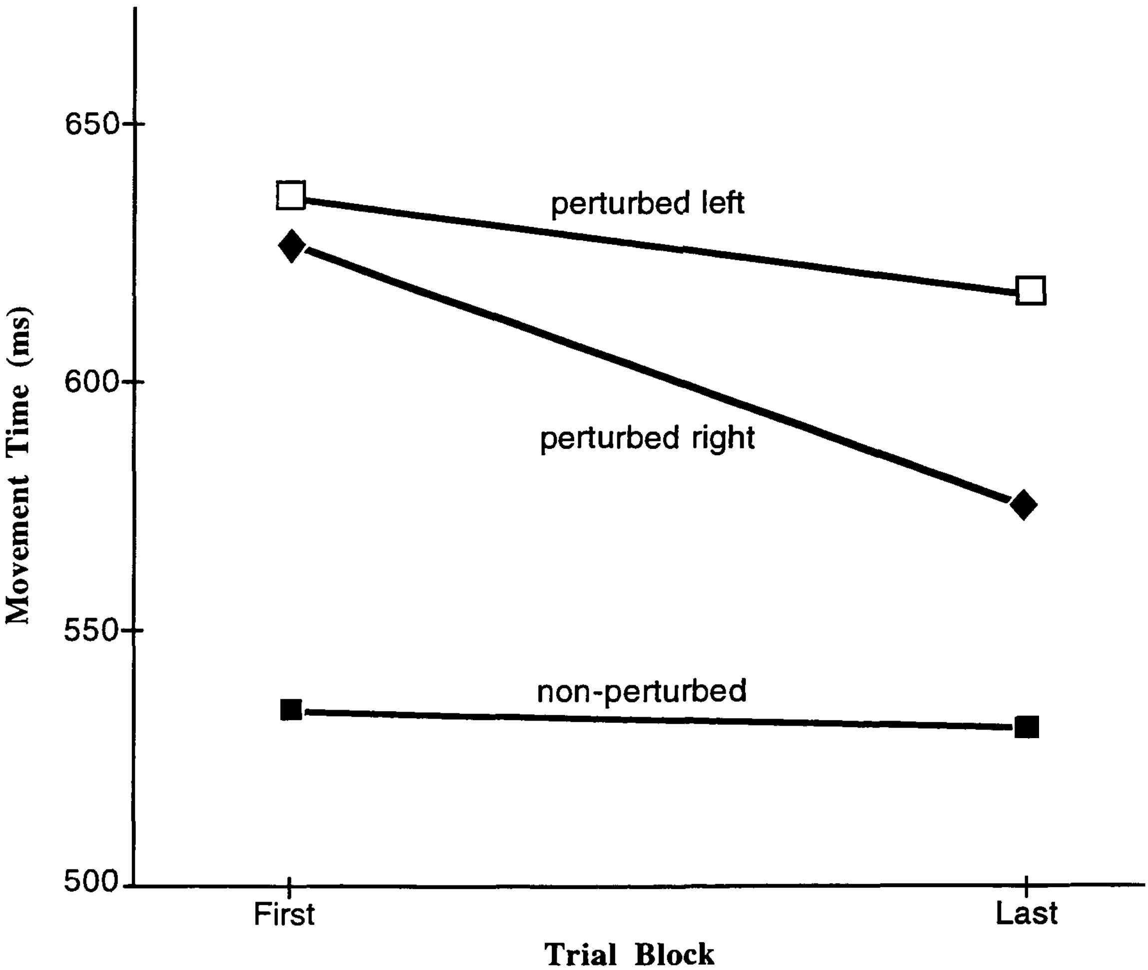

Figure 26. Movement times revealed an interaction of perturbation and trial block. MT for trials perturbed to the right decreased with practice, while the times for non-perturbed trials and those perturbed to the left showed little change between the first and last block of trials.

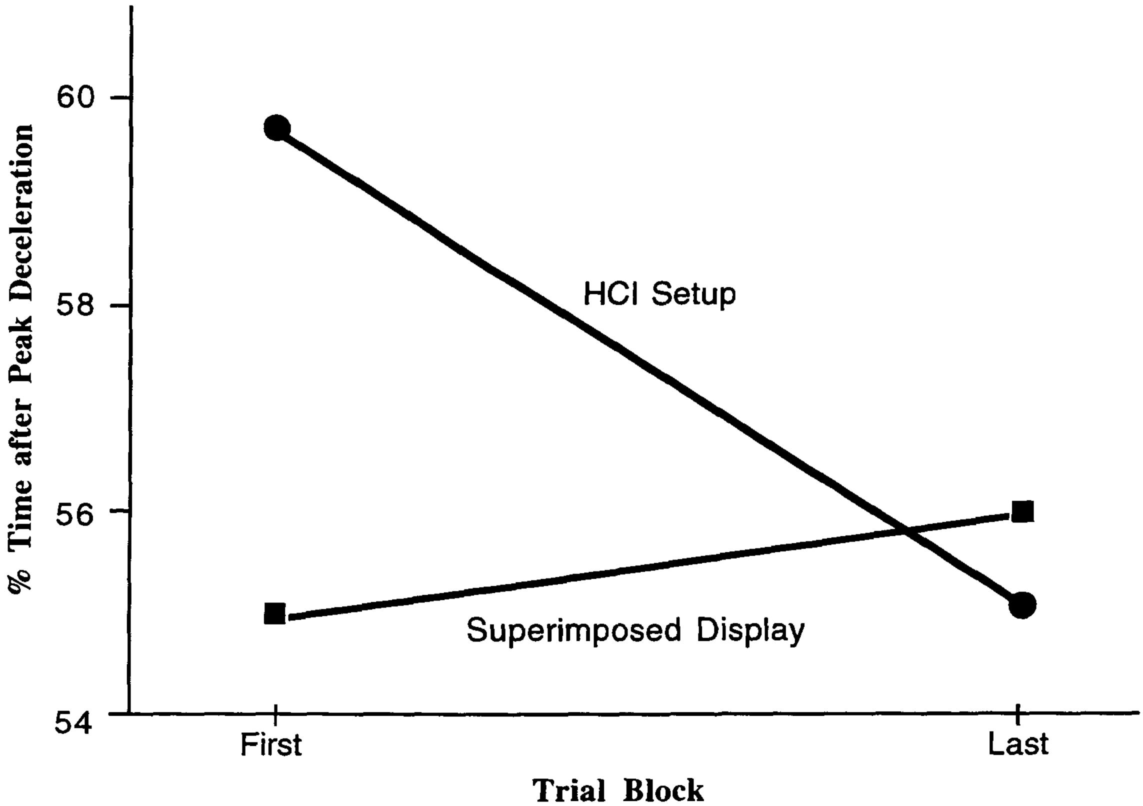

Figure 27. The interaction of display and trial block for nonperturbed and perturbed trials to the centre target. The proportion of time spent in deceleration to the target (%TAPD) decreased from the first to the last trial block in the HCI setup, but remained relatively constant for the superimposed display.

TABLE 13.

Summary of main effects of trial block on nonperturbed and perturbed trials to the centre target

Trial Block	First	Last	$F_{1,7}$	p <
MT (ms)	599	575	12.50	.0095
PA (mm/s/s)	12930	12485	1.48	.2633
TPA (ms)	88.9	90.6	.40	.5471
DIRPA (degrees)	-1.4°	-.79°	.25	.6332
PV (mm/s/s)	1130	1100	4.33	.0760
TPV (ms)	158	160	.23	.6484
DIRPV (degrees)	-.13°	-1.6°	2.12	.1889
%TAPD	57.3	55.6	3.11	.1213

Note that in the direction analysis discussed in the previous section, there were no significant effects due to trial block. Subjects exhibited stable performance over the course of the experiment for nonperturbed trials, having done a number of such trials during the practice session. This suggests that trial block effects in the perturbation analysis are due to changes in performance over time in perturbed trials, a condition in which the subjects received no practice before starting the experiment. An interpretation of this combination of effects and interactions is that subjects adapted or accommodated over the course of the experiment to the difficulties of correcting their trajectories when targets were perturbed to the right, but not when targets were perturbed to the left. The adaptation was only evident in the first block using the HCI setup; by the last block of trials, deceleration profiles became similar in both the HCI and superimposed display conditions. While statistically significant, these overall effects of trial block on MT were fairly small, on the order of 25 ms.

Table 14 summarises the main effects for display, which are concentrated on the initial movement phase. Movement time was slower by 86 ms with the HCI setup compared to the superimposed display. In the superimposed display condition, the initial acceleration and velocity peak occurred sooner, and with greater magnitudes than in the HCI setup. The small difference due to the display in %TAPD did not reach significance, and was mainly due to the block by display interaction discussed above. These effects indicate that the overall movement trajectory was slower in the HCI setup than with the superimposed display, but the difficulty of reaching targets, and hence the shape of the velocity profile, was similar with both displays.

The effects of target perturbation are summarised in Table 15. Overall movement times were significantly greater for perturbed trials (627 ms left, 601 ms right) than for nonperturbed trials (533 ms) to the centre target. The extra time was spent mainly in deceleration to targets, as evidenced by greater values for %TAPD (59.9 left, 57.2 right) in perturbed trials than in nonperturbed trials (52.2) to the centre target. In contrast to the large effect of the perturbation on the deceleration phase of the movement, there were no significant effects of the perturbation on the timing and magnitude of the initial velocity and acceleration peaks. Overall, this indicates that there was no early correction or modification

TABLE 14.

Summary of main effects of display on nonperturbed and perturbed trials to the centre target

Display	HCI	Superimposed	$F_{1,7}$	$p <$
MT (ms)	630	544	11.98	.0105
PA (mm/s/s)	10431	14482	32.96	.0007
TPA (ms)	91.4	88.1	.81	.3989
DIRPA (degrees)	-.70°	-1.5°	.19	.6750
PV (mm/s/s)	1029	1201	20.06	.0029
TPV (ms)	173	145	16.67	.0047
DIRPV (degrees)	-.85°	-.94°	.01	.9439
%TAPD	57.4	55.5	2.16	.1854

TABLE 15.

Summary of main effects of perturbations: nonperturbed trials to the centre target, perturbed trials to the right, and perturbed trials to the left

Target	Nonperturbed	Perturbed Left	Perturbed Right	$F_{2,14}$	p <
MT (ms)	533	627	601	7.25	.0069
PA (mm/s/s)	12700	12701	12719	.00	.9965
TPA (ms)	90.9	89.5	88.8	1.34	.2943
DIRPA (degrees)	-.55°	+.37°	-3.1°	4.24	.0362
PV (mm/s/s)	1118	1113	1114	.02	.8231
TPV (ms)	173	173	145	1.37	.2865
DIRPV (degrees)	-.11°	+.89°	-3.5°	4.30	.0349
%TAPD	52.2	59.9	57.2	15.23	.0003

to the scaling of the velocity profile on perturbed trials, and differences in timing between perturbed and nonperturbed trials were only apparent in the final phase of the movement. There was no display by perturbation interaction ($F_{2,14} = .10$, $p > .9$) as predicted; the extra time taken in perturbed trials was clearly a result of the increased proportion of time spent decelerating to targets in both display conditions.

There was, however, some evidence of early corrections in response to the perturbation in directional measures. Perturbations caused significant small effects on the direction of movement of the distal IRED as early as the time of the first velocity peak and even the time of the initial acceleration peak. These effects are illustrated schematically in Figure 28, using an exaggerated scale to emphasise the pattern of corrections in both display conditions.

Consider that in perturbed trials, the target position was perturbed at movement onset as detected by the computer system, but the latency for updating the graphics screen was at least 17 ms (one frame of video). This means that the target did not appear to be shifted until at least 17 ms after the start of the pointing movement. Thus, the directional effects evident at the time of peak acceleration (about 90 ms in this experiment) must have been manifested in less than 73 ms from the time of the visual cue that the target had shifted. To illustrate these effects, we selected representative trials from one subject where the measures for DIRACC and DIRVEL showed a pattern similar in magnitude to the overall effects in Table 15. The spatial paths for these trials are plotted in Figure 29 to illustrate how small differences in the direction of the movement are apparent in the initial phase. The magnitude of the corrections (range 7.4° from left to right) is significant, but smaller than the actual range of directions to the two extreme target positions (more than 20° from left to right). This indicates that directional corrections early in the movement did not compensate completely for the perturbed target position, and that additional corrections were required in the deceleration phase to reach the target in its final position. Contrary to the experimental hypothesis, there was no evidence that these fast corrections occurred only in the superimposed display condition. The display by position interactions for DIRACC ($F_{2,14} = .46$, $p > .6$) and DIRVEL ($F_{2,14} = 2.57$, $P > .1$) did not approach significance,

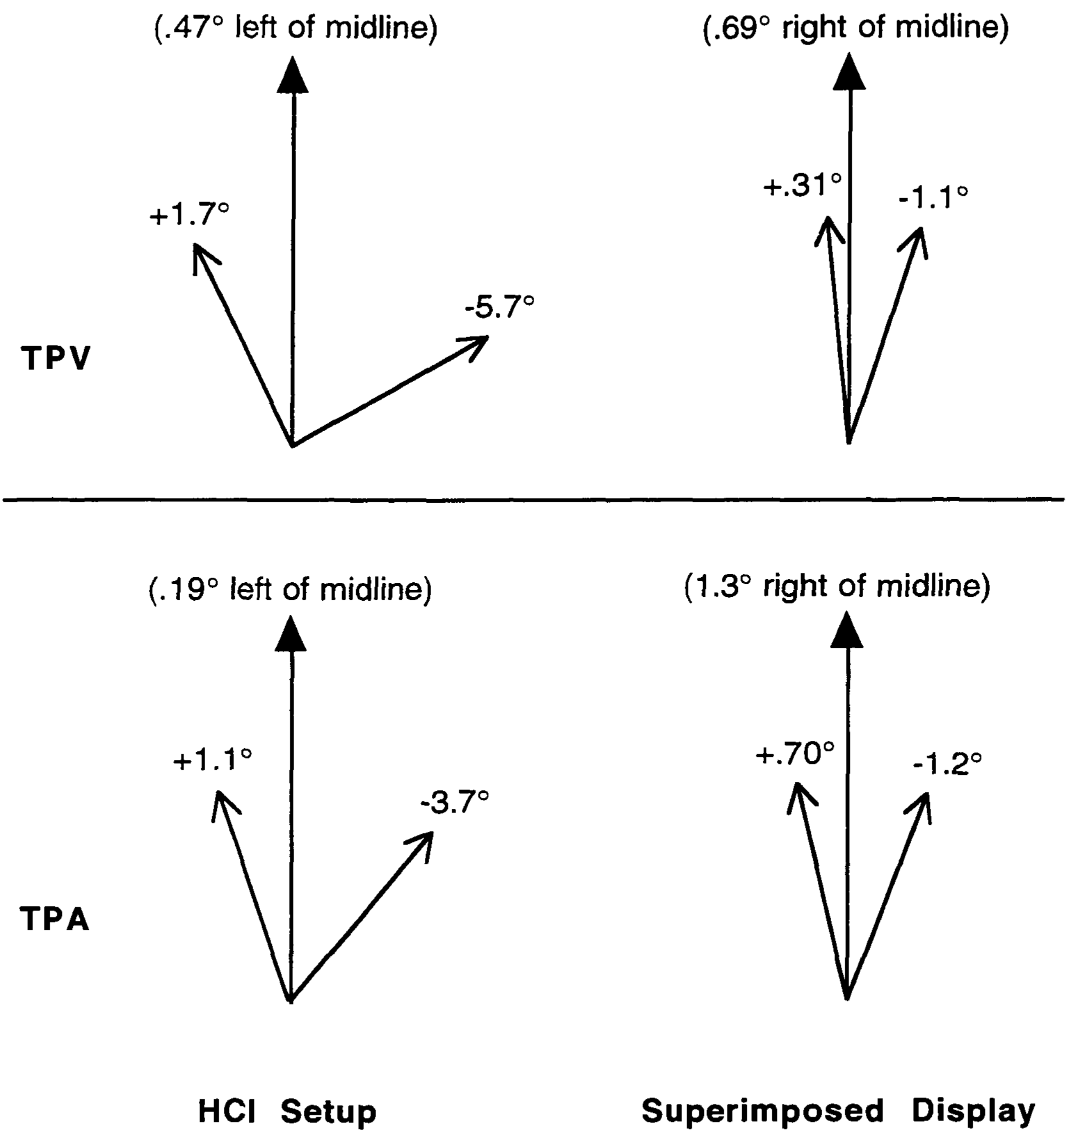

Figure 28. Summary of directional measures (tangent to the spatial path on a horizontal plane) for nonperturbed and perturbed trials. The solid arrow heads show reference values (in degrees from the body midline) for nonperturbed trials to the centre target. Open arrow heads to the left and right use an exaggerated scale to illustrate the relative magnitude of corrections to the direction of movement for perturbed trials, with respect to the reference direction for nonperturbed trials. The top diagram represents corrections at the time of the initial velocity peak (TPV, mean 159 ms). The bottom diagram shows corrections at the time of the first acceleration peak (TPA, mean 90 ms). Note the pattern of corrections is similar for the HCI and superimposed displays.

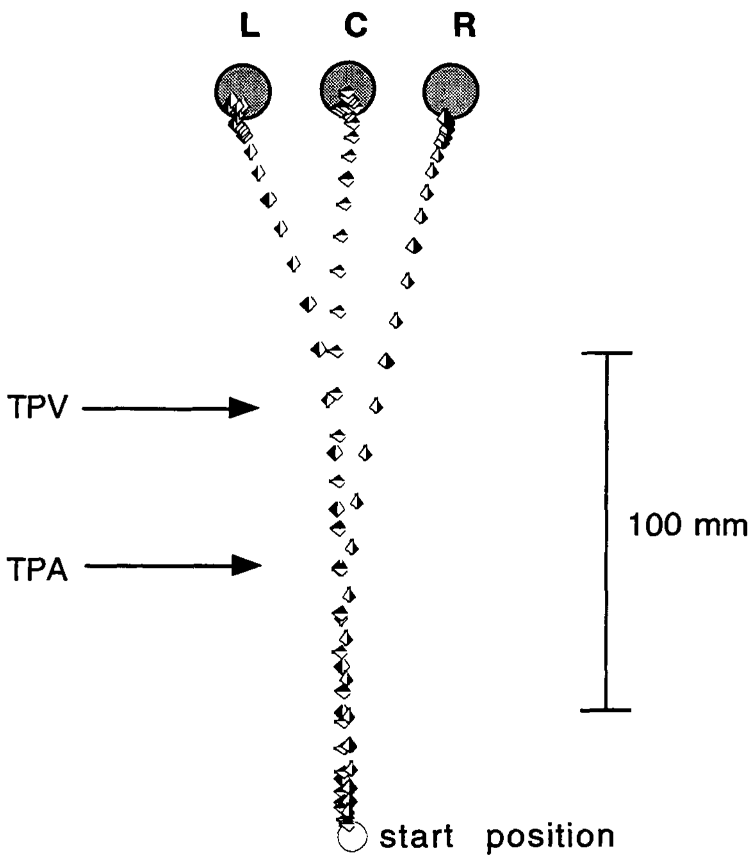

Figure 29. Movement paths in the x-y plane for representative trials in the HCI setup, shown from the point of view of the subject: (C) to the centre target position, (L) to a target initially presented in the centre position, and subsequently shifted to the left, and (R) to a target initially presented in the centre position and subsequently shifted to the right at movement onset. Each mark in the path represents the position from data sampled at 60 Hz. The analysis showed that changes in movement direction were measurable as early as the time of the first acceleration peak (TPA). Here, differences in the spatial path are clearly evident by the time of peak speed (TPV) of the hand movement.

and the pattern of corrections in directional measures illustrated in Figures 28 and 29 was similar for both the HCI setup and superimposed display conditions.

Discussion

The effects on movement kinematics when target position was perturbed were clearly different from the small systematic effects that occur normally for different target directions. As a result, the effects on kinematics in the perturbation analysis can be attributed to corrections or adjustments made to compensate for error caused by the perturbation, rather than to characteristic differences due to movement direction. In addition, the direction analysis revealed no significant differences comparing the first and last trial blocks. This indicates that subjects did not substantially change their strategy for nonperturbed trials over the course of the experiment as a result of encountering perturbations.

Trial block effects in the perturbation analysis were mainly due to subjects' adaptation to the perturbed trials, for which they had received no practice, in the HCI setup. One difference between the displays which could account for this interaction was the position of the work space relative to the subject's body. In the HCI setup, movements of the right hand even to the leftmost target still fell in ipsilateral space, to the right of the body midline, and movements to the centre target were approximately in line with the subject's shoulder. In contrast, with the superimposed display hand movements to the left targets crossed the body midline to contralateral space, and movements to the centre target were along the body midline. Thus, the fact that movements to the left targets were slower only with the superimposed display could be due to the increased difficulty of reaching into contralateral space, rather than the result of differences in the relationship between hand space and display space. Similarly, the adaptation in responses to targets perturbed to the right, but not for targets perturbed to the left over the course of the experiment may be due to the increased difficulty of making corrections to the movement trajectory in the direction of contralateral space. Further experiments, using a larger number of perturbed trials in each direction should address this possibility.

Two types of corrective response to the perturbations can be distinguished in comparison to normal trials. First, in the initial phase of the movement there were smooth corrections in the direction of movement without alteration in the timing or magnitude of the first acceleration or velocity peaks. The latency for this smooth directional response was very short; the estimate of less than 83 ms suggests a direct corrective mechanism requiring very little additional processing time. This result can be attributed to motor error being estimated, or dynamically evaluated, during the acceleration phase in the normal course of movement, rather than a mechanism based on a reorganisation of the motor response. The magnitude of directional corrections during the acceleration phase, however, was smaller than required to completely compensate for the change in target position.

The second type of response occurred later in the movement when homing in on perturbed targets. Here the deceleration phase took significantly longer, during which the movement trajectory exhibited large changes in direction in order to reach the new target position. The timing of this response correlated with subject's reports of their perceptual experience with perturbed targets; they felt that the target jumped to the left or right just before they were about to home in on the final position, when in fact the position had changed when they had covered less than 5% of the total distance to the target. These corrections can be attributed to visual comparison of the target and pointer positions. For a similar task, Komilis et al. (1993) concluded that there are two distinct mechanisms for controlling hand movements directed at visual targets: a visuomotor loop which makes crude corrections based on initial error in hand trajectory, and a refined mechanism based on retinal error between the target and pointer. The results of the present experiment support this view: the fast corrections were crude in that they did not completely compensate for the error in movement direction; but in all cases the target was reached by further corrections in the deceleration phase.

The results raise a question as to how how initial errors in hand trajectory are detected when a subject points to targets using our HCI setup. Contrary to the experimental hypothesis, there was no evidence of an interaction between the two display conditions, such that fast corrections in the initial phase of the movement were observed only for the superimposed display. In fact, the magnitude of direction changes in response to perturbations was greater in the HCI setup than for the superimposed display (see

Figure 28). This result is at odds with the dominant theory that the mechanism for fast corrections is based on a neural map which relates gaze, directed to a location in the work space, directly to the limb position required to reach that location. Were this the case, there would be no evidence of fast corrections with the HCI setup, where gaze was directed away from the goal location for the hand towards the graphics display.

One possibility is that eye movements are coupled directly to the signals that control the hand during rapid movements, as suggested by Goodale et al. (1986). Although we did not measure gaze in this experiment, the experimental system was designed to enforce a particular pattern of eye movements, and these were confirmed by the experimenter for each subject while they were performing blocks of trials. Gaze was fixed on the start mark until the target was presented, when it was shifted to the target position well before (greater than 400 ms) the movement was initiated. This, and the fact that subjects reported that they were aware of the perturbations indicates that their movement started well after the initial saccade to the target. When the target position was perturbed, subjects altered their gaze to the left or right accordingly. Cunningham (1989) has performed experiments with a system similar to our HCI setup, and has found evidence that visuomotor maps may be based on polar coordinates, based on relatively independent representations of direction and distance. In our experiment, in this coordinate system the relative change in direction (left, right) of perturbed visual targets is highly compatible with the change in direction of hand movements (left, right) required to correct for error in the initial phase of the movement.

Although other studies have suggested that vision of the hand and arm is not necessary to correct aiming errors when a target is viewed directly, reafference from the moving limb may still be important in the absence of gaze information specifying target position for the hand in the work space, as is the case with our HCI setup. Whether or not reafference from the moving pointer, analogous to the moving limb, contributes to this error signal could be determined by further study.

We can propose three experimental conditions to determine more precisely the source of visual information in the HCI setup leading to the fast corrections observed in the present study. One approach would be to remove vision of the pointer during the acceleration phase of the movement. In this case, for perturbed trials, only retinal or gaze information about the target could contribute to the error signal caused by the sudden shifting of target position. If there was no evidence of fast corrections for perturbed trials, then we could conclude that reafference from the pointer was contributing significantly to the error signal. A second approach would be to rotate the mapping from hand space to display space by an amount equal to the shift in direction of the target on perturbed trials. In this case, retinal and gaze information specifying target position would be shifted, but visual information about the direction of movement of the pointer relative to the target would be unchanged for perturbed trials. Fast corrections observed here would in fact be miscorrections, and would support the hypothesis that eye movements are the source of error information for directional corrections during the acceleration phase. The third approach would be to rotate the mapping between hand and display space at movement onset by an amount equal to the change in target direction in the present study. This type of perturbation, in which the viewed position of the target is not shifted, would create an error in direction which could only be determined by comparison of the movement of the pointer relative to the target. In this case, fast corrections would indicate that reafference from the moving pointer was contributing significantly to control of the movement in the acceleration phase.

In the present study, direction to the target was perturbed without changing the distance to be covered. Fast corrections in response to the perturbation had no effect on the timing or magnitude of the initial velocity peak, a result consistent with Experiment 1 which demonstrated that the scaling of kinematic features in the initial phase of a pointing movement are a function of the distance to be covered, and independent of the varying difficulty of accurately controlling the final position to reach small or large targets. If, as suggested above, direction and extent are coded separately in visuomotor maps, allowing humans to adapt to situations like the HCI setup in this experiment, we wonder if perturbing the distance to a target at movement onset would also result in fast corrections of a different kind than were observed when perturbing direction. This has been investigated in grasping physical objects (Castiello et al., 1991), but not in pointing on a computer display.

In all of the experiments in this thesis, subjects did not use a physical input device;

the experimental system tracked the index finger as a pointer directly. This was done in order to control or eliminate as much as possible some undesirable characteristics of typical input devices — lags, nonlinearities, granularity — which might confound the results, as well as to study a movement which represents a natural gesture useful in a wide range of graphical computer interfaces. It was surprising to see evidence of extremely fast visuomotor control processes in our HCI setup, and this result has implications for research in control and coordination of movements in virtual environments. First, it is important to pursue this line of investigation by detailed analysis of performance, in addition to typical usability approaches where task performance with particular input devices is compared, and preferences or opinions are solicited from users. Users of a system are not always aware of processes by which they achieve goal-directed movements, particularly in a less familiar, more abstract environment such as our HCI setup. Second, the latency of sensorimotor control processes evident in this experiment is smaller than the common wisdom about visual reaction times; they are generally considered to be much greater than 100 ms.

This short latency sensorimotor control emphasises a fundamental difference in the ways of interacting with a virtual environment. For scene viewing, where the user's movements are not directly represented and communication takes the form of discrete actions representing choices (e.g. pressing a button to change something in the scene), display update rates and lags may not affect performance to a great degree. In contrast, for visually coupled systems, where the user's movements are represented in order to give the illusion that they are interacting with a virtual world which appears to be an analogue to some physical reality, movement planning and control processes will be sensitive to decorrelation of actively generated sensory information in relation to the movement. Whether the virtual environment represents a hand directly in working space, as with the superimposed display, or indirectly, as in the HCI setup, movement control processes still respond quickly to small variations in visuo-spatial information. The technical difficulties of maintaining an accurate representation of the position of a moving limb can be clarified with a simple example. In our experimental system, the overall lag is somewhat less than 25 ms from sensing the hand position to rendering on the monitor screen. Although this value is fairly small compared to a commercial system employing a data glove, typical goal-directed hand movements reach a peak speed of about 1 metre per second. With the superimposed display, this lag results in a displayed error of about 25 mm in finger position, more than the width of the target (24 mm) to which the subjects were pointing.

In summary, the results demonstrated evidence of fast sensorimotor control processes correcting errors in the direction of hand movements to a target, even when subjects were not viewing the goal position directly. This suggests that these processes can be based on information other than a combination of retinal and proprioceptive signals to map gaze directly to a goal position for the hand and arm. Further study is required to fully understand how eye movements and motion of the pointer in our HCI setup contribute to fast corrections in the initial phase of a pointing movement. These results also suggest that a model of human-computer interaction (using pointing devices) based on error signals derived from visual comparison of cursor and target position needs to be augmented to take into account these fast correction processes. For graphical displays and virtual environments where the motion of a pointing device or the hand is continuously represented, understanding the exact nature of these fast processes is important for system design. On the one hand, visually representing important features of the movement profile can contribute to speed and accuracy in a pointing task; on the other hand, misrepresenting the same features due to lags or inaccuracy in spatial registration can cause difficulties. In either case, control and coordination of pointing movements is affected by sensorimotor processes of which the user is largely unaware.

Experiment 4: Non-linear Control-display Mappings

As introduced in Experiment 1, one important design issue in HCI is the proper selection of control-display gain for a particular configuration of input device, graphics display, and set of tasks (Buck,1980). For devices which transduce motion or position, higher gain offers the advantage of a reduced footprint; a smaller work space is required to reach all available locations on a display. In addition, results from previous experiments suggest that pointing performance will actually be improved by increasing the control-display gain to make the work space for the hand as small as practically possible.

Of course, there are disadvantages of using higher gains; hand tremor, limited device resolution, granularity, and accuracy and stability of the motion and position sensors all compromise our ability to reliably select a target as small as a single pixel. A common solution to these problems, used in most commercial systems, is a two-step gain or so-called "powermouse". In this design, at slow speeds the gain is reduced in order to make the device more controllable to reach small targets, and at faster speeds associated with distance-covering, the gain is increased. According to Jellinek & Card (1990), the real advantage of such systems is the reduced footprint, as they could demonstrate no overall advantage in pointing speed with various powermouse settings. Flach, Hagen & O'Brien (1987) used what amounted to a continuous variation in gain over the course of positioning movements with a joystick (log space mapping) and did find an advantage, both in terms of time and also in the number of corrective movements, for reaching a small target.

There is some support for the idea that varying the gain as a function of movement speed would improve performance, given our understanding of pointing movement phases. For example, compare a constant unity gain mouse to a power mouse with a low speed gain of unity, and a high speed gain of two. The powermouse could reduce the total distance travelled by the hand by up to one half, while leaving the target width as an accuracy constraint for the hand unchanged compared to the constant gain mouse. Not only should this arrangement decrease the effective A, but also lower the ID for the hand, for a given A and W combination on the display. A possible drawback is that discrete gain changes amount to a perturbation, and have been shown to disrupt the characteristic smooth velocity profile of a pointing movement. The resulting compensatory reacccelerations and iterative corrections (Graham et al., 1993) can actually increase the time to reach a target. How do humans adapt when the mapping (control-display gain) between hand motion and the motion of the pointer on the display is changed during the course of a movement? Does such a strategy really offer an advantage in terms of improved performance?

Experiment 4 was designed to test two non-linear control-display mappings (a discrete gain change, and a continuous gain change) which offer to reduce the size of the work space for the hand, and to examine how movement kinematics of both the hand and cursor changed to compensate for distortion in the relationship between hand and display motions. One prediction is that characteristics of the movement trajectory which result from the movement plan should stay the same for the hand, but be distorted for the cursor in non-linear compared to linear mapping conditions. In particular, the timing and magnitude of the initial velocity peak should be a function of movement distance for the hand in all conditions, as demonstrated in previous experiments. On the other hand, features governed by visual comparison of the target and cursor position should stay the same for the cursor, but be distorted for the hand in non-linear compared to linear mapping conditions. Thus the deceleration characteristics for the cursor should remain similar in all conditions.

Method

Subjects

Six subjects with normal or corrected vision were paid to participate in the experiment. All had experience using computer systems with pointing devices, and preferred to use their right hand in this context. The subjects were naive as to the purpose of the experiment.

Apparatus

The experimental setup was the same as in Experiment 1, using the standard HCI setup. At the start of each experiment session, we defined a work space for the hand having exactly the same dimensions (350 by 280 mm) as the display area of the computer monitor. Both three-dimensional position data from IREDs on the hand, and projected coordinates which controlled the position of the pointer on the display were saved in computer files for subsequent analysis. Three different control-display mappings (illustrated in Figure 30) were contrasted, each characterised by gain as a function of hand speed in the plane of the table top:

1) Constant Gain. In this condition, gain was set to unity, with the result that each millimetre of hand movement (in the plane of the table top) by the subject resulted in one millimetre of pointer movement on the display. In this condition, the position of the finger over the work space was mapped directly to a corresponding position on the monitor screen.
2) Discrete Gain Change. In this condition gain was unity when the finger moved slowly (less than 200 mm/s). Gain was doubled when the finger moved quickly (faster than 200 mm/s) so that each millimetre of hand movement resulted in two millimetres of pointer movement.
3) Continuous Gain Change. In this condition, gain varied continuously as a function of the speed of finger movement according to the relation:

$$\text{gain} = (\text{finger speed}/v_0)^p \qquad (12)$$

The parameter v_0 represents a slow speed criterion, and was set to 200 mm/s for this experiment. The exponent p was set to 0.756471, so that gain was doubled at the fast movement criterion of 500 mm/s.

The fast movement criterion was chosen to represent a typical peak velocity for the range of conditions in this experiment, after examining the range of PV values for subjects in Experiments 1 and 2. The slow movement criterion was chosen to be similar to, but somewhat greater than 100 mm/s, used in previous studies (Jellinek & Card, 1990; Trankle & Deutschmann, 1991) employing a mouse with variable gain, because the results of Experiment 1 and 2 indicated that 200 mm/s represented an upper bound for subjects during the deceleration phase, using the hand for pointing in the standard HCI setup.

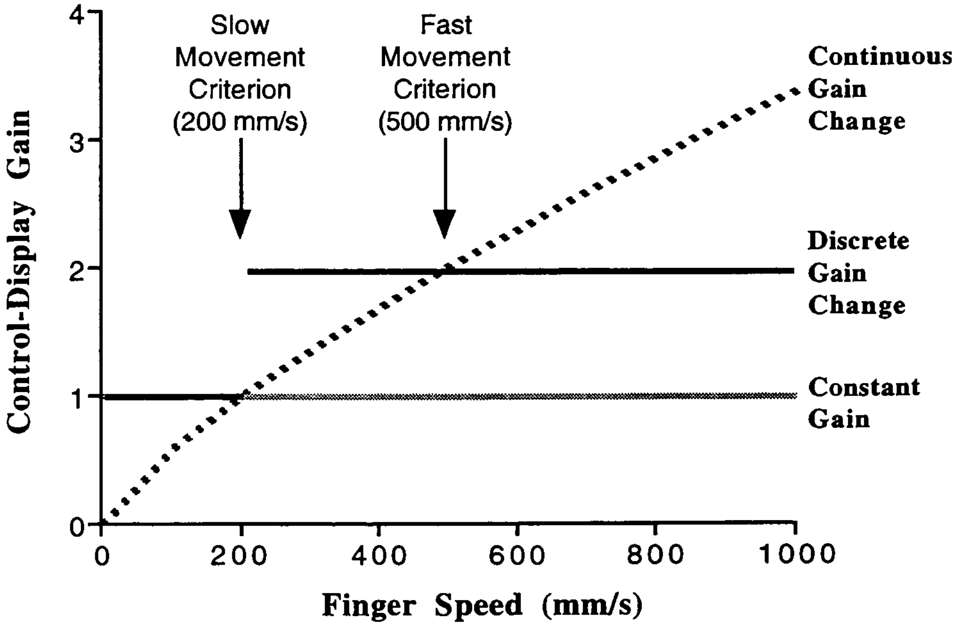

Figure 30. Gain as a function of hand speed in the plane of the table for the three control-display mappings used in the experiment. The control condition (grey horizontal line) used a constant unity gain. The discrete gain change mapping (two black lines) used unity gain for the slow portion of movements (speed less than 200 mm/s), and doubled the gain for the faster portion of movements. In the continuous gain change mapping (dashed line), gain increased smoothly with speed. For comparison with the other conditions, the continuous mapping was chosen to be equal to the discrete mapping gain at the slow and fast movement criterion values. Previous work showed that 200 mm/s represents an upper bound for speed in the deceleration phase, and that 500 mm/s represents a typical peak velocity for the range of conditions in this experiment.

Procedure

At the start of each trial, the subject rested the hand on the table, positioning the index finger so that the tip of the red pointer was held inside a start mark (5 mm diameter white circle) on the display. The experimenter viewed an identical display on a separate monitor placed in the corner of the room, and initiated a trial by pressing a button, causing white circular target to be presented. While looking at the monitor screen, the subject lifted the hand and pointed with the index finger to a place on the table, causing the tip of the red pointer to land anywhere inside the target on the display. The subject was told not to rush to initiate the movement, but to move quickly, and to make a small correction if they missed the target. Four movement amplitudes on the display (A = 37.5, 75, 150, 300 mm) were combined with four target diameters (W = 3, 6, 12, 24 mm) for a total of sixteen combinations with IDs varying from 1.6 to 6.6 bits, with all movements made in the same direction, away from the subject parallel to the body midline.

Data for each subject were collected during a single session lasting about one and one half hours. For each control-display mapping, the subject performed a block of up to 64 practice trials to different amplitude-width combinations. The practice session was stopped when the subject expressed that they were comfortable with the setup, and after they had reached four targets in a row quickly without making large corrections. After practice, the subject performed a block of 192 trials to all amplitude-width combinations in random order, so that there were twelve trials for each A-W combination for analysis for a particular control-display mapping. After a short rest period, the subject continued the experiment using a different control-display mapping. The order of presentation of the three control-display mappings was counterbalanced across the six subjects.

Data analysis

As in Experiment 1, three-dimensional position data from the distal index finger IRED, and two-dimensional position data from the pointer on the display were analysed to produce the following dependent measures: movement time (MT), magnitude and timing of the first velocity peak for the hand (PV, TPV) and cursor (CPV, CTPV), proportion of time after the first deceleration peak for the hand (%TAPD) and cursor (C%TAPD), and effective target width (We). For this experiment, two additional measures were also calculated:

1) SUBM - The number of corrective submovements was counted from the first negative peak in the hand acceleration profile, using a method similar to that in (Walker et al., 1993). A submovement was defined as a positive zero-crossing in the acceleration profile which continued to a criterion value of 100 mm/s/s.
2) DIST - The actual distance of the hand movement was defined as the straight-line distance from the start position to the position at the end of movement.

Means for each dependent variable were calculated from the last ten good trials in each condition, and used for an analysis of variance with repeated measures on a subjects (6) by mapping (3) by movement amplitude (4) by target width (4) design. In addition, multiple correlations were calculated using the two-part model of Equation 5 to characterise movement time as a function of A and W for each control-display mapping.

Results

The main effects from the ANOVAs on all dependent measures for control-display mapping, movement amplitude, and target width are summarised in Tables 16, 17, and 18 respectively. F and p values are also given for each effect; these will not be included in the text. The analysis contrasted the control condition (constant gain) with the two non-linear mappings (discrete gain change, continuous gain change).

Control-display Mapping Effects

Table 16 indicates that there were statistically significant and pronounced effects of control-display mapping on all dependent measures. Movement times (MT) in the constant and discrete gain change conditions were similar (696, 683 ms), but were significantly greater for the continuous gain change mapping (782 ms). The first velocity peak for the cursor (CTPV) occurred earlier by more than 25 ms in the non-linear mapping conditions than with constant gain. Similarly, the first peak for the hand (TPV) occurred earlier by about 40 ms. The magnitude of the first peak for the hand (PV) was dramatically reduced with non-linear mappings (from 514 to 325 and 250 mm/s), but the magnitude of the

TABLE 16.

Main effects for control-display mapping on cursor and hand kinematics

Mapping	Constant	Discrete	Continuous	$F_{2,10}$	p <
MT (ms)	696	683	782	16.99	.001
CPV (mm/s)	474	535	535	2.73	.110
CTPV (ms)	194	170	165	19.86	.001
C%TAPD	56.8	62.1	70.5	42.67	.001
We (mm)	9.27	11.11	10.95	19.54	.001
PV (mm/s)	514	325	250	114.42	.001
TPV (ms)	180	153	137	41.41	.001
%TAPD	62.6	68.8	75.3	76.36	.001
SUBM	.59	.71	1.06	11.68	.002
DIST (mm)	141.0	71.4	55.5	1493.83	.001

cursor peak (CPV) was not significantly different across mappings. The earlier peaks for the non-linear mapping conditions suggest that increased movement times were not a result of subjects simply slowing down their movements to adapt to the increased gain. The invariance of CPV across conditions suggests that one element of the movement plan is to vary the peak speed for the hand to achieve an ideal, visually determined speed for the cursor.

The extra movement time for the non-linear conditions was accounted for by significant differences in the proportion of time spent decelerating to the target for both the cursor (C%TAPD) and hand (%TAPD). Along the two-dimensional path of the cursor, the proportion increased from 56.8 per cent for constant gain, to 62.1 and 70.5 per cent for the discrete and continuous gain change mappings. Along the three-dimensional path of the hand, a similar pattern was evident, the proportion increasing from 62.6 with constant gain to 68.8 and 75.3 per cent for the non-linear mappings. The increase in time decelerating to the target was not the result of smoothly extending this phase of the movement, however. The average number of corrective submovements (SUBM) increased significantly from .59 in the control condition, to .71 for the discrete gain change, and 1.06 for the continuous gain change mappings. Increased difficulty of homing in on targets with non-linear mappings was also evidenced by increases in effective target width (We) compared to the constant gain condition (11.11 and 10.95 mm, compared to 9.27 mm).

As expected, the actual distance between start to end points for the hand (DIST) was significantly reduced for the non-linear mappings. With constant gain, the mean was 141 mm, identical to the predicted value of 140.6 mm which would be observed if subjects always aimed for the centre of the circular targets. Actual distance was almost half (71.4 mm) for the discrete gain change condition, indicating that subjects were in fact taking advantage of the doubling of gain in the distance-covering phase of movements. This was reduced to almost one third (55.5 mm) for the continuous gain change mapping, indicating that subjects were taking advantage of much higher gains for distance-covering in this condition.

The ANOVA revealed significant two-way interactions between control-display mapping and amplitude for CTPV and TPV. These are summarised on the left side of

Figure 31, which shows that the systematic increase in time to the first velocity peaks (CTPV and TPV) for larger amplitudes was less for the continuous gain change mapping than for the constant and discrete gain change conditions. This effect could be due to the subjects' attempts to control maximum cursor velocities for faster movements in the continuous gain change mapping, resulting in a different scaling of the movement program as a function of amplitude. On the right of Figure 31, we have plotted the same times as a function of the actual distance moved by the hand. First, note that the spread of the curves due to control-display mapping is reduced. This result supports the idea that it is actual distance to be covered by the hand which is used to form the movement plan for the initial phase, rather than the visually perceived distance (A) on the display. Second, the effect of the interaction is somewhat reduced when the data are viewed in this manner, more so for the timing of the hand velocity peak than for the cursor.

A base trajectory representation for the velocity profile predicts that the timing and magnitude of the initial velocity peak will be highly correlated. This is evident in Figure 32, where a strong linear relationship is illustrated between these measures for the hand (PV and TPV), and also for the cursor (CPV and CTPV). For the hand, the individual correlations are almost perfect, with $r^2 > .99$ for all control-display mappings. For the cursor, the correlations are also strong, with r^2 of .85 for the continuous gain change, .96 for the discrete gain change, and .91 for the constant gain mappings. We suggest that Figures 31 and 32 reveal how subjects adapt their distance-covering strategy to the different control-display mappings: the base trajectory is always scaled to the actual distance they expect to move, given their experience with a particular mapping. The height (PV) and the width (TPV) of the initial portion of the profile covary because overall size of the velocity profile is determined as a function of the distance the hand is required to move. The ratio between height (PV) and width (TPV) is altered somewhat to adapt to different control-display mappings, but the relationship is strongly linear within all conditions.

Amplitude Effects

The main effects for movement amplitude on all kinematic measures are summarised in Table 17 . As predicted by Fitts' law, movement time increased for larger

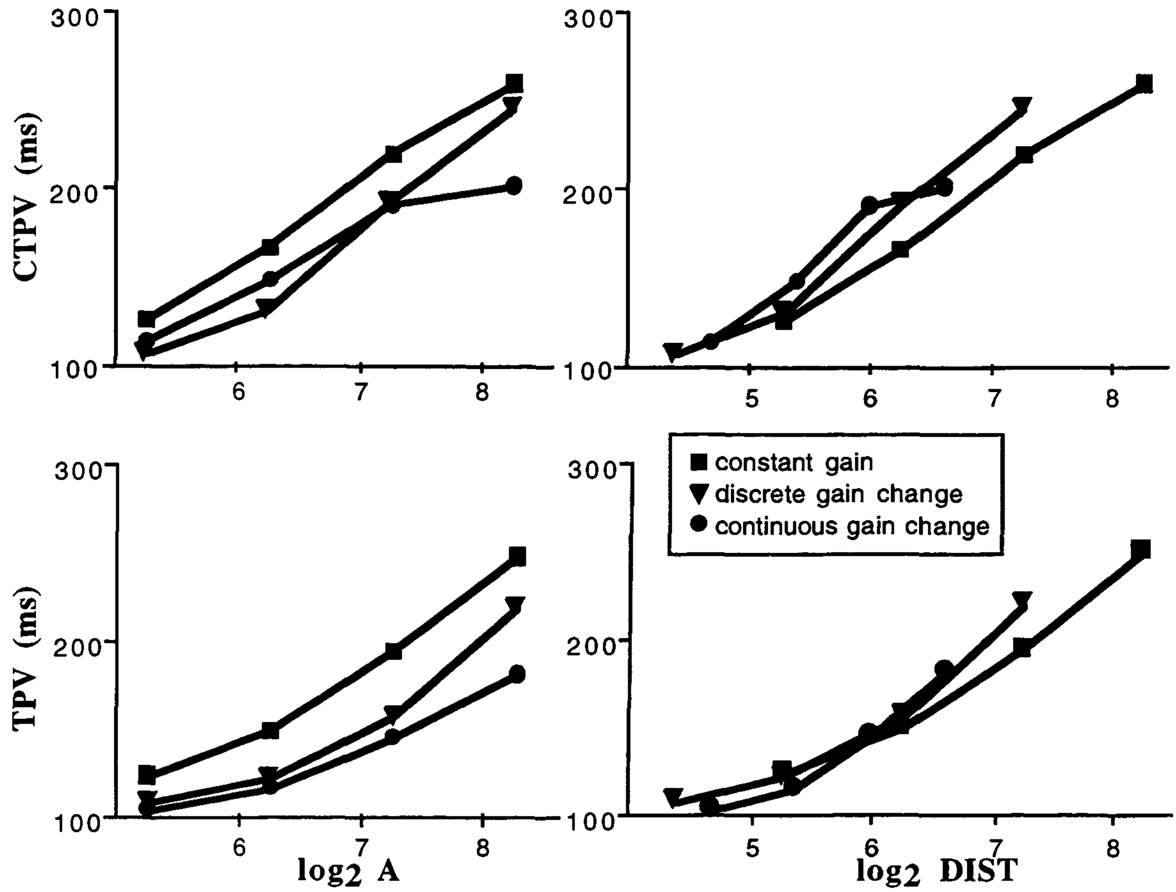

Figure 31. Interactions between control-display mapping and movement amplitude as evidenced by the timing of the initial velocity peaks for the cursor (top) and the hand (bottom). On the left, as movement amplitude on the display increased, time to the first velocity peak for the hand increased more rapidly in the constant and discrete gain change conditions than in the continuous gain change condition. The timing of the peaks versus the log of actual distance covered by the hand (DIST) is shown on the right. Using DIST, there is still some evidence of an interaction, but the magnitude of the main effect for control-display mapping is less pronounced.

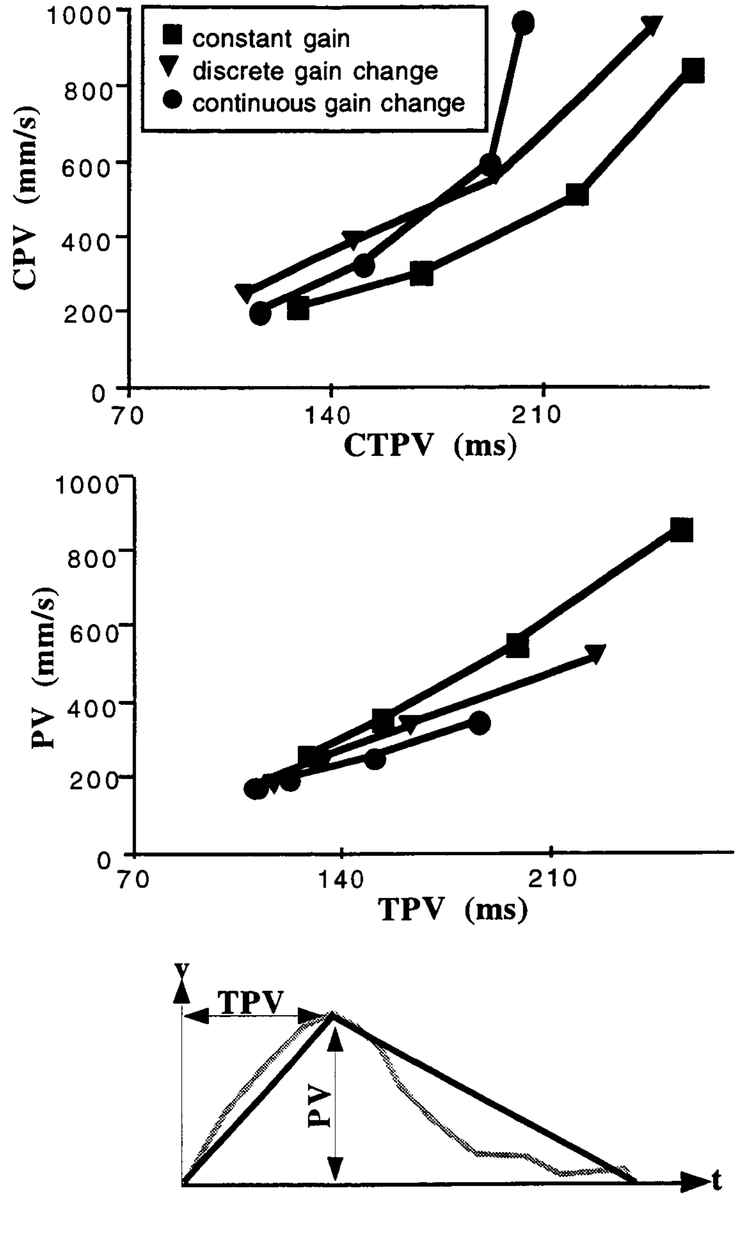

Figure 32. The relationship of the magnitude and timing of the first velocity peak for different control-display mappings. Measures derived from the three-dimensional hand trajectory (middle) show a strong linear relationship for each mapping. Cursor measures (top) are highly correlated, but the linear relationship is not as evident. At the bottom is a typical velocity profile (grey) and a schematic representation of its base trajectory representation. The shape (skew) is characteristic for a given target width. The overall size is scaled by the actual distance that the hand must move to cover to reach the target.

TABLE 17.

Main effects for movement amplitude (A) on cursor and hand kinematics

Amplitude (mm)	**37.5**	**75**	**150**	**300**	**$F_{3,15}$**	**p <**
MT (ms)	533	630	755	965	376.31	.001
CPV (mm/s)	231	333	565	931	197.45	.001
CTPV (ms)	116	150	202	236	138.98	.001
C%TAPD	65.1	63.5	60.1	63.7	4.12	.026
We (mm)	9.3	10.1	11.5	11.0	6.34	.005
PV (mm/s)	214	269	386	583	463.10	.001
TPV (ms)	114	130	166	217	69.17	.001
%TAPD	69.6	70.3	68.2	67.6	1.30	.312
SUBM	.68	.69	.78	.99	10.21	.001
DIST (mm)	27.9	51.1	95.8	181.1	4712.94	.001

amplitudes and for smaller targets. Systematic amplitude effects were evident in the timing and magnitude of the initial velocity peaks for both the cursor (CPV and CTPV) and the hand (PV and TPV), the peaks occurring later and with increased magnitude for greater distances. Amplitude showed no significant effect on the proportion of time decelerating to the target for the hand (%TAPD), but did show a significant small effect for the cursor (C%TAPD). The effect was not orderly, with movements of the second largest amplitude (150 mm) spending a smaller proportion of time in deceleration than either the smallest (37.5 mm) or largest (300 mm) amplitude. Effective target width (We) varied slightly with movement amplitude (9.3 to 11 mm) with values near the actual mean target width of 11.25 mm. Amplitudes also had an effect on the number of corrective submovements (SUBM), which ranged from .68 for the shortest distance, to .99 for the longest distance.

Width Effects

The main effects for target width are outlined in Table 18. The results showed no significant effect of width on the timing of the first velocity peak for both the cursor (CTPV) and the hand (TPV). The magnitude of the peak showed a significant but small increase for larger targets for both the cursor (CPV, from 492 to 533 ms) and the hand (PV, from 351 to 375 ms). In contrast, proportion of time spent decelerating to targets increased systematically as target width decreased, from 55.6 to 69.2 per cent for the cursor (C%TAPD), and from 62.6 to 73.8 per cent for the hand (%TAPD). Effective target width (We) ranged from 5 to 19.5 mm, demonstrating a slight range effect compared to the actual extreme values of 3 mm for the smallest and 24 mm for the largest targets. The number of corrective submovements (SUBM) decreased from 1.04 for the smallest, to .54 for the largest target. The magnitude of the effect of target width on SUBM was similar to that for movement amplitude.

Fitting Movement Times with the Two-part Model

It is evident from the movement time (MT) data in Tables 17 and 18 that the effect for changes in amplitude (A) was much more pronounced than for width (W), even though the ratio of the largest to smallest A, and largest to smallest W was identical (3 bits). This

TABLE 18.

Main effects for target width (W) on cursor and hand kinematics

Width (mm)	**3**	**6**	**12**	**24**	$F_{3,15}$	p <
MT (ms)	837	766	690	590	168.11	.001
CPV (mm/s)	492	514	521	533	4.94	.014
CTPV (ms)	172	178	179	175	4.05	.028
C%TAPD	69.2	65.7	62.0	55.6	98.58	.001
We (mm)	5.0	6.2	11.5	19.5	345.18	.001
PV (mm/s)	351	361	366	375	20.57	.001
TPV (ms)	156	159	155	156	0.95	.442
%TAPD	73.8	70.9	68.4	62.6	79.02	.001
SUBM	1.04	.82	.75	.54	25.91	.001
DIST (mm)	89.7	89.6	89.2	87.6	7.32	.003

result is consistent with Experiments 1 and 2. For the three control-display mappings in the present experiment, the two part models based on Equation 5 explain at least 97 per cent of the variance in movement time:

$$\text{Constant:} \quad MT = -105 + 153 \log_2 A - 77 \log_2 W \qquad (R^2 = .97) \qquad (13)$$

$$\text{Discrete:} \quad MT = 9 + 131 \log_2 A - 73 \log_2 W \qquad (R^2 = .97) \qquad (14)$$

$$\text{Continuous:} \quad MT = 149 + 138 \log_2 A - 101 \log_2 W \qquad (R^2 = .98) \qquad (15)$$

Note that in each case, MT is more sensitive to changes in amplitude than width. One explanation would be that subjects were failing to take advantage of the relaxed accuracy constraint of a larger target to speed up their movement. To answer this question, we modelled MT based on effective target width (We) rather than actual width, giving:

$$\text{Constant:} \quad MT = -72 + 162 \log_2 A - 111 \log_2 We \qquad (R^2 = .97) \qquad (16)$$

$$\text{Discrete:} \quad MT = 97 + 136 \log_2 A - 107 \log_2 We \qquad (R^2 = .97) \qquad (17)$$

$$\text{Continuous:} \quad MT = 133 + 147 \log_2 A - 113 \log_2 We \qquad (R^2 = .94) \qquad (18)$$

Although in this case MT is somewhat more sensitive to We, the pattern of coefficients is similar to the models based on W, indicating that subjects were taking advantage of larger targets. For predicting performance, a model based on the actual target dimensions is preferable as long as it effectively describes the subjects' performance in a task. For this reason, we continue to discuss the results for MT with respect to W on the display.

In Figure 33 we illustrate how the models of Equations 13, 14, and 15 can be visualised in a parameter space. Points for data well characterised by Fitts' law will fall on the line $b_1 = b_2$. Points which fall above the line indicate tasks for which movement time is more sensitive to changes in target width than movement amplitude. For these tasks, faster movement times are predicted when the scale of hand movements is increased. Points from the present study fall below the Fitts' law line, indicating that movement time will decrease when the scale of hand movements is reduced. The point for the discrete gain change mapping falls very close to the point for constant gain, indicating that the "powermouse" strategy provides no obvious advantage in terms of sensitivity to changes in MT due to

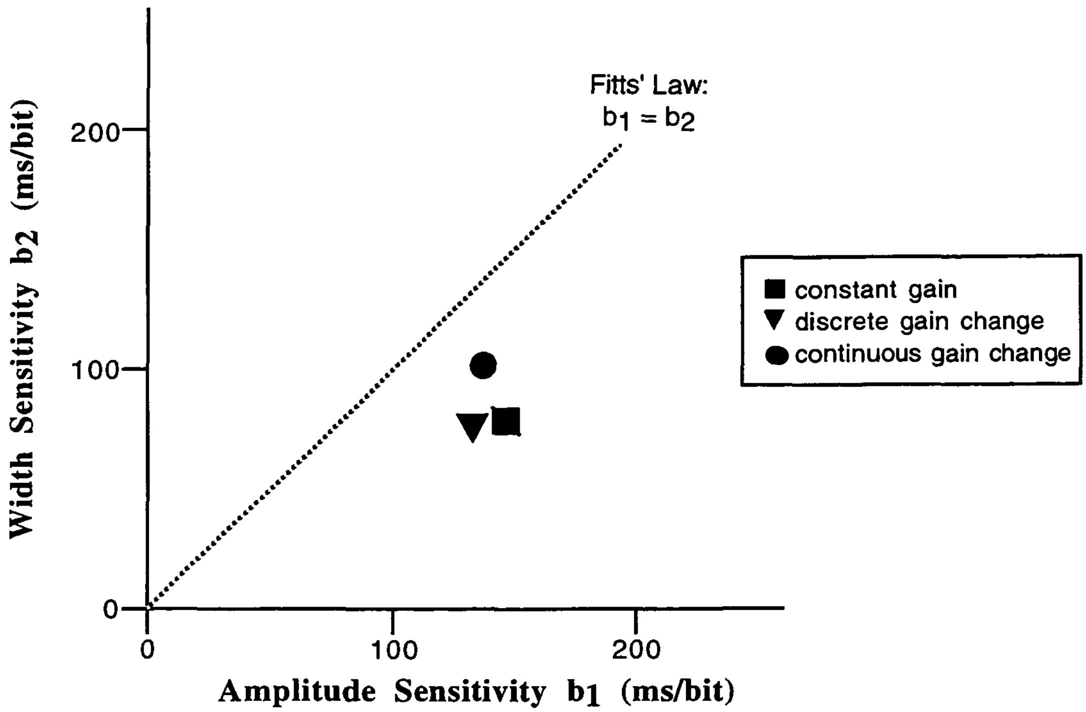

Figure 33. Parameter-space description for the two part model of Equation 5. relating changes in movement time to changes in amplitude and width. Points are plotted to characterise each of the three control-display mappings in the experiment. As observed in Experiments 1 and 2, points fall below the Fitts' law line, indicating that performance for a given index of difficulty will be faster for shorter distances. The vertical displacement of the continuous gain change condition relative to the other points indicates increased difficulty in controlling deceleration to the target, compared to the constant and discrete gain change mappings.

changes in A or W. The point for continuous gain change falls above the others, indicating this mapping is more sensitive to changes in W, probably due to increased difficulty in locating smaller targets.

Discussion

In general, the non-linear mappings seem to be disruptive in terms of movement control, with any distance-covering advantage paid for with increased time, number of submovements, and presumably difficulty, in decelerating to the target. The strong systematic relationship between actual movement distance for the hand and the timing and magnitude of initial velocity peaks indicates that the ballistic phase of a pointing movement is stereotypic, and based on a plan with simple features. The results of the present experiment suggest that peak velocity for the hand is selected to limit the speed of the visual representation of the movement, whereas timing of the first velocity peak is simply a function of distance to be covered by the hand. This strategy is effective with natural, three-dimensional pointing movements to physical targets, but may generate more variability when gain is changed during the initial phase of the movement; the resulting error must be dealt with in the deceleration phase. This property of the data might be revealed through further analysis of the variability of kinematic features in the different control-display conditions. The results support a base trajectory representation of the velocity profile for the hand for a particular control-display mapping, but the shape is dramatically affected by the type of mapping. Contrary to our experimental hypotheses, it seems that the features of the initial phase of the movement are determined both visually (from the display) and kinesthetically (from the hand), and that deceleration characteristics are distorted for both the hand and cursor for non-linear mappings.

Although they reduced distance covered by the hand, neither non-linear mapping demonstrated any significant advantage in performance in our pointing task. The slight increase in performance (13 ms) for the discrete gain change mapping, where gain increased to two for distance covering, over the constant gain of one was not as great as would have been achieved by simply using a constant gain of two. This effect can be calculated from Equation 13 by reducing both A and W for the hand by a factor of two. In this case, the predicted improvement would be 153 - 77 = 76 ms. This emphasises an important point illustrated by the position of the models in the parameter space of Figure 33: in all conditions performance will benefit from reducing the scale of hand movements, or equivalently, from increasing control-display gain. This result might not have been observed using a typical (less expensive) pointing device, but our experimental system has an extremely fine sensing resolution (≤ .1 mm), so that tremor is the limiting factor for controlling the display with free hand gestures at higher gains. Given the difficulties subjects encountered in adapting to the non-linear control-display mappings, and the disruption of their ability to decelerate to targets, we would recommend the strategy of increasing pointing device resolution to allow the use of higher constant gain in order to improve performance, rather than altering the relationship between hand and cursor dynamics.

General Discussion

As techniques evolve to support more realistic and direct forms of human-computer interaction, understanding the nature and limitations of human sensorimotor control processes will become increasingly important. The four experiments described in this thesis address the question of why mediated pointing is difficult from different perspectives. In this section, we will consider the results as a whole, both in terms of theory of how movement planning and coordination is affected by the constraints of the real and virtual environments, and also in terms of implications for analysis and design of interfaces for human-computer interaction.

Overview

Experiment 1 provided evidence that movement planning is largely based in hand space, with the implication that modelling and analysis of performance in mediated pointing should be based on the hand actions required for a task, regardless of the size of the display. The results of the other experiments support this view, although in Experiment 4 it was suggested that the visually determined peak velocity of the pointer in display space, in non-linear mapping conditions, affected the way in which planning parameters were adjusted to adapt to the changes in control-display gain. Even so, the systematic relationship between magnitude and timing of the initial velocity peak under non-linear mappings was more clearly evident for hand kinematics than for pointer kinematics.

Experiment 2 revealed that asymmetric effects of movement amplitude and target width on movement profiles were inherent features of the pointing task, and not due to difficulties or constraints in mediated pointing. This confirmed the need to resort to a more detailed two-part model of movement times, since gain (scaling A and W for the hand) effects on movement times were large, and of similar magnitude to target width effects, even when a subject pointed to physical targets in the work space.

Experiment 3 did not find support for the hypothesis that mediated pointing is difficult because the constraints of the HCI setup compromise the motor system's ability to make fast corrections. The results demonstrated two types of corrective processes, both of which were robust even when the spatial relationship between visual information and the work space for the hand is abstracted.

Experiment 4 demonstrated that kinematic features of the initial phase of movement are stereotypic for the hand, and that movement trajectories are based on a simple plan. The parameters of the plan for hand trajectory are adjusted to compensate for distortions in spatial mapping and control-display gain.

The combined results suggest that modelling and analysis of performance in mediated pointing must be based on hand actions required for a task, regardless of the nature of the display. The detailed analyses of hand actions for pointing on a computer display presented here suggest several strategies for engineering analysis and interface design. For example, we predicted that pointing performance will be improved by reducing the scale of hand movements; thus future effort should be focused on improving the sensing resolution of input devices for manual interaction. The two-part model for predicting movement time effectively characterises human performance, and provides a more detailed description of behaviour than Fitts' law. It captures scaling effects which are even more pronounced than accuracy constraints in physical pointing. In addition, the model parameter space predicts performance changes when hand movements are scaled, and characterises tasks and devices in terms of difficulties in the two movement phases.

The experiments also provide a detailed view of performance when pointing with an unencumbered hand, a benchmark against which pointing devices can be compared. The measurement and analysis techniques for identifying and separating direction, amplitude, and width effects will prove useful for predicting target location while a movement is ongoing, and can be applied to the design of an intelligent interface, employing precomputation to respond to object selection.

These points are considered in more detail in the following sections. First, a detailed explanation of scaling effects is provided, and the implications for interface design and how to measure human performance in pointing are discussed. With the conclusion that systematic effects of task constraints are based in hand space, we further discuss how these effects are separated in terms of movement phases. Fast corrections in the initial phase of pointing movements, and their implications for viewing, versus interacting with virtual environments are considered. Further discussion covers the derivation of the two-part model, its relationship to movement phases, and its use as an analysis tool to measure and distinguish difficulties in movement planning and control. Finally, we discuss how the real-time measurement of characteristic features of movement profiles and variation of features as a function of task constraints may provide advance information about object selection, and how to use this information for precomputation in a predictive interface.

Asymmetry of Amplitude and Width Effects

In Experiment 1 we concluded that modelling and analysis of human performance in mediated pointing should be based on understanding the hand actions required for a task, regardless of the scale of the display. The results demonstrated that in mediated pointing, examining kinematics of the cursor on the screen produced confusing results when gain is considered, while there was no such confusion in the analysis of corresponding hand movements. At first glance, this result seems counterintuitive, but is so only because of an assumption inherent in using the index of difficulty in Fitts' law as a predictive, rather than a descriptive model. Although the concept of ID captures the main feature of movement times as a function of speed-accuracy tradeoffs, it does so only for aggregate data. An important detail is lost in this simplification: empirical results demonstrated that changing movement distance by some factor affects movement time to a greater extent than does changing target width by the same factor. That is, the magnitude of the A effect is greater than the magnitude of the W effect on movement times. This asymmetry is marked in all the experimental results in this thesis, and most exaggerated for the natural version of the task, pointing at physical targets in the work space with full vision of the hand in Experiment 2. The asymmetry for physical pointing, summarised in Equation 9, is repeated here:

$$MT\ (ms) = -24 + 98 \log_2 A - 32 \log_2 W \qquad (R^2 = .96) \qquad (9)$$

which indicates that movement time increases by 98 ms for each doubling of distance, but decreases only by 32 ms as target size is doubled. Another way to understand the significance of the asymmetry is to consider the equivalent model (to Equation 9) in terms of index of difficulty and a scale factor:

$$MT\ (ms) = -25 + 32\ ID + 64 \log_2 A \qquad (19)$$

which reveals that, for pointing at physical targets in the workspace with the index finger, the scale or size of the hand space affects movement times (64 ms/bit) to a greater extent than the index of difficulty (32 ms/bit).

If we wish to restrict a model to describe a single effect over a wide range of movement amplitude and target width combinations, it would be more effective to capture the scale of the movements rather than the index of difficulty. For example, by a particular choice of conditions designed to emphasise the effects of scale, we can produce the result illustrated in Figure 34, where movement time actually decreases as a function of the index of difficulty. Although this example seems contrived, it is simply an exaggerated version of the confusing effects observed in the analysis of cursor kinematics with different gains from Experiment 1. Here the effect of changing gain was to scale hand movements even though A and W on the display were unchanged, and the scaling affected movement time, as well as all the other features of the movement trajectory.

What is the source of the asymmetry? One explanation is based on differences in the way large and small scale movements over the workspace are coordinated in order to move the tip of the index finger to a target. The smallest distance that was studied (in Experiment 1) was 19 mm, which can almost be covered by index finger flexion and extension. In contrast, the greatest distance of 300 mm requires the use of more, and larger segments including the shoulder. Rather than a single motor system which can be well characterised by a simple measure such as the index of performance derived from Fitts' law, the involvement of more limb segments means that performance of different motor systems is being measured at different distances (Card, MacKinlay & Robertson, 1991). Despite this difficulty, our results demonstrate that the main features of movement kinematics show considerable regularity regardless of the scaling of movements, but in a manner that would be difficult to account for using an approach based on a simple model derived from linear control systems or information theory. It may be possible to develop a detailed explanation of the effects using kinematic analysis of other limb segments when they move in conjunction with the finger and pointer on the display. Nevertheless, the

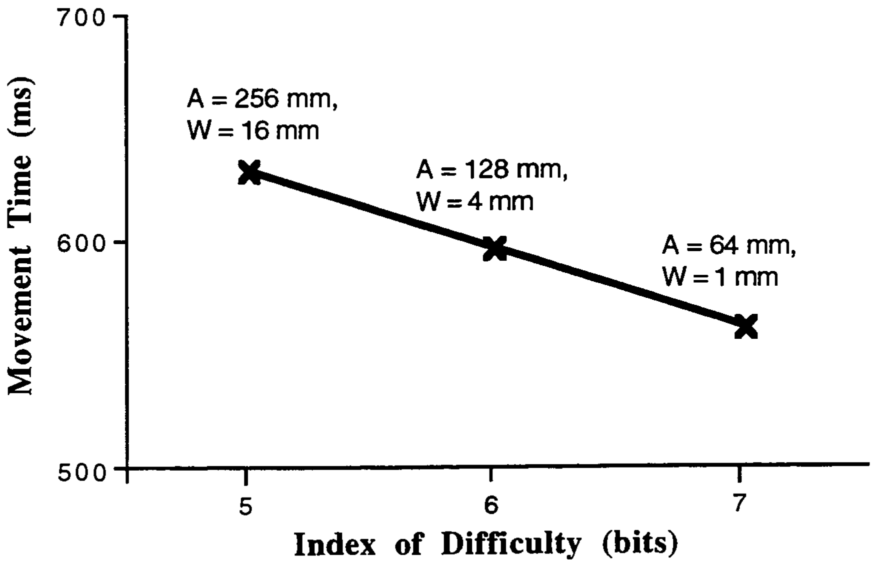

Figure 34. Predicted movement times as a function of index of difficulty (ID) from the model of movement times for physical pointing in Experiment 2. Exaggerating the decrease in target width (W) compared to the decrease in movement amplitude (A) leads to increasing ID as the movements are scaled down. In this case, movement times will actually decrease as ID increases, which would lead to a negative slope for Fitts' law.

two-part model and systematic A and W effects on movement trajectories reported here indicate that we can effectively consider the motor system as a black box, with certain properties of scaling, in the context of HCI tasks for scales from a few to several hundred millimetres.

In Experiments 1 and 4, in conditions in which there was not a unity gain, or linear mapping, the hand space analysis always provided a simpler, clearer picture of systematic effects than the analysis in display space. This supports the idea that hand movements are planned and controlled in hand space, and that visual information about task constraints — direction, distance, and target size — is transformed to body-centred coordinates for coordinating reaching to a target. The implication for empirical studies of HCI is that, given a specific input device, task, mapping, and display, some predictions can be made about how performance will vary based on restricted manipulations of these constraints as represented on the display. As a rule, however, these predictions cannot be generalised to other devices and systems. In other words, any change in the mapping will change the underlying model.

Software development techniques for graphical interfaces have emphasised the concept of device independence. That is, elemental interactions with a display are programmed without regard to the hardware details of input devices, or even the general category of device. With this approach, the programmer's model of the interface cannot distinguish pointing using step keys, joystick, graphics tablet, or mouse as input devices; all are considered equivalent. Dimensioning of objects is usually pixel-based, or done relative to an arbitrary coordinate system relating one or more display surfaces. Although device independence is an important concept for simplifying software development for interfaces, the concept creates difficulties for predicting how design choices will affect human performance at the level of the basic elements of interaction. To take the extreme position, we would advocate that these design tradeoffs can be best understood using the concept of display independence; that is, a view that disregards the physical details of the display (within reasonable limits) and uses measures related to a body or hand-centred coordinate system.

Movement Phases

In general the data from all experiments supports the idea that hand movement trajectories can be understood in terms of two phases. The initial phase is characterised by the timing and magnitude of the initial velocity peak. As suggested by previous studies, results here show that timing and magnitude of the initial peak are highly correlated, suggesting that they are the result of a single planning parameter: movement distance for the hand. Although the functional relationship of magnitude and timing showed an adaptation for the non-linear mappings in Experiment 4, this relationship remained strongly linear. The adaptation can not be easily attributed to a feedback process based on motion or the pointer on the display, as the relationships between A on the display and the magnitude and timing of the pointer velocity peak (CPV and CTPV) were complex. Even under these conditions, the systematic linear relationships between A for the hand, and the timing and magnitude of the velocity peak for the hand (PV and TPV) were preserved. The adaptation may be an attempt to keep pointer velocity in a range which preserves the illusion of apparent motion on a raster display. That is, at a 60 HZ refresh rate there is a maximum distance between successive images of the pointer beyond which the display appears as a stroboscopic series of images, rather than a single object in motion.

In our analyses we chose to use per cent time after peak deceleration (%TAPD), rather than simply the proportion of time after the first velocity peak, to describe the precision effect, or skew of the velocity profile in pointing to smaller targets. We found that %TAPD shows a clearer separation of phases (in terms of separable A and W effects) than in previous studies which divided the movement phases exactly at the time of the first velocity peak (TPV). Our data showed a gap of from 20 to 60 ms between the initial velocity peak and the first deceleration peak. Using a simple division of the acceleration phase from the deceleration phase by TPV shows separable amplitude effects on the initial phase, but leaves amplitude and width confounded in the proportion of time in the final movement phase. Conversely, using a simple division of the first phase from the second at the time of peak deceleration shows a separable width effect, but leaves width and amplitude confounded in the time spent in the initial phase. By not attributing the interval between TPV and the first deceleration peak to only distance covering or deceleration to the target, a much clearer picture emerges. We conjecture that in this interval, movement control switches from the planned impulse (feedforward) to a strategy of controlled deceleration by comparisons of the relative positions of the target and pointer.

As illustrated in Figure 35, the base trajectory representation of velocity profiles (introduced in Experiment 4) captures these effects schematically, although in this simplified representation the final phase of the movement is represented as time after the initial velocity peak. This diagram had been suggested for use with data normalised for movement time (to illustrate amplitude scaling effects), or normalised for peak velocity (to illustrate the precision effect). It is quite informative also to visualise actual data from the experiments this way. Profiles from all the experiments are plotted in Figures 36 to 38 to clearly illustrate the A and W effects on hand movement trajectories under a variety of conditions. A striking feature of the comparison of these profiles under a wide range of conditions is the similarity of the effects of speed-accuracy tradeoffs on movement kinematics for the hand regardless of control-display mapping.

Visually-Based Corrections

In terms of Woodworth's movement phases, vision contributes in different ways to each phase. The initial phase is based on a movement plan, given the visually determined direction and distance, in the work space, of the target from the start position. In deceleration, vision is used on-line to control the approach to the target with corrections to the trajectory based on an error signal derived from comparison of the limb and target positions. The discussion in the previous sections offers some evidence that both these mechanisms are affected by the constraints of a virtual environment, and offers a partial explanation of why and how mediated pointing is more difficult than natural pointing at targets in the work space. As introduced in Experiment 3, recent studies have demonstrated the involvement of faster and more direct mechanisms to correct the trajectory of a moving limb in aiming and pointing to a location, as well as reaching to grasp an object. Based on Jeannerod's (1991) model of the neural basis for fast corrections, it seemed likely that these were incompatible with the constraints of the standard HCI setup in which gaze is not directed at the goal location for the hand. This being the case, we would

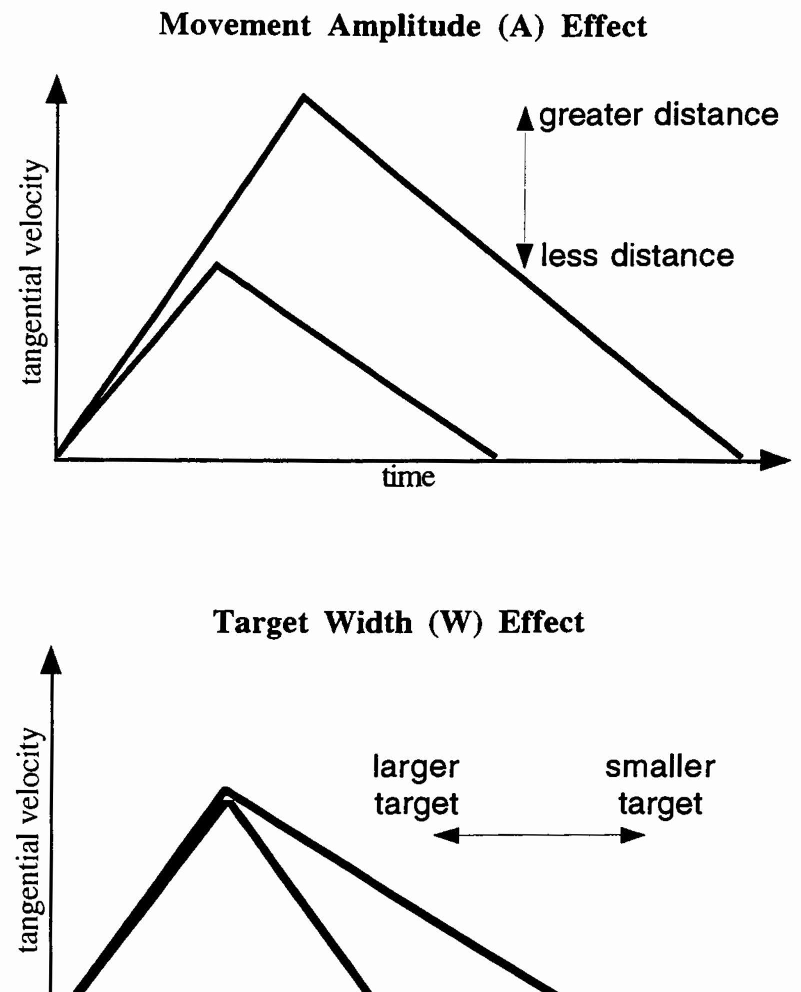

Figure 35. Base trajectory representation of velocity profiles introduced in Experiment 4. The features of the trajectory are determined by movement time (length of the base of the triangle), timing magnitude of the initial velocity peak (height of the triangle), and proportion of time in the deceleration phase (skew of the triangle). Above, movement amplitude affects the overall size, without affecting the shape. Below, target width affects the shape, without altering the position of the apex of the triangle.

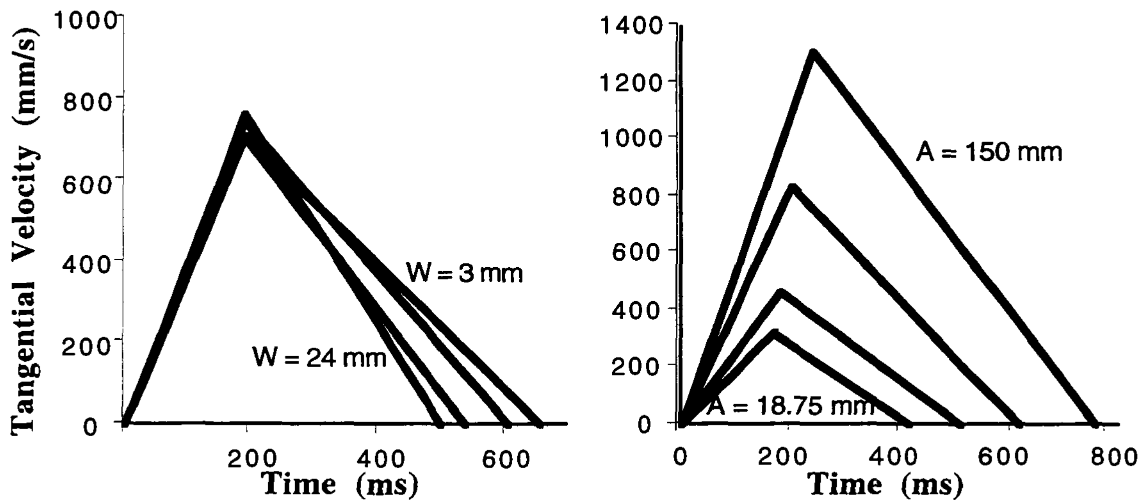

Figure 36. Experiment 1: schematic view of amplitude and width effects on hand velocity profiles. Points are actual data, not normalised, from means across subjects for a subset of conditions in the experiment..

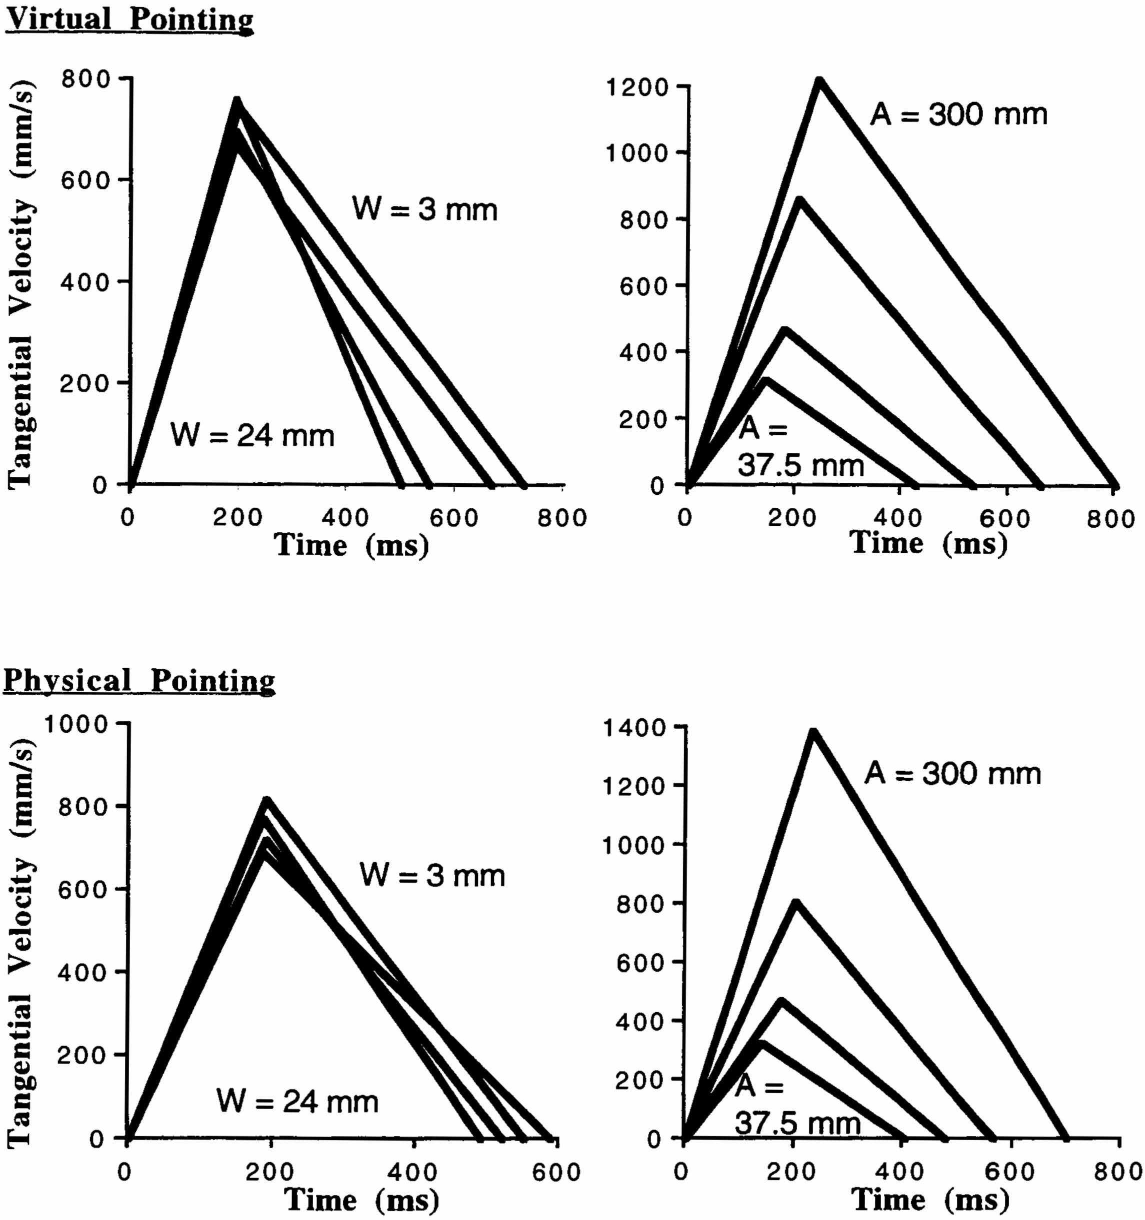

Figure 37. Experiment 2: schematic view of amplitude and width effects on hand velocity profiles from both virtual and physical pointing. Points are actual data, not normalised, from means across subjects.

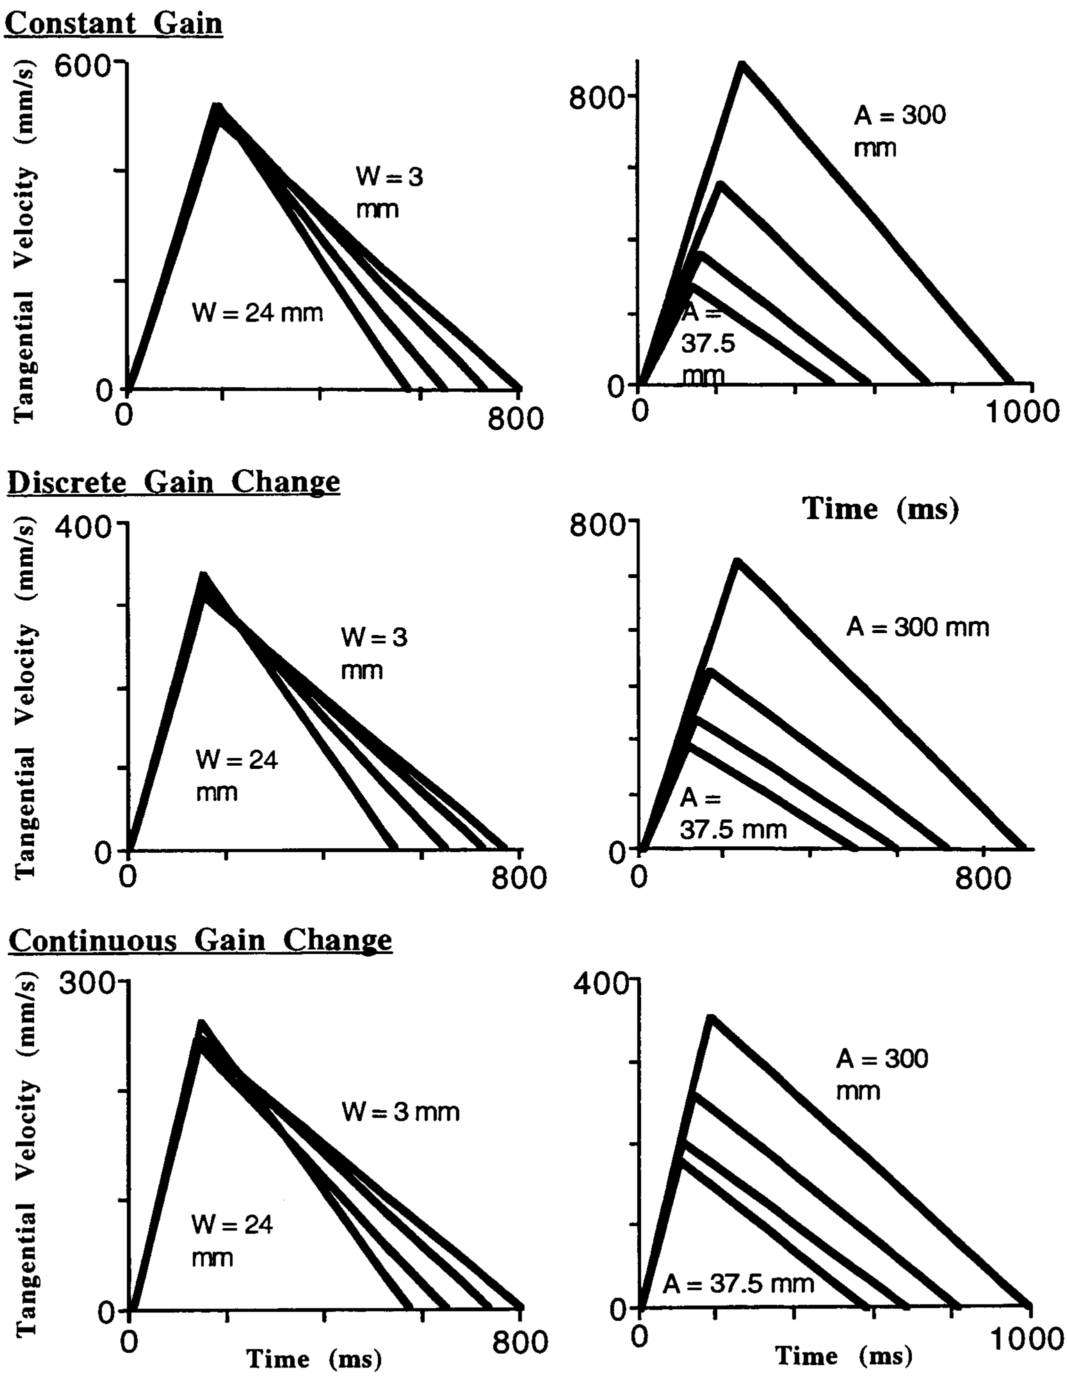

Figure 38. Experiment 4: schematic view of amplitude and width effects on hand velocity profiles in constant gain, discrete gain change, and continuous gain change conditions. Points are actual data, not normalised, from means across subjects.

be provided with another way in which mediated pointing differs from, and is more difficult than the natural version of the task. Contrary to our predictions, the results did not support this view, and demonstrated evidence for fast corrections even in our HCI setup. We have determined that this issue requires more study, and have spent some effort in the discussion of Experiment 3 to outline how the the contribution of gaze, eye movements, and visual reafference from the motion of the pointer can be determined by further experimentation. Whatever our confusion about the exact processes which produce fast corrections, the fact that they are evident even in a highly abstract version of a pointing task has important implications for performance requirements of virtual and augmented reality systems.

Our view of movement phases emphasises that what is commonly referred to as "eye-hand coordination" involves several mechanisms, each of which utilises visual information about the environment in a unique manner. Since Trevarthen (1968) proposed multiple mechanisms of vision in primates, there is evidence for at least two such channels. For perturbed trials in Experiment 3, although subjects were aware of their corrections to the movement trajectory in the deceleration phase, they were generally not aware of the perturbation or of their corrective response in the initial phase of the movement. Processes that modify the trajectory of hand movements but which operate below the level of awareness are suggested in other studies which have examined the effects of distorted visual representations of movement on hand trajectories. In one example (Nielsen, 1968) subjects reported that they felt a force dragging their limb to one side when the movement trajectory of the hand was visually misrepresented, even though the apparatus exerted no force. In another example, subjects in a study by Held and Durlach (1987) interpreted the effect of a small lag (30 to 60 ms) in a visually coupled system as the sensation of "dragging the hand through a viscous medium", whereas the image was perceived to be dissociated from their hand for longer lags. All of the above results suggest that users' perceptions are sometimes dissociated from their motor responses when these responses are required to compensate for discordant sensory information. The forces felt were clearly generated in the subjects own hands and arms as a result of visual capture, but we propose they were attributed to the environment because they involve sensorimotor mechanisms which operate below the level of awareness.

The requirements for controlling lags, and for maintaining accurate position tracking in visually coupled systems may be entirely different when users interact with, rather than simply view, a virtual environment. For example, Ware and Balakrishnan (1994) showed that lags in head-tracking had less effect than lags in the representation of limb position when subjects reached to a location in an augmented reality environment. Using an immersive display with head tracking, both lags and lack of sensing resolution result in a misrepresentation of the direction of targets (with respect to an eye-centred coordinate system) when the head is moved. Gentilucci et al. (1994) demonstrated systematic direction errors due to this misrepresentation even when the relative alignment of the image of the limb and target location were correctly rendered. When the position of the limb is also displayed in the virtual environment, errors due to lag are compounded, so that neither the absolute position of target and limb, nor their relative locations are accurately represented. Given that reaching to a location involves a natural synergy of eye, head, arm, and hand movements, our results suggest that lags even as small as 20 or 30 ms may affect eye-hand coordination if the head is moved, and perhaps even if there is little head movement when gaze is already fixed on a target location.

When a user interacts with such a virtual environment, head motion will create errors in target direction and a misrepresentation of the position and the direction of limb. Motor responses to correct for these errors are initiated extremely quickly, even though the user's perception in terms of scene viewing may be entirely satisfactory. Experiments using displacing prisms demonstrate that humans can adapt to a fixed error (e.g. Redding & Wallace, 1994) using a number of strategies. It is not clear how they adapt when the error changes dynamically, and is coupled to their own natural movements. Given that fast corrective responses occur below the level of awareness, it seems unlikely that users will be able to adapt completely to lags in virtual and augmented reality environments, or develop an ability to coordinate movements in an entirely natural manner.

Two-Part Model

Our two-part model was derived from one proposed by Welford (1968, p. 157) to account for asymmetrical amplitude and width effects in movement time data from repetitive tapping with a pencil. The original formulation was given as:

$$MT = a \log_2 A/W_0' + b \log_2 W_0'/W$$

$$= a \log_2 A - b \log_2 W + (b - a) \log_2 W_0' \qquad (20)$$

Here, W_0' is related to a theoretical error distribution of the end points of ballistic movements, but cannot be measured directly. The model also relies on the assumption that if this distribution could be measured, it would be independent of movement distance (A), as if ballistic movements depend on an absolute appreciation of the goal position, regardless of A. This assumption is in opposition to other views, one example being the concept of impulse variability (Schmidt et al., 1979), which would predict that end point variability increases with increasing distance. Given the assumption that W_0' is independent of A, the final term in Equation 20 becomes a constant, and we can derive a two-part model in suitable form for fitting to experimental data using multiple linear regression. Introduced in Experiment 1, this model is repeated here:

$$MT = a + b_1 \log_2 A - b_2 \log_2 W \qquad (5)$$

Note that the coefficient a in Equation 20 becomes b_1, and b becomes b_2. We feel the form of Equation 5 is more suitable for fitting to empirical data, and also provides an easier interpretation in terms of a high-level view of scaling effects. However, the units for the coefficients are confusing; strictly speaking b_1 and b_2 should be in milliseconds per log millimetre, even though they are clearly milliseconds per bit in the original formulation of Equation 20. Throughout, for purposes of discussion we have chosen ms/bit because this supports their interpretation as indices of performance, and helps for comparisons with the original concept of the index of performance in Fitts' law.

For the experiments in this thesis it is not possible to determine a value for W_0' in the pointing task, nor to verify the assumption that this value is independent of movement amplitude. Nevertheless, we found that Equation 5 provides the necessary detail for an accurate description of data from Experiments 1, 2 and 4. In addition, the relationship of the amplitude and width coefficients are helpful for understanding performance in terms of separable effects in the movement phases, as we will discuss below. Our choice of the two-part model is neither an indication that Fitts' law is incorrect, nor an attempt to account for a feature in our data such as a negative intercept. Rather, the choice is pragmatic since the asymmetry in our results was pronounced, and needed to be captured in an effective manner.

As introduced in Experiment 2, the parameter space for the model provides a method for visualising the magnitudes of the amplitude and width effects on movement time, and also for predicting effects when the scale of hand movements is altered. Figure 39 shows the model parameters for different conditions in Experiments 1, 2, and 4. The reader is reminded that although within-experiment comparisons of points on the graph are valid, an entirely different group of subjects participated in each experiment, making between experiment comparisons less accurate. In addition, we have included a point derived from data on three-dimensional pointing with a stylus in MacKenzie et al. (1987). The absolute position of a point in the space gives information about the index of performance derived from a particular data set. Considering the coefficients to be measures of sensitivity, the horizontal distance of a point from the origin show the sensitivity of movement time to changes in distance. Similarly, the height of a point on the graph gives the sensitivity of movement time to changes in target width. If these are roughly equal, Fitts' law gives a good account of the data, and it will fall somewhere on the line $b_1 = b_2$. For a point on the Fitts' law line, increased distance from the origin indicates increased sensitivity to changes in the index of difficulty, and hence a smaller index of performance.

Points below the Fitts' law line in the region $b_1 > b_2$ indicate a condition for which movement time is more sensitive to changes in distance than to changes in target width. All the points from Experiments 1, 2, and 4 fall in this region. The distance of a point from the Fitts' law line indicates how marked the asymmetry for A and W effects is in a data set. The greater the asymmetry, the greater will be the effect on movement time when the dimensions of hand movements for a task are scaled up or down. In the region below the line, reducing the distance (A) reduces the movement time, while the corresponding increase in movement time for reducing target width (W) in the same proportion is not as

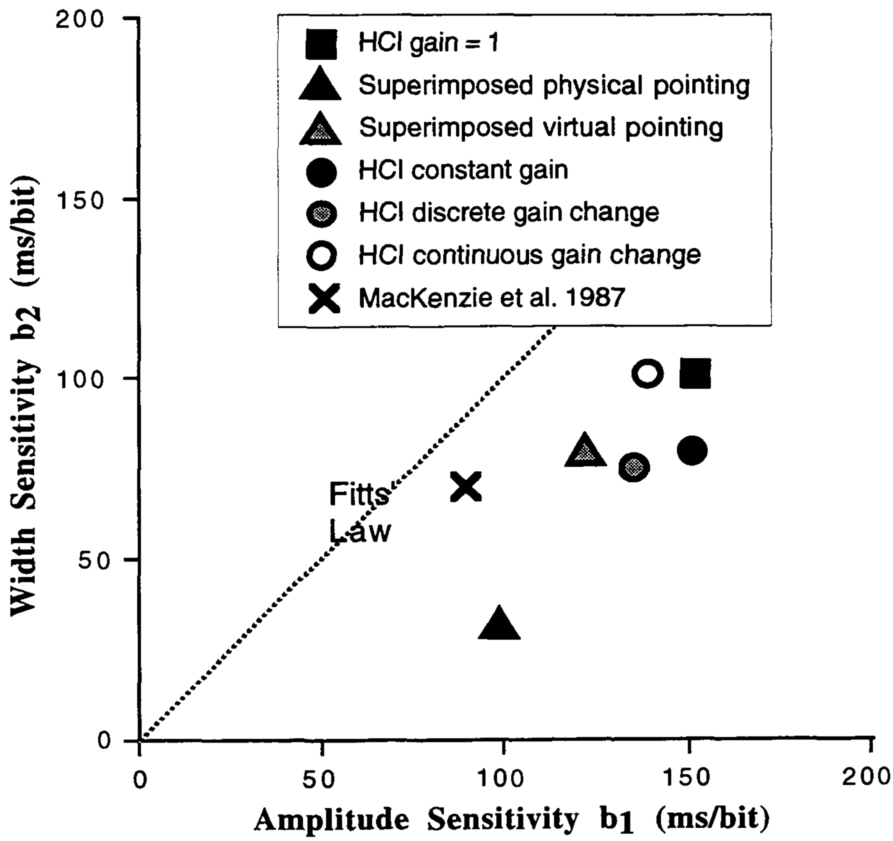

Figure 39. Points in the parameter space for models from Experiments 1, 2, and 4. Also included is a point from data in MacKenzie et al. (1987) for three-dimensional pointing to physical targets with a stylus.

great. In this case, the net effect of reducing A and W in concert is to reduce movement times. For points above the Fitts' law line, where $b_1 < b_2$, the opposite effect will be noted. In summary, for points which fall near the Fitts' law line scaling A and W will have little effect on movement times. Points which fall below indicate that "smaller is better", and will benefit from reducing the scale of hand movements, or equivalently increasing control-display gain. Points which fall above indicate that "bigger is better", and will benefit from increased scale, or reduced control-display gain. Thus, we can tell at a glance that there is evidence that performance in three-dimensional discrete pointing, whether at physical targets in the workspace or mediated by a computer and graphics display, will benefit from being scaled down. That is, target widths and distances for the hand should be reduced to the smallest practical values in order to reduce movement time.

Does this property of the data from our studies indicate an aggregate effect, or does it also apply at the level of individuals? We examined within subject means from all experiments and found that despite individual differences, all produced a point below the Fitts' law line on the graph. For example, we present Figure 40 illustrating the results of modelling individual subject data for the participants in Experiment 1. Each individual is more or less sensitive to changes in movement distance and accuracy constraints in the task, as indicated by the distance from the origin for their point in the parameter space. Despite these differences, all points fall well below the line to indicate that reducing the scale of hand movements in the work space will improve their performance.

Going back to Welford's derivation of Equation 20, another way to make comparisons is with respect to the contribution of movement phases to the model coefficients. The horizontal axis represents distance covering, which is more related to movement planning and features of the initial phase, whereas the vertical axis represents accuracy constraints, related to visuomotor control decelerating to the target. The relative positions of points can be interpreted in terms of comparative difficulties with each of the movement phases. For example, in Figure 39 the points for superimposed physical pointing and three-dimensional pointing in MacKenzie et al. (1987) are similar in the horizontal axis, with a sensitivity of about 100 ms/bit which agrees with Fitts' (1954) and other studies which followed his paradigm. The vertical displacement of the point for

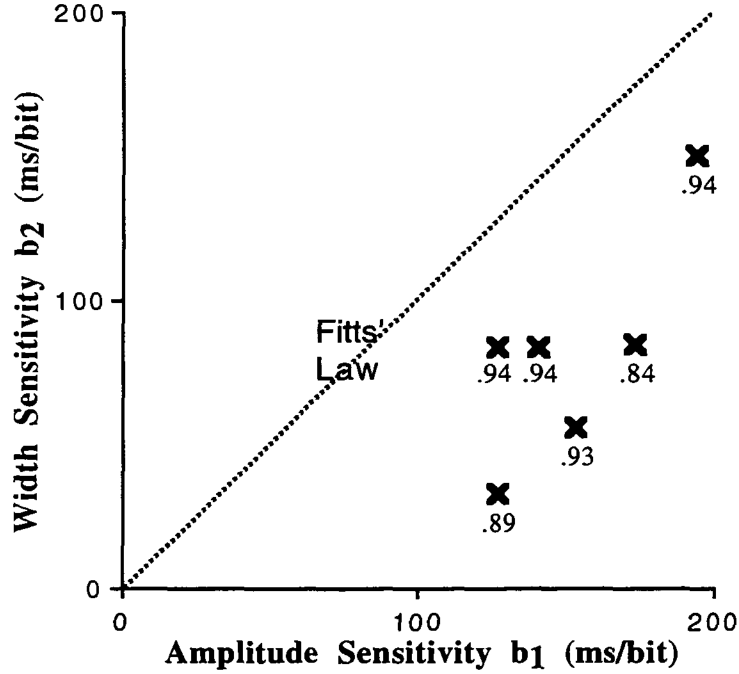

Figure 40. Points in the parameter space of the two part model for individual subjects in Experiment 1. Each point is annotated with the adjusted R^2 from the multiple regression which determined the model coefficients. Note that even at the individual level, they fall below the Fitts' law line, indicating that for all subjects reducing the scale of hand movements will reduce movement time for an equivalent index of difficulty.

MacKenzie et al. from the point for Experiment 2 indicates increased sensitivity to accuracy constraints, perhaps as a result of controlling a stylus, rather than the index finger for which small distal adjustments can be used to correct the final position.

All the data for pointing in the HCI setup falls to the right and above data for pointing at physical targets with vision of the hand (Figure 39). We propose that the points are to the right because it is more difficult to form a movement plan when the visual location of a target does not match the goal location for the hand. In this condition, planning must take into account the spatial translation from hand space to display space. The vertical displacement of the HCI data, relative to the point for physical pointing with the index finger indicates increased difficulty controlling deceleration to the target location. In Experiment 2 we suggested two explanations for this displacement: perhaps the lack of visual information about the height of the finger above the table makes it difficult to use contact with the surface to stop at the target location; or possibly the lag in the superimposed display increases the difficulty of decelerating to smaller targets.

The statements above are made with the caveat that accurate comparisons can only be done for data from a single group of subjects, since we used repeated measures experimental designs throughout. Nevertheless, we propose that this form of analysis may be used to compare and classify devices, tasks, and systems as to their effects on the two movement phases. In practical terms, model parameters can be determined by a small-scale experiment which combines two relatively extreme values for movement distance (A) with two extreme values of target width (W).

Given the detail with which data is reported in previous studies, and the fact that most HCI research has not examined simultaneous variations of A and W, it is difficult to determine whether the asymmetry in amplitude and width effects is typical of discrete positioning tasks in HCI for all devices. Results from Card et al. (1978) show similar effects for an isometric (force controlled) joystick for a less varied range of targets and distances. Data presented in Walker et al. (1993) also indicates little effect on movement times due to accuracy constraints compared to movement distance. In Experiment 4 we proposed a general strategy for improving performance with mediated pointing; the same strategy may also apply to other devices used for discrete positioning if, as a preliminary examination of other studies suggests, they also produce points below the Fitts' law line.

The relationship of the two part model parameters indicates that performance will be improved by reducing the scale of hand movements, or increasing the control-display gain. At present the noise, sensing resolution, and stability of pointing devices provided with commercial systems limits the maximum value of control-display gain which can be used because users still require fine positioning accuracy for targets as small as a single pixel. Pointing performance can be enhanced by improving the fidelity with which motions can be sensed, allowing the work space for hand movements to be reduced to the point where the limitations for locating small targets are due to human attributes, such as tremor, rather than properties of the sensing technology itself.

Predicting Target Location

The two part model of movement time is predictive in the sense that it characterises the average performance over a number of positioning movements under given conditions. Given our understanding of regularities in the kinematic features of movement trajectories, some of the computational techniques used to derive dependent measures for analysis in all four experiments may also be useful for real-time processing. For example, some of these techniques were employed in developing the system used in Experiment 3, in which perturbations were introduced at movement onset. The experiment system also used a limited form of gesture recognition to parse pointing movements, and allowed a subject to complete a block of self-paced trials without experimenter intervention. Here we propose a strategy for target prediction in real time to improve the response time of an interface.

In direct manipulation a user often indicates an action by reaching and selecting an object (target) on the display. Before an object is activated, there is often some significant amount of precomputation that could be done to help prepare for the subsequent interaction. For example, selecting a file icon in the traditional desktop metaphor often means that a word processing program must be launched and the file loaded. In this case, predicting the target location in advance would allow these actions to be initiated well in advance (several hundreds of milliseconds) of the actual object selection.

Typical movement trajectories for pointing at objects in the work space follow a relatively straight path to the goal location. The directional effects reported in Experiment 3 demonstrate that the movement direction (tangent to the spatial path) at the time of the first velocity peak was a function of the direction of the target from the starting point. Using a three point smoothing filter and three-point differentiation algorithm, it is possible to reliably determine the initial velocity peak within two or three samples of its occurrence. As demonstrated in previously in this chapter, the two important features of the initial velocity peak, timing and magnitude, are highly correlated, and are systematically related to the distance of a hand movement in computer pointing. Given an estimate of the direction and extent of the hand trajectory and a control-display mapping, a region for potential targets on the display can be selected.

The potential target must be converted to an accuracy constraint for the hand (W) by applying the inverse of the control-display mapping, or by some other heuristic in the case that the mapping is non-linear (and thus is not uniquely invertible). Our results suggest that W is highly correlated to the proportion of movement time after the first deceleration peak (%TAPD). The timing of this peak can be estimated within a few samples of its occurrence by calculating changes in the magnitude of the resultant velocity profile. Even given the lag of several samples required to identify this point in the movement trajectory in real time, our results indicate that about half of the movement time would still remain for typical A and W combinations. At this point, not only has a potential target been identified, but %TAPD combined with the absolute timing of the first deceleration peak can be used to predict movement time, and thus time remaining for precomputation before the target is reached. This time would vary from approximately 200 to 500 ms for the range of pointing conditions in all experiments in this thesis.

The accuracy of this technique will depend on a number of factors. Subjects demonstrated individual differences in the relationship of the timing and magnitude of the initial velocity peak to movement distance, but invariably demonstrated a strong correlation. For this technique to work, the systematic relationships of kinematic features to A and W would have to be identified for each individual user. These parameters could be determined in advance, or accumulated by an "intelligent" system in the background as it observes and measures key features of a user's interaction patterns. Variability inherent in individual movements, as opposed to mean values of the features used in our experimental analysis, will make exact prediction impossible. In impoverished virtual environments, however, there will be little choice of targets, making object prediction simpler, and if the variability were quantified for each user the technique would still provide a lower bound on the estimate of time left for precomputation.

Conclusion

We feel that the four experiments described in this thesis provide some insight as to how theories and models from motor control research can be applied to the study of human-computer interaction. In addition, some results have been produced which will be useful for analysis of human performance in HCI, and for the design of interactive systems. Part of the motivation for choosing a task in which a subject controlled a pointer on a two-dimensional display with a natural, three-dimensional hand movement, as opposed to a typical computer input device, was to examine how and why mediated pointing is difficult. The experiments were designed to gain an understanding of how the visual representation of a user's actions by a computer system affected movement planning and control, compared to pointing directly at targets in the work space. A somewhat surprising conclusion is that although the behaviour of the visual representation was affected by our experimental manipulations, the fine features of hand trajectories were consistently related to constraints for hand movement, not on visual information from the computer display. With this insight we realise that the choice of the task for the hand is what determined the results in our studies. Although the detailed analysis of the fine features of movement trajectories as represented on a computer display produces some interesting results, this cursor-based approach can also provide misleading answers because it does not measure directly the aspects of human behaviour which are coordinated and controlled in mediated pointing. Much more can be learned from a human-centred approach, and the application of useful theories and models from motor control to pointing on a computer display can best be done by taking a hand-centred approach to the study of human-computer interaction.

Bibiography

Abbs, J.H., and Gracco, V.L. (1984). Control of complex motor gestures: Orofacial muscle responses to load perturbations of lip during speech. Journal of Neurophysiology, 51(4), 705-723.

Abrams, R.A. (1994). Eye-hand coordination: Spatial localization after saccadic and pursuit eye movements. Journal of Motor Behaviour, 26(3), 215-224.

Abrams, R.J., and Pratt, J. (1993). Rapid aimed limb movements: differential effects of practice on component submovements. Journal of Motor Behaviour, 25(4), 288-298.

Adams, J.A (1989). Human Factors Engineering. Macmillan.

Adelstein, B.D., Johnston, E.R, and Ellis, S.R. (1992). A testbed for characterizing dynamic response of virtual environment spatial sensors. Proceedings of UIST '92, Monterey, November, 15.1-15.7.

Aglioti, S., DeSouza, J.F.X., and Goodale, M.A. (1995). Size-contrast illusions deceive the eye but not the hand. Current Biology, 5(6), 679-685.

Albert, A.E. (1982). The effect of graphics input devices on performance in a cursor positioning task. Proceedings of the Human Factors Society - Twenty-sixth Annual Meeting, 54-58.

Amaya, K., Calvert, T. and MacKenzie, C.L. (in preparation). Posture estimation of a human hand from a limited number of markers.

Arbib, M.A. (1980). Interacting schemas in motor control. In G. Stelmach and J. Requin (Eds.), Tutorials in Motor Behaviour (pp.71-81). Amsterdam: North Holland.

Arbib, M.A. (1981). Perceptual structures and distributed motor control. In V.B. Brooks (Ed.), Handbook of Physiology -The Nervous System, II Motor Control (pp. 1449-1480). Bethesda, MD: American Physiology Society, 1449-1480.

Arbib, M.A. (1985). Schemas for the temporal control of behaviour. Human Neurobiology,4, 63-72.

Arbib, M.A., Iberall, T., and Lyons, D. (1985). Coordinated control programs for movements of the hand. In A.W.Goodwin and I. Darian-Smith (Eds.), Hand Function and the Neocortex (pp. 111-129). Berlin: Springer-Verlag.

Arthur, K.W., Booth, K.S., and Ware, C. (1993). 3D task performance in fish tank virtual worlds. ACM Transactions on Information Systems, Special Issue on Virtual Worlds, 11(3), 239-265.

Baecker, R.M., Grudin, J., Buxton, W.A.S., and Greenberg, S. (Eds.). (1995). Readings in Human-Computer Interaction: Toward the Year 2000. San Francisco: Morgan Kaufmann.

Baecker, R.M. and Buxton, W.A.S. (Eds.). (1987). Readings in Human-Computer Interaction. Reading MA: Addison Wesley.

Berkinblit, M.B., and Feldman, A.G. (1988). Some problems of motor control. Journal of Motor Behaviour, 20(3), 369-373.

Blouin, J., Bard, C., Teasdale, N., and Fleury, M. (1993). On-line versus off-line control of rapid aiming movements. Journal of Motor Behaviour, 25(4), 275-279.

Bolas, M. (1994). Human factors in the design of an immersive display. IEEE Computer Graphics and Applications, 28(1), 55-59.

Bolt, R.A. (1987). The integrated multi-modal interface. Transactions of the Institute of Electronics, Information and Communication Engineers (Japan), J70-D(11), 2017-2025.

Bolt, R.A., and Herranz, E. (1992). Two-handed gesture in multi-modal natural dialog. Proceedings of UIST '92, Monterey, California, November.

Bootsma, R.J., Marteniuk, R.G., MacKenzie, C.L., and Zaal, F.T.J.M. (1994). The speed-accuracy trade-off in manual prehension: effects of movement amplitude, object size and object width on kinematic characteristics. Experimental Brain Research, 98, 535-541.

Boritz, J., Booth, K.S., and Cowan, W.B. (1991). Fitts' law studies of directional mouse movement. Proceedings of Graphics Interface '91, Toronto, 216-223.

Bowen, H.M. & Guinness, G.V. (1965). Preliminary experiments on keyboard design for semi-automatic mail sorting. Journal of Applied Psychology, 49, 194-198.

Brenner, E., and Smeets, J.B.J. (1995). Moving one's finger to a visually specified position: target orientation influences the finger's path. Experimental Brain Research, 105, 318-320.

Brenner, E., and Smeets, J.B.J. (1994). Different frames of reference for position and motion. Naturwissenschaften, 81, 30-32.

Bruner, J. (1966). Towards a Theory of Instruction. New York: W.W Norton and Company.

Bryson, S. (1995). Effects of lag and frame rate on various tasks. Proceedings of SPIE '95, 16.1-16.11.

Bryson, S. (1995). Measurement and calibration of static position data from 3D trackers. Proceedings of SPIE '95, 13.1-13.11.

Buck, L. (1980). Motor performance in relation to control-display gain and target width. Ergonomics, 23, 579-589.

Buck, L. (1979). The sensory basis of target location. Quarterly Journal of Experimental Psychology, 31, 111-120.

Buxton, W. (1990) The natural language of interaction: A perspective on nonverbal dialogues. In Laurel, B. (Ed.), The Art of Human-Computer Interface Design (pp. 405-416). New York: Addison-Wesley.

Cannon, D.J. (1994). Experiments with a target-threshold control theory model for deriving Fitts' law parameters for human-machine systems. IEEE Transactions on Systems, Man, and Cybernetics, 24(8), 1089-1098.

Card, S.K., English, W.K, and Burr, B.J. (1978). Evaluation of mouse, rate-controlled joystick, step keys, and text keys for text selection on a CRT. Ergonomics, 21(8), 601-613.

Card, S.K., MacKinlay, J.D., and Robertson, G.G. (1990). The design space of input devices. Proceedings of CHI '90, 117-124.

Card, S.K., MacKinlay, J.D., and Robertson, G.G. (1991). A morphological analysis of the design space of input devices. ACM Transactions on Information Systems, 9(2), 99-122.

Card, S.K., Moran, T.P., and Newell, A. (1980). The keystroke-level model for user performance time with interactive systems. Communications of the ACM, 23(7), 396-410.

Card, S.K., Moran, T.P., and Newell, A. (1983). The Psychology of Human-computer Interaction. Hillsdale NJ: Lawrence Erlbaum.

Carlton, L. G. (1992). Visual processing time and the control of movement. In L. Proteau and D. Elliot (Eds.), Vision and Motor Control (pp. 3-31). Amsterdam: Elsevier Science Publishers.

Carnahan, H. and Marteniuk, R.G. (1991). The temporal organization of hand, eye, and head movements during reaching and pointing. Journal of Motor Behaviour, 23(2), 109-119.

Carnahan, H., Goodale, M.A., and Marteniuk, R.G. (1993). Grasping versus pointing and the differential use of visual feedback. Human Movement Science, 12, 219-234.

Carroll, J.M. Ed. (1991). Designing Interaction: Psychology at the Human-Computer Interface. Cambridge: Cambridge University Press.

Castiello, U, Paulignan, Y., and Jeannerod, M. (1991). Temporal dissociation of motor responses and subjective awareness. Brain, 114, 2639-2655.

Charteris, J., Cooper, L.A., and Bruce, J.R. (1976). Human Kinetics: A conceptual model for studying human movement. Journal of Human Movement Studies, 2, 233-238.

Chung, J.C. (1992). A comparison of head-tracked and non-head-tracked steering modes in the targeting of radiotherapy treatment beams. Proceedings of CHI '92, 193-196.

Clark, J. (1995). The easy-access mac. MacUser, January, 119-121.

Cohen, O., Meyer, S., and Nilsen, E. (1992). Studying the movement of high tech rodentia: pointing and dragging. Proceedings of CHI '92, 135-136.

Cole, K.J. and Abbs, J.H. (1983). Intentional responses to kinesthetic stimuli in orofacial muscles: implications for the coordination of speech movements. Journal of Neuroscience, 3(12), 2660-2669.

Cole, K.J. and Abbs, J.H. (1987). Kinematic and electromyographic responses to perturbation of a rapid grasp. Journal of Neurophysiology, 57(5), 1498-1510.

Conrad, R. & Longman, D.J.A. (1965). Standard typewriter versus chord keyboard - an experimental comparison. Ergonomics, 8, 77-88.

Cooper, L.A., and Munger, M.P. (1993). Extrapolating and remembering positions along cognitive trajectories: uses and limitations of analogies to physical motion. In N. Eilan, R. McCarthy, B. Brewer (Eds.), Spatial Representation (pp. 112-131). Oxford: Blackwell.

Craik, K.J.W. (1947). Theory of the human operator in control systems: I. The operator as an engineering system. British Journal of Psychology, 38, 56-61.

Crossman, E.R.F.W., and Goodeve, P.J. (1983). Feedback and control of hand movement and Fitts' law. Quarterly Journal of Experimental Psychology, 35A, 251-278.

Cunningham, H.A., and Vardi, I. (1990). A vector-sum process produces curved aiming paths under rotated visual-motor mappings. Biological Cybernetics, 64, 117-128.

Cunningham, H.A. (1989). Aiming error under transformed spatial mappings suggests a structure for visual-motor maps. Journal of Experimental Psychology, 15(3), 493-506.

Cunningham, H.A., and Welch, R.B. (1994). Multiple concurrent visual-motor mappings: implications for models of adaptation. Journal of Experimental Psychology: Human Perception and Performance, 20(5), 987-999.

Cutkosky, M.R. and Hyde, J.M. (1993). Manipulation control with dynamic tactile sensing. Proceedings of International Symposium on Robotics Research, Hidden Valley, PA.

De Graef, P., and D'Ydewalle, G. (1992). Local and global constraints on the identification of objects in scenes. Canadian Journal of Psychology, 46(2), 489-508.

Deering, M. (1995). HoloSketch: A virtual reality sketching/animation tool. ACM Transactions on Computer-Human Interaction, 2(3), 220-238.

Deering, M. (1992). High resolution virtual reality. Computer Graphics, 26, 195-202.

Deutsch, D. (1983). The generation of two isochronous sequences in parallel. Perception and Psychophysics, 34(4), 331-337.

Diaper, D. (1989). Task Analysis for Human-Computer Interaction. New York: Wiley.

Dillon, T.W. (1995). Mapping the discourse of HCI researchers with citation analysis. SIGCHI Bulletin, 27(4), 56-62.

Dix, A., Finlay, J., Abowd, G., and Beale, R. (1993). Human-Computer Interaction. New York: Prentice Hall.

Douglas, S.H., and Mithal, A.K. (1994). The effect of reducing homing time on the speed of a finger-controlled isometric pointing device. Proceedings of CHI '94, 411-416.

Dretske, F. (1990). Seeing, believing and knowing. In D.N. Osherson, M.S. Kosslyn, & J.M. Hollerbach (Eds.) Visual Cognition and Action (pp. 129-148). Cambridge, Mass: MIT Press.

Drucker, S.M., and Zeltzer, D. (1994). Intelligent camera control in a virtual environment. Proceedings of Graphics Interface '94, May, Banff, Alberta, Canada.

Earnshaw, R.A., Gigante, M.A., and Jones, H. (1993). Virtual Reality Systems. New York: Addison-Wesley.

Elliot, D., Carson, R.G., Goodman, D., and Chua, R. (1991). Discrete vs. continuous control of manual aiming. Human Movement Science, 10, 393-418.

Elliott, D., and Roy, E.A. (1981). Interlimb transfer after adaptation to visual displacement: patterns predicted from the functional closeness of limb neural control centres. Perception, 10, 383-389.

Elliott, D., Roy, E.A., and Allard, F. (1982). The influence of a secondary task and type of feedback on adaption to lateral displacement of the visual array. Psychological Research, 44, 343-353.

Ellis, S.R. (1994). What are virtual environments? IEEE Computer Graphics and Applications, 28(1), 17-22.

English, W.K., Englebart, D.C., and Berman, M.L. (1967). Display-selection techniques for text manipulation. IEEE Transactions on Human Factors in Electronics, 8, 5-15.

Ewing, J., Mehrebanzad, S., Sheck, S., and Scheiderman, B. (1986). An experimental comparison of a mouse and arrow-jump keys for an interactive encyclopedia. International Journal of Man-Machine Studies, 23(1), 29-45.

Feldman, A.G. (1986). Once more on the equilibrium-point hypothesis for motor control. Journal of Motor Behaviour, 18(1), 17-54.

Fitts, P.M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. Journal of Experimental Psychology, 47(6), 381-391.

Fitts, P.M., and Peterson, J.P. (1964). Information capacity of discrete motor responses. Journal of Experimental Psychology, 67(2), 103-112.

Flach, J.M., Hagen, B.A, O'Brien, D., and Olson, W.A. (1987). Alternative displays for discrete movement control. Report of Dept. of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign.

Foley, J.D., van Dam, A., Feiner, S.K., and Hughes, J.F. (1991). Computer Graphics: Principles and Practice. New York: Addison-Wesley.

Foley, J.D., Wallace, V.W., and Chan, P. (1984). The human factors of graphics interaction techniques. IEEE Computer Graphics and Applications, November, 13-48.

Fowler, B., Taylor, M., and Prolier, G. (1987). The effects of hypoxia on reaction time and movement time components of a perceptual-motor task. Ergonomics, 1987, 30(10), 1475-1485.

Fowler, B., White, P.L., Holness, D.E., Wright, G.P., and Ackles, K.N. (1987). The effects of inert gas narcosis on the speed and accuracy of movement. Ergonomics, 1982, 25(3), 203-212.

Friedman, M., Starner, T., and Pentland, A. (1992). Device synchronization using an optimal linear filter. Proceedings of CHI '92, 57-62.

Gentilucci, M., Daprati, E., Toni, I., Chieffi, S., and Saetti, M.C. (1995). Unconscious updating of the grasp program. Experimental Brain Research, 105, 291-303.

Gentilucci, M., Chieffi, S., Scarpa, M., and Castiello, U. (1992). Temporal coupling between transport and grasp components during prehension movements: effects of visual perturbation. Behavioural Brain Research, 47, 71-82.

Gentilucci, M., Jeannerod, M., Tadary, B., and Decety, J. (1994). Associating visual and kinesthetic coordinates during pointing movements. Experimental Brain Research, 102, 359-366.

Gentner, D.R., Grudin, J. & Conway, E. (1980). Finger movements in transcription typing. (Technical Report no. 8001). San Diego: University of California, Center for Human Information Processing.

Georgopoulos, A.P., Kalaska, J.F. and Massey, J.T. (1981). Spatial trajectories and reaction times of aimed movements: Effects of practice, uncertainty and change in target location. Journal of Neurophysiology, 46, 725-743.

Ghilardi, M.F., Gordon, J. and Ghez, C. (1995). Learning a visuomotor transformation in a local area of work space produces directional biases in other areas. Journal of Neurophysiology, 73(6), 2535-2539.

Gibson, J.J. (1966). The senses considered as perceptual systems. Boston: Houghton Mifflin.

Gibson, J.J. (1979). The ecological approach to visual perception. Boston: Houghton Mifflin.

Gillan, D.J., Holden, K., Adam, S., Rudisill, M., and Magee, L. (1990). How does Fitts' law fit pointing and dragging? Proceedings of CHI '90, 227-234.

Gillespie, B. and Cutkosky, M. (1993). Interactive dynamics with haptic display. Advances in Robotics, Mechatronics, and Haptic Interfaces DSC-Vol. 49, 65-72.

Gobel, M., Luczak, H., Springer, J., Hedicke, V., and Rotting, M. (1995). Tactile feedback applied to computer mice. International Journal of Human-Computer Interaction, 7(1), 1-24.

Goodale, M.A., Milner, A.D., Jakobson, L.S., and Carey, D.P. (1991). A neurological dissociation between perceiving objects and grasping them. Nature, 349, 154-156.

Goodale, M.A., Pelisson, D., and Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. Nature, 320, 748-750.

Goodwin, N.C. (1975). Cursor positioning on an electronic display using lightpen, lightgun, or keyboard for three basic tasks. Human Factors, 17(3), 289-295.

Graham, E.D. (1996). Determining perceived lag in an augmented reality system. Manuscript in preparation.

Graham, E.D. (in press). Virtual pointing on a computer display: Non-linear control-display mappings. Proceedings of Graphics Interface '96.

Graham, E.D. and MacKenzie, C.L. (in press). Physical versus virtual pointing. Proceedings of CHI '96, Vancouver, April.

Graham, E.D. and MacKenzie, C.L. (1995a). Pointing on a computer display. CHI '95 Conference Companion, Denver, Colorado, May, pp. 314-315.

Graham, E.D. and MacKenzie, C.L. (1995b). Mapping pointing movements in hand space to a computer display. Journal of Sport and Exercise Psychology, 17(supp.), p. 56.

Graham, E.D., MacKenzie, C.L., Calvert, T., & Crawford, J. (1993). Human-computer interface design for effective manual communication. Abstract and poster, Proceedings of Third Annual IRIS/PRECARN Conference, Ottawa, June, p.98.

Guthrie, W. (1995). An overview of portable GUI software. SIGCHI Bulletin, 27(1), 55-69.

Haggard, P., and Wing, A. (1991). Remote responses to perturbation in human prehension. Neuroscience Letters, 122, 103-108.

Haller, R., Mutschler, H., and Voss, M. (1984). Comparison of input devices for correction of typing errors in office systems. Proceedings of INTERACT 84, 218-223.

Hammontree, M.L., Hendrickson, J.J., and Hensley, W. (1992). Integrated data capture and analysis tools for research and testing on graphical interfaces. Proceedings of CHI '92, 431-432.

Held, R., and Durlach, N. (1987). Telepresence, time delay, and adaptation. NASA Conference Publication 10032, 28-1 to 28-16.

Held, R. (1965). Plasticity in sensory-motor systems. Scientific American, 213(5), 84-94.

Henderson, J.M. (1992). Object identification in context: The visual processing of natural scenes. Canadian Journal of Psychology, 46(3), 319-341.

Hill, W.C., and Hollan, J.D. (1992). Pointing and visualization. Proceedings of CHI '92, 665-666.

Hinckley, K., Pausch, R., Goble, J.C., and Kassell, N.F. (in press). Design hints for spatial input. Proceedings of UIST '94.

Hirota, K., and Hirose, M. (1993). Development of surface display. IEEE Virtual Reality Annual International Symposium, 256-262.

Hocherman, S. (1993). Proprioceptive guidance and motor planning of reaching movements to unseen targets. Experimental Brain Research, 95, 349-358.

Hoffman, E.R., Tsang, K.K., and Mu, A. (1995). Data-entry keyboard geometry and keying movement times. Ergonomics, 38(5), 940-950.

Hogan, N. (1984). An organizing principle for a class of voluntary movements, Journal of Neuroscience, 4(11), 2745-2754.

Humphrey, G.K., and Khan, S.C. (1992). Recognizing novel views of three-dimensional objects. Canadian Journal of Psychology, 46(2), 170-190.

Iberall, T. (1987). A neural model of human prehension. Unpublished doctoral dissertation, University of Massachusetts.

Jackson, S.R., Jackson, G.M., Harrison, J., Henderson, L. and Kennard, C. (1995). The internal control of action and Parkinson's disease: A kinematic analysis of visually-guided and memory-guided prehension movements. Experimental Brain Research, 105, 147-162.

Jackson, S.R., Jackson, G.M., Rosicky, J., and Posner, M.I. (in preperation). Are non-relevant objects represented in working memory?: The effect of non-target objects on reach and grasp kinematics.

Jacob, R.K., and Sibert, L.E. (1992). The perceptual structure of multidimensional input device selection. Proceedings of CHI '92, 211-218.

Jacob, R.K., Sibert, L.E., McFarlane, D.C., and Mullen, M.P. (1994). Integrality and separability of input devices. ACM Transactions on Computer-Human Interaction, 1(1), 3-26.

Janin, A.L., Mizell, D.W., and Caudell, T.P. (1993). Calibration of head-mounted displays for augmented reality applications. IEEE Virtual Reality Annual International Symposium, 246-255.

Jeannerod, M., Arbib, M.A., Rizzolatti, G., and Sakata, H. (1995). Grasping objects: the cortical mechanisms of visuomotor transformation. Trends in Neurosciences, 18(7), 314-320.

Jeannerod, M. (1984). The timing of natural prehension movements. Journal of Motor Behaviour, 16, 235-254.

Jeannerod, M. (1986). The formation of finger grip during prehension. A cortically mediated visuomotor pattern. Behavioural Brain Research, 19, 99-116.

Jellinek, H.D., and Card, S.K. (1990). Powermice and user performance. Proceedings of CHI '90, 213-220.

Johansson R.S., Riso, R., Hager, C. and Backstrom, L. (1992). Somatosensory control of precision grip during unpredictable pulling loads: Changes in load force amplitudes. Experimental Brain Research, 89, 181-191.

Johansson, R.S., Hager, C., and Riso, R. (1992). Somatosensory control of precision grip during unpredictable pulling loads: Changes in load force rate. Experiment Brain Research, 89, 192-203.

Kabbash, P. and Buxton, W. (1995). The "prince" technique: Fitts law and selection using area cursors. Proceedings of CHI '95, Denver, Colorado, 273-279.

Kabbash, P., Buxton, W., and Sellen, A. (1994). Two-handed input in a compound task. Proceedings of CHI '94, Boston, Mass, April, 417-423.

Kalawsky, R.S. (1993). Critical aspects of visually coupled systems. In R.A. Earnshaw, M.A. Gigante, and H. Jones (Eds.), Virtual Reality Systems (pp. 203-212). London: Academic Press.

Kalish, M.L. (1994). Idiosyncratic errors in visually directed reaching. Journal of Motor Behaviour, 26(3), 296-300.

Kamienska-Zyla, M. (1993). Ergonomics evaluation of the work of VDT operators in Poland. Applied Ergonomics, 24(6), 432-433.

Kantowitz, B.H., and Sorkin, R.D (1990). Human Factors: Understanding People-system Relationships. New York: John Wiley & Sons.

Karat, J., McDonald, J., and Anderson, M. (1984). A comparison of selection techniques: Touch panel, mouse and keyboard. Proceedings of INTERACT 84, 149-153.

Kay, A. (1990). User interface: a personal view. In Laurel, B. (Ed.) The Art of Human-Computer Interface Design (pp. 191-208). New York: Addison-Wesley.

Keating, J.G., Martin, T.A., Mink, J.W, and Thatch, W.T. (1995). Gaze direction plays a role in guiding a natural but not an instrumented wrist movement. Abstracts for Neuroscience, November, San Diego, p. 422.

Keele, S.W., and Posner, M.I. (1968). Processing of visual feedback in rapid movements. Journal of Experimental Psychology, 71(1), 155-158.

Kennedy, R.S., Hettinger, L.J, and Lilienthal, M.G.(1987). Simulator sickness. In G. Crampton (Ed.) Motion and Space Sickness (pp. 317-341). Boca Raton FL: CRC Press.

Kenyon, R.V., and Afenya, M.B. (1995). Training in virtual and real environments. [WWW Document], URL: http://cubs.eecs.uic.edu/ppt.html.

Kinchla, R.A. (1992). Attention. Annual Review of Psychology, 43, 711-742.

Kirsh, D., and Maglio, P. (1994). On distinguishing epistemic from pragmatic action. Cognitive Science, 18, 513-549.

Klatzky, R., McCloskey, B., Doherty, S., Pellegrino, J., and Smith, T. (1987). Knowledge about hand shaping and knowledge about objects. Journal of Motor Behaviour, 19, 187-213.

Kim, W.S., Ellis, S.R., Mitchell, E.T., Hannaford, B. and Stark, L.W. (1987). Quantitative evaluation of perspective and stereoscopic displays in three-axis manual tracking tasks. IEEE Transactions on Systems, Man, and Cybernetics, SMC-17(1), 61-72.

Komilis, E., Pelisson, D., and Prablanc, C. (1993). Error processing in pointing at randomly feedback-induced double-step stimuli. Journal of Motor Behaviour, 25(4), 299-308.

Koons, D.B. (1994). Capturing and interpreting multi-modal descriptions with multiple representations. Presented at AAAI Symposium on Intelligent Multi-Modal Multi-Media Interface Systems, Stanford, March.

Kosslyn, S.M. (1990). Mental Imagery. In D.N. Osherson, M.S. Kosslyn, & J.M. Hollerbach (Eds.) Visual Cognition and Action (pp. 73-98). Cambridge, Mass: MIT Press.

Kozak, J.J., Hancock, P.A., Arthur, E.J., and Chrysler, S.T. (1993). Transfer of training from virtual reality. Ergonomics, 36(7), 777-784.

Langolf, G.D., Chaffin, D.B., and Foulke, J.A. (1976). An investigation of Fitts' law using a wide range of movement amplitudes. Journal of Motor Behaviour, 8, 113-128.

Latta, J.N., and Oberg, D.J. (1994). A conceptual virtual reality model. IEEE Computer Graphics and Applications, 28(1), 23-29.

Laurel, B. (Ed.). (1990). The Art of Human-Computer Interface Design. New York: Addison-Wesley.

Lederman, S.J., and Klatzky, R.L. (1987). Hand movements: A window into haptic object recognition. Cognitive Psychology, 19, 342-368.

Lederman, S.J., and Klatzky, R.L. (1990). Hand classification of common objects: knowledge-driven exploration. Cognitive Psychology, 22, 421-459.

Liang, J., Shaw, C., & Green, M. (1991). On temporal-spatial realism in the virtual reality environment. Proceedings of UIST '91, Hilton Head, November, 14.1-14.7.

Lippman, A. (1981). And seeing through your hand. Proceedings of SID, 22(2), 103-107.

Long, J. and Whitefield, A. (1989). Cognitive Ergonomics and Human-Computer Interaction. Cambridge: Cambridge University Press.

Machover, C., and Tice, S.E. (1994). Virtual Reality. IEEE Computer Graphics and Applications, 28(1), 15-16.

MacKenzie, C.L. (1985). Structural constraints on timing in human finger movements. In D. Goodman, R. B. Wilberg, & I. M. Franks (Eds.), Differing Perspectives in Motor Learning and Control (pp. 157-173). Elsevier Science Publishers: North Holland.

MacKenzie, C.L. (1992). Constraints, phases and sensorimotor processing in prehension. In G.E. Stelmach and J. Requin (Eds.), Tutorials in Motor Behaviour II (pp.371-398). Elsevier Science, 371-398.

MacKenzie, C.L., and Iberall, T., (1994). The Grasping Hand. Elsevier Science: Amsterdam.

MacKenzie, C.L., Marteniuk, R.G., Dugas, C., and Eickmeier, B. (1987). Three-dimensional movement trajectories in Fitts' task: Implications for motor control. The Quarterly Journal of Experimental Psychology, 39A, 629-647.

MacKenzie, S. (1989). A note on the information-theoretic basis for Fitts' law. Journal of Motor Behaviour, 21,(3), 91-139.

MacKenzie, S. (1991). Fitts' law as a research and design tool in human-computer interaction. Human-Computer Interaction, 7, 91-139.

MacKenzie, S. (1992). Movement time prediction in human-computer interfaces. Proceedings of Graphics Interface '92, 140-150.

MacKenzie, S., and Buxton, W. (1992). Extending Fitts' law to two-dimensional tasks. Proceedings of CHI '92, 219-226.

MacKenzie, S., Sellen, A., and Buxton, W. (1991). A comparison of input devices in elemental pointing and dragging tasks. Proceedings of CHI '91, 161-166.

MacKnik, S.L., Fisher, B.D., and Bridgeman, B. (1991). Flicker distorts visual space constancy. Vision Research, 31(12), 2057-2064.

Maggioni, C. (1993). A novel gestural input device for virtual reality. IEEE Virtual Reality Annual International Symposium, 118-124.

Marr, D. (1982). Vision. New York: W. H. Freeman.

Marteniuk, R., and MacKenzie, C. (1990). Invariance and variability in human prehension: Implications for theory development. In M. Goodale (Ed.), Vision and Action: The Control of Grasping (pp. 49-64). Norwood, NJ: Ablex, 49-64.

Marteniuk, R., Leavitt, J., MacKenzie, C., and Athenes, S. (1990). Functional relationships between grasp and transport components in a prehension task. Human Movement Science, 9, 149-176.

Marteniuk, R., MacKenzie, C., Jeannerod, M., Athenes, S., and Dugas, C. (1987). Constraints on human arm movement trajectories. Canadian Journal of Psychology, 41, 365-378.

Mathieson, B. (1989). The effects of hypoxia and brightness on the control of aimed movements: Implications for a two process model. Unpublished master's thesis, York University, Toronto.

McKenna, M. (1992). Interactive viewpoint control and three-dimensional operations. Proceedings of CHI '92, 53-56.

Meilach, D.Z. (1993). 3-D stereo systems. Computer Pictures, November-December, 44-51.

Meyer, K., Applewhite, H.L., and Biocca, F.A. (1992). A survey of position trackers. Presence, 1(2), 173-200.

Michaels, C.F. & Carello, C. (1981). Direct perception. Englewood Cliffs, NJ: Prentice Hall.

Mitchell, D.E., and O'Hagan, S. (1972). Accuracy of stereoscopic localization of small line segments that difffer in size or orientation for the two eyes. Vision Research, 12, 437-454.

Murakami, T., and Nakajima, N., (1994). Direct and intuitive input device for 3-D shape deformation. Proceedings of CHI '94, 465-470.

Nielsen, T.I. (1963). Volition: a new experimental approach. Scandinavian Journal of Psychology, 4, 225-230.

Niguma, G. and Heinrichs, M. (1995). Virtual Hand. Unpublished project report, Faculty of Applied Science, Simon Fraser University, Burnaby, BC.

Norman, D.A. (1988). The Psychology of Everday Things. New York: Doubleday.

Norman, D.A. and Draper, S.W. Eds. (1986). User Centered System Design. Hillsdale, NJ: Lawrence Erlbaum.

Noyes, J. (1983). Chord Keyboards. Applied Ergonomics, 14, 55-59.

O'Regan, J.K. (1992). Solving the real mysteries of visual perception: The world as an outside memory. Canadian Journal of Psychology, 46(3), 461-488.

Olson, W.A., and Flach, J.M. (1986). An exploration of a Fitts' law task on a logarithmically scaled screen. Proceedings of the Twenty-Second Conference on Manual Control, Dayton, Ohio: Air Force Wright Aeronautical Laboratories, 165-173.

Osherson, D.N., Kosslyn, S.M., and Hollerbach, J.M. (Eds.) (1990). Visual Cognition and Action. Cambridge, Mass: MIT Press.

Paillard, J. (1980). The multichanneling of visual cues and the organization of a visually guided response. In G.E. Stelmach and J. Requin (eds.), Tutorials in Motor Behaviour (pp. 259-277). Amsterdam: North Holland.

Paillard, J. (1982). The contribution of peripheral and central vision to visually guided reaching. In Analysis of Visual Behaviour (pp. 367-385). Cambridge, Mass: MIT Press.

Paulignan, Y., Jeannerod, M., MacKenzie, C., and Marteniuk, R. (1991a). Selective perturbation of visual input during prehension movements, 2. The effects of changing object position. Experimental Brain Research, 87, 407-420.

Paulignan, Y., MacKenzie, C., Marteniuk, R., and Jeannerod M. (1990). The coupling of arm and finger movements during prehension. Experimental Brain Research, 79, 431-435.

Paulignan, Y., MacKenzie, C., Marteniuk, R., and Jeannerod M. (1991b). Selective perturbation of visual input during prehension movements, 1. The effects of changing object position. Experimental Brain Research, 83, 502-512.

Pausch, R., Vogtle, L., and Conway, M. (1992). One dimensional motion tailoring for the disabled: A user study. Proceedings of CHI '92, 405-411.

Pelisson, D., Prablanc, C., Goodale, M.A., and Jeannerod, M. (1986). Visual control of reaching movements without vision of the limb. II. Evidence of fast unconscious processes correcting the trajectory of the hand to the final position of a double-step stimulus. Experimental Brain Research, 62, 303-311.

Polit, A., and Bizzi, E. (1979). Characteristics of motor programs underlying arm movements in monkeys. Journal of Neurophysiology, 42(1), 183-194.

Popovic, M., and Popovic, D. (1994). A new approach to reaching control for tetraplegic subjects. Journal of Electromyography and Kinesiology, 4, 242-253.

Prablanc, C., Echallier, J.F., Komilis, E., and Jeannerod, M. (1979). Optimal response of eye and hand motor systems in pointing at a visual target. II Static and dynamic visual cues in the control of hand movements. Biological Cybernetics, 35, 183-187.

Prablanc, C., Pelisson, D., and Goodale, M.A. (1986). Visual control of reaching movements without vision of the limb: I. Role of retinal feedback of target position in guiding the hand. Experimental Brain Research, 62, 293-302.

Rabbitt, P.M.A, Vyas, S.M., and Fearnley, S. (1975) Programming sequences of complex responses. In P.M.A. Rabbitt and S. Dornic (Eds.), Attention and Performance V (pp. 395-417). New York: Academic Press.

Ratz, H.C. & Ritchie, D.K. (1961). Operator performance on a chord keyboard. Journal of Applied Psychology, 45, 303-308.

Rayner, K. & Pollatsek, A. (1992). Eye movements and scene perception. Canadian Journal of Psychology, 46(3), 342-376.

Reason, J. (1990). Human Error. Cambridge: Cambridge University Press.

Redding, G.M., and Wallace, B. (1993). Adaptive coordination and alignment of eye and hand. Journal of Motor Behaviour, 25(2), 75-88.

Redding, G.M., and Wallace, B. (1994). Effects of movement duration and visual feedback on visual and proprioceptive components of prism adaptation. Journal of Motor Behaviour, 26(3), 257-266.

Reichman, R. (1986). Communication paradigms for a window system. In D.A. Norman and S.W. Draper (Eds.), User Centered System Design (pp. 285-314). Hillsdale, NJ: Lawrence Erlbaum.

Robinett, W. and Rolland, J.P. (1993). A computational model for the stereoscopic optics of a head-mounted display. In R.A. Earnshaw, M.A. Gigante, and H. Jones (Eds.), Virtual Reality Systems (pp. 51-75). London: Academic Press.

Rose, C. (1985). Inside Macintosh, Volume II. Reading, MA: Addison Wesley.

Rosenbaum, D., Marchak, F., Barnes, H., Vaughan, J., Slotta, J., and Jorgenson, M. (1990). Constraints for action selection: overhand versus underhand grips. In M. Jeannerod (Ed.), Attention and Performance XIII (pp. 321-342). Hillsdale, NJ: Erlbaum.

Rosenberg, L.B. (1993). The effect of interocular distance upon operator performance using stereoscopic displays to perform virtual depth tasks. IEEE Virtual Reality Annual International Symposium, 27-30.

Rosenberg, L.B., (1993). Virtual fixtures: Perceptual tools for telerobotic manipulation. IEEE Virtual Reality Annual International Symposium, 76-82.

Rossetti, Y., Stelmach, G., Desmurget, M., Prablanc, C., & Jeannerod, M. (1994). The effect of viewing the static hand prior to movement onset on pointing kinematics and variability. Experimental Brain Research, 101, 323-330.

Rossetti, Y., Tadary, B., & Prablanc, C. (1994). Optimal contributions of head and eye postions to spatial accuracy in man tested by visually directed pointing. Experimental Brain Research, 97, 487-496.

Rubine, D. (1991). Specifying gestures by example. Computer Graphics, 25(4), 329-337.

Rummelhart, D.E. & Norman, D.A (1982). Simulating a skilled typist: A study of skilled cognitive motor performance. Cognitive Science, 6, 1-36.

Scheifler, B., and Krikorian, D. (1988). Manual page for xset, X-Windows Version 11, Massachusetts Institute of Technology.

Schmandt, C. (1982). Interactive three-dimensional computer space. SPIE Processing and Display of Three-Dimensional Data, 367, 155-160.

Schmidt, R.A., Zelaznik, H., Hawkins, B., Frank, J.S., and Quinn, J.T. (1979). Motor output variability: A theory for the accuracy of rapid motor acts. Psychological Review, 86(5), 415-451.

Schneiderman, B. (1982). The future of interactive systems and the emergence of direct manipulation. Behaviour and Information Technology, 1, 237-256.

Schneiderman, B. (1992). Designing the User Interface: Strategies for Human-Computer Interaction. Reading MA: Addison Wesley.

Schutz, R.W., Smoll, F.L., and Gessaroli, M.E. (1983). Multivariate statistics: a self-test and guide to their utilization. Research Quarterly for Exercise and Sport, 54(3), 255-263.

Sears, A., and Shneiderman, B. (1991). High precision touchscreens: design strategies and comparisons with a mouse. International Journal of Man-Machine Studies, 34, 593-613.

Shackel, B. and Richardson, S.J. (Eds.). (1991). Human Factors for Informatics Usability. Cambridge: Cambridge University Press.

Shepp, B.E., and Ballesteros, S. (Eds.). (1989). Object Perception, Structure and Process. Hillsdale, New Jersey: Lawrence Erlbaum.

Sheridan, M.R. (1984). Planning and controlling simple movements. In M.M. Smythe and A.M. Wing (Eds.) The Psychology of Human Movement (pp. 47-82). London: Academic Press.

Shimojo, M., Sato, S., Seki, Y., and Takahashi, A. (1995). A system for simultaneous measuring grasping posture and pressure distribution. Proceedings of ICRA '95, May, Nagoya, Japan.

Sivak, B. & MacKenzie, C.L. (1990). Integration of visual information and motor output in reaching and grasping: The contributions of peripheral and central vision. Neuropsychologia, 28(10), 1095-1116.

Smets, G., Gaver,. W.W., Overbeeke, C.J., and Stappers, P.J. (1993). Designing in virtual reality: perception-action coupling and form semantics. Proceedings of the International Conference on Event Perception and Action, Vancouver, 11-12.

Smith, K.U., and Sussman, H. (1969). Cybernetic theory and analysis of motor learning and memory. In E. Bilodeau and I. Bilodeau (Eds.), Principles of Skill Acquisition (pp. 103-139). New York: Academic Press.

Smith, W.M., Smith, K.U., Stanley, R., and Harley, W. (1956). Analysis of performance in televised visual fields: Preliminary report. Perceptual and Motor Skills, 6, 195-198.

Soechting, J.F. (1984). Effect of target size on spatial and temporal characteristics of a pointing movement in man. Experimental Brain Research, 54, 212-132.

Soechting, J.F., and Flanders, M. (1989). Errors in pointing are due to approximations in sensorimotor transformations. Journal of Neurophysiology, 62(2), 595-608.

Soechting, J.F., and Flanders, M. (1989). Sensorimotor representations for pointing to targets in three-dimensional space. Journal of Neurophysiology, 62(2), 582-594.

Sollenberger, R.L., and Milgram, P. (1993). Effects of stereoscopic and rotational displays in a three-dimensional path-tracing task. Human Factors, 35(3), 483-499.

Spijkers, W.A.C., and Lochner, P. (1994). Partial visual feedback and spatial end-point accuracy of discrete aiming movements. Journal of Motor Behaviour, 26(3), 283-295.

Stammers, R.B., and Bird, J.M. (1980). Controller evaluation of a touch input air traffic data system: An indelicate experiment. Human Factors, 22(5), 581-589.

States, R.A. (1995). Comparison of three methods for locating joint centers during planar arm motion. Abstracts for Neuroscience, November, San Diego, 424.

Stelmach, G.E., Castiello, U., and Jeannerod, M. (1994). Orienting the finger opposition space during prehension movements. Journal of Motor Behaviour, 26(2), 178-186.

Strong, G. (1995). New directions. Interactions, January, 69-81.

Sturman, D.J., and Zeltzer, D. (1993). A design method for "whole-hand" human-computer interaction. ACM Transactions on Information Systems, 11(3), 219-238.

Sutherland, I.E. (1963). Sketchpad: A man-machine graphical communication system (Technical Report 296) M.I.T. Lincoln Laboratory, Lexington, MA.

Takeuchi, A., and Naito, T. (1995). Situated facial displays: Towards social interaction. Proceedings of CHI '95, Denver, Colorado, May, 450-455.

Taubes, G. (1994). Taking the data in hand - literally - with virtual reality. Science, 265, 884-886.

Teal,. S.L., and Rudnicky, A.I. (1992). A performance model of system delay and user strategy selection. Proceedings of CHI '92, 295-305.

Thimbleby, H. (1990). User Interface Design. New York: Addison-Wesley.

Thorisson, K.R. (1994). Simulated perceptual grouping: An application to human-computer interaction. Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society, Atlanta, Georgia, July.

Trankle, U. and Deutschmann, D. (1991). Factors influencing speed and precision of cursor positioning using a mouse. Ergonomics, 34(2), 161-174.

Trevarthen, C.B. (1968). Two mechanisms of vision in primates. Psychologische Forschung, 31, 299-337.

Turvey, M.T. (1977). Contrasting orientations to the theory of visual information processing. Psychological Review, 84(1), 67-88.

Ullman, S. (1980). Against direct perception. Behavioural and Brain Sciences, 3, 373-415.

Vanden Abeele, S., Delreux, V., Crommelinck, M., and Roucoux, A. (1993). Role of eye and hand initial position in the directional coding of reaching. Journal of Motor Behaviour, 25(4), 280-287.

Venolia, D. (1993). Facile 3D direct manipulation. Proceedings of INTERCHI '93, 31-36.

Verfaille, K. (1992). Variant points of view on viewpoint invariance. Canadian Journal of Psychology, 46(2), 215-235.

Walker, N. (1990). A comparison of selection times from walking and pull-down menus. Proceedings of CHI '90, 221-225.

Walker, N., Meyer, D.E., Smelcer, J.B. (1993). Spatial and temporal characteristics of rapid cursor-positioning movements with electromechanical mice in human-computer interaction. Human Factors, 35(3), 431-458.

Walker, N., Smelcer, J.B., and Neilsen, E. (1991). Optimizing speed and accuracy of menu selection: A comparison of walking and pull-down menus. International Journal of Man-Machine Studies, 35, 871-890.

Wallace, S.A., and Weeks, D.L. (1988). Temporal constraints in the control of prehensile movement. Journal of Motor Behaviour, 20(2), 81-105.

Ward, M., Azuma,. R., Bennett, R., Gottschalk, S., and Fuchs, H. (1992). A demonstrated optical tracker with scalable work area for head-mounted display systems. Proceedings of CHI '92, 43-52.

Ware, C., and Balakrishnan, R. (1994). Reaching for objects in VR displays: Lag and frame rate. ACM Transactions on Computer-Human Interaction, 1(4), 331-356.

Ware, C., Arthur, K., and Booth, K.S. (1992). Fishtank virtual reality. Proceedings of INTERCHI '93, 31-37.

Welford, A.T. (1968). Fundamentals of Skill. London: Methuen.

Wellner, P. (1993). Computer-augmented environments: Back to the real world. Communications of the ACM, 36(7), 23-27.

Westerink, J.H.D.M, and van der Reek, K. (1995). Pointing in object-oriented environments: appreciation versus performance. SIGCHI Bulletin, 27(1), 41-44.

Wexelblat, A. (1994). Natural gesture in virtual environments. Proceedings of ACM Conference on Virtual Reality Systems and Technology, Singapore, August.

Wickens, C. (1992). Engineering Psychology and Human Performance. Merrill.

Wing, A., Turton, A., and Fraser, C. (1986) Grasp size and accuracy of approach in reaching. Journal of Motor Behaviour, 18(3), 245-260.

Weiser, M. (1993). Some computer science issues in ubiquitous computing. Communications of the ACM, 36(7), 75-83.

Woodworth, R. (1899). The accuracy of voluntary movement, Psychological Review Monograph Supplement, 3.

Yuille, A.L. & Ullman, S. (1990). Computational theories of low-level vision. In D.N. Osherson, M.S. Kosslyn, & J.M. Hollerbach (Eds.) Visual cognition and action (pp. 5-40). Cambridge, Mass: MIT Press.

Zelaznik, H.N., Hawkins, B., and Kisselburgh, L. (1983). Rapid visual feedback processing in single-aiming movements. Journal of Motor Behaviour, 15(3), 217-236.

Zeltzer, D., and Johnson, M.B. (in press). Virtual actors and virtual environments: Defining, modeling and reasoning about motor skills. In L. MacDonald and J. Vince (Eds.), Interacting with Virtual Environments. New York: John Wiley & Sons.

Zeltzer, D., and Swetlana, G. (in press). Skillbuilder: designing motor skills for virtual actors. Proceedings of ACM SIGGRAPH '95.

Zhai, S., Milgram, P., and Drascic, D. (1992). An evaluation of four 6 degree-of-freedom input techniques. Proceedings of INTERCHI '93, 124-125.