HEWLETT PACKARD

Direct Manipulation and Other Lessons

David Frohlich
Personal Systems Laboratory
HP Laboratories Bristol
HPL-96-152
November, 1996

direct manipulation, interfaces, human computer interaction, mixed mode, graphical

Direct manipulation refers to an interface design philosophy which originated in the early 1980's and now dominates the creation of modern software packages. In this chapter I update this philosophy in the light of recent studies, theories and interface innovations in this area. The main lesson of these developments is that a manual mode of interaction does not always lead to direct or usable interfaces, and that conversational and mixed mode forms of interaction should be more widely considered as ways of extending the current HCI paradigm.

Internal Accession Date Only
Published as Chapter 8 in *Handbook of Human-Computer Interaction* by M. Helander, T. Landauer and P. Prabhu, North Holland, Elsevier Science Publishers B.V., 1997.

Chapter 8 in M. Helander, T. Landauer & P. Prabhu (Eds.) *Handbook of human-computer interaction*. North Holland: Elsevier Science Publishers B.V., 1997

Direct manipulation and other lessons

David M. Frohlich
Hewlett-Packard Laboratories, Bristol
August 1996

The understanding of direct manipulation interfaces is complex. There are a host of virtues and vices. Do not be dismayed by the complexity of the arguments. To our mind, direct manipulation interfaces provide exciting new perspectives on possible modes of interactions with computers...We believe that this direction of work has the potential to deliver important new conceptualizations and a new philosophy of interaction.

Hutchins, Hollan & Norman 1986, pp123-4.

1. INTRODUCTION

What comes to mind when you think of direct manipulation interfaces? Maybe an icon in the shape of a folder or filing cabinet, a pop down menu or two, sliders and gauges, mice and windows? In short, the modern graphical interface! This impression is confirmed by searching the human computer interaction literature for articles containing the term 'direct manipulation'. Matching articles often use the term very broadly to refer to all things graphical.

In this chapter I want to move beyond this simple conception of direct manipulation as a graphical interaction style, to the more complex notion of direct manipulation as a design philosophy - as indicated in the opening quote from Hutchins, Hollan & Norman. This means I will *not* be reviewing the best ways to design screens, icons, menus or even graphical interactions in general, as these topics are all covered separately and in depth in other chapters of this volume. Instead I will return to the original definition of direct manipulation introduced by Shneiderman (1982, 1983) and to the early debate it sparked about *why* certain manual/graphical forms of interaction are more attractive than command line interfaces. The rest of the chapter will try to update this debate in the light of recent critiques and findings on direct manipulation, and of other relevant developments in the HCI field. It concludes with a re-evaluation of the direct manipulation philosophy in the modern context, and an attempt to resolve aspects of the ongoing debate which is now fueled by concerns over the negative effects of graphical standards (Buxton 1993, Card 1995), and the rise of interface agents (Streeter 1995, Wooldridge and Jennings 1995).

Thus the main thrust of the chapter will be to guide the reader through the many twists and turns of the arguments and findings surrounding direct manipulation as a philosophy, in order to draw out some of the most valuable lessons they teach us about usability and design. The approach is similar to that recommended by Wixon & Good (1987) who argue for a deeper approach is similar to that recommended by Wixon & Good (1987) who argue for a deeper analysis of the dimensions of usability underlying proposed interface categories. Many of the resulting lessons *do* relate to the design of graphical interactions, but mainly at the level of deciding when to employ graphical rather than linguistic techniques at the interface. 'Other lessons' relate to the selection and use of conversational and mixed mode forms of interaction. In this sense the chapter can be seen as a compliment to others on graphical forms of interaction, and a particular companion to the chapter by Gould on The Design of Usable Systems (Chapter 35).

2. WHAT IS DIRECT MANIPULATION?

2.1 Original definitions and claims

Shneiderman (1982, 1983) first used the term direct manipulation to refer to an emerging class of highly usable and attractive systems of the day, including display editors, early desktop office systems, spreadsheets, CAD systems and video games. These systems had graphical interfaces which allowed them to be operated 'directly' using manual actions rather than typed instructions. Actions were mediated by special function keys, screen displayed menus and pointing devices such as mice, pens and joysticks. Such systems seemed to change the entire paradigm for human-computer interaction from *dialogue* to *manipulation* by utilizing what the programmers knew was a visual *language* in a way users believed to be an entire interactive *world.*

This illusion was the real 'magic' of the new approach which Shneiderman summarised in the following design principles:
(1) Continuous representation of the object of interest.
(2) Physical actions or labelled button presses instead of complex syntax.
(3) Rapid incremental reversible operations whose impact on the object of interest is immediately visible.

Thus display editors showed users a full page of text exactly as it would appear when printed out ('What You See Is What You Get'), instead of in abbreviated form punctuated with formatting commands (Principle 1). Modifications to this 'object of interest', *the page*, could be made directly in a number of ways, including typing new text at any selected point and watching the old text shuffle up to make room, highlighting and moving blocks of text with gestures and function buttons, and removing characters in situ at the press of a button (Principles 2 and 3). Office systems such as the Xerox Star (Canfield, Smith, Irby, Kimball,Verplank 1992) selected *documents* as their key 'objects of interest' and presented them to users as document icons on a virtual desktop instead of filenames in a list (Principle 1). Users could then open, print, file or delete documents by moving their icons around the desktop in relation to other familiar office objects such as folders, filing cabinets, wastebaskets, and printers (Principle 2). Such movements were continuously controlled by the user and had direct visible consequences on the documents involved (Principle 3). The story was similar for spreadsheets, CAD systems, video games, and early visual programming tools (e.g. Iseki & Shneiderman 1986). All presented some key set of objects in visual form which users could modify, extend, combine and otherwise manipulate.

Shneiderman claimed a number of usability benefits for these kinds of systems - the last of which has been added in a more recent publication (Shneiderman 1992). These include:
(1) Learnability
(2) Enhanced expert performance
(3) Memorability
(4) Fewer error messages
(5) Better feedback
(6) Reduced anxiety
(7) Increased control
However, from the earliest articles, Shneiderman has always acknowledged that the *quality* of the selected graphic representation that is critical to the direct manipulation effect. Simply finding a graphical alternative to a command-based interaction, in itself, will not suffice. The representation must be meaningful and accurate to users as determined by appropriate testing and evaluation. This list of benefits should therefore be seen as an idealised list representing the potential of the approach if executed well.

2.2 Modern direct manipulation interfaces

Thirteen years later Shneiderman is still a strong advocate of the direct manipulation approach. Speaking at a panel discussion at CHI '95 he pointed to modern developments in *information visualisation* as representing where direct manipulation has got to today and indicating where it is going in the future (Streeter 1995). Much progress has been made in extending the sheer amount of information that can be displayed to users by using 3 dimensions, stereo displays, multiple perspectives and layers, and novel data structures (e.g. Card, Robertson & Macinlay 1991, Harrison, Ishii, Vicente, & Buxton 1995, Strausfeld 1995). At the same time there has been an extension of input devices and techniques for manipulating visual information, with head, eye, hand and body movements, often in combination with novel sliders, menus or gestural languages (e.g. Ahlberg & Shneiderman 1994a, Foley 1987, Jacob 1991, Kurtenbach, Sellen & Buxton 1993, Venolia 1993, Zimmerman, Smith, Paradiso, Allport & Gershenfeld 1995). In effect, the *scope* for finding effective graphical representations of application 'objects' has never been greater, and there is much momentum behind the approach to do so (see Chapter 12 for a more extended discussion).

A particularly compelling example of a modern direct manipulation interface is the *Dynamic Homefinder* system (Ahlberg, Williamson & Shneiderman 1992, Shneiderman 1994, Williamson & Shneiderman 1992). This presents an interface to a database of houses for sale by displaying a map peppered with dots corresponding to individual homes (see Figure 1). Users constrain the range of dots that appear on the map by moving sliders for properties such as cost, number of bedrooms and distance to school or work. This system embodies all three of Shneiderman's defining principles for direct manipulation, but implementation of the rapid/incremental/reversible principle (3) is particularly impressive. The distribution of dots on the map changes continuously as each slider is moved left or right, allowing the user to explore the effects of an entire range of property settings in a sweep of the hand. A similar technique is used in the interface to a database of films (Ahlberg & Shneiderman 1994b). It is difficult to see how such 'dynamic queries' could be done at all, let alone better, in a standard query language.

FIGURE 1 ABOUT HERE

3. WHAT MAKES MANIPULATION DIRECT?

The first extensive critique of direct manipulation was written by Hutchins, Hollan & Norman (1986). The authors agreed with Shneiderman that certain graphical interfaces appear more attractive and usable than their command-based counterparts, but sought to explain when and why. In effect they were the first to ask "What makes manipulation direct?".

Their answer is rather complicated and abstract, and has to do with how the user thinks about the role of the computer in the interaction, and with the effort required to get it to do what you want. These two dimensions of *engagement* and *distance* are said to work together to conjure up the impression of directness at the interface. In short, Hutchins et al argue that:

DIRECTNESS = ENGAGEMENT + DISTANCE

To expand, engagement refers to the perceived locus of control of action within the system (after Laurel 1986). A critical distinction is whether users feel themselves to be the principle actors within system or not (see Figure 2). In systems based on a conversational style of interaction the locus of control appears to reside with a 'hidden intermediary' who executes linguistic commands such as <delete filename> on the users behalf. This is like shouting through to someone in another room who then performs the requested action and shouts back a response (Figure 2a). This interaction is indirect in the sense that the user is not directly engaged with the actual objects of interest such as documents. In systems based on a graphical style of interaction the locus of control appears to reside with users themselves who manipulate the objects of interest within a model world. This is like reaching into the world yourself to carry out the action (Figure 2b). Users act on the world in 'first-person' and are therefore said to be directly engaged with the system.

FIGURE 2 ABOUT HERE

Distance refers to the mental effort required to translate goals into actions at the interface and then evaluate their effects (after Norman 1986). Norman explains this with reference to the 'gulf' between goals and actions shown in Figure 3. This gulf is spanned through cycles of interaction in which the user does something and the system responds. Certain sorts of mental calculations have to be made by the user at the beginning of each cycle to work out what to do and bridge the 'Gulf of Execution'. Other sorts of calculations have to be made at the end of each cycle to work out what has happened and bridge the 'Gulf of Evaluation'. Systems and interfaces which make it easier for users to make these calculations are said to be more direct to operate. Use of the desktop metaphor used in the Xerox STAR system is a good example of this. By representing objects of interest as documents within a model world the designers ensured more than direct engagement with the system. They also invoked a set of ready-made expectations in users about what kinds of things could be done with documents (e.g. create, move, copy, open, delete) and what this might look like if they were. This might mean that users don't have to 'think so hard', or at least learn so much, compared to learning and using a command language for the same operations.

FIGURE 3 ABOUT HERE

Figure 1

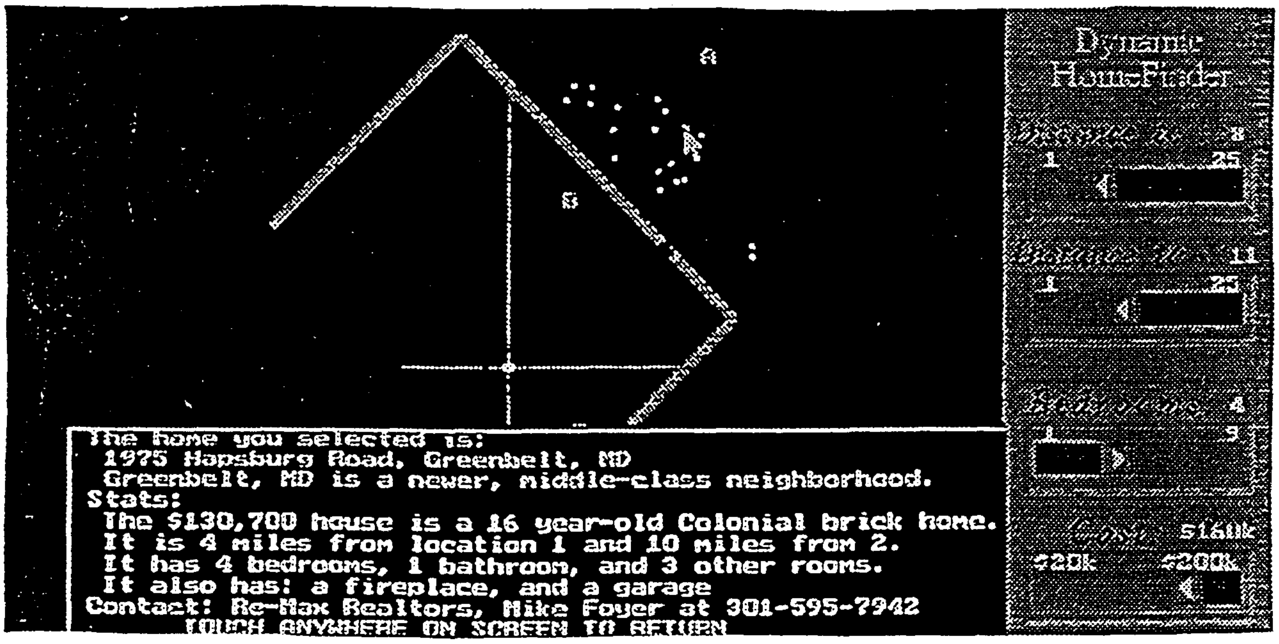
Dynamic HomeFinder
$20k
$200k
The home you selected is:
1975 Hapsburg Road, Greenbelt, MD
Greenbelt, MD is a newer, middle-class neighborhood.
Stats:
The $130,700 house is a 16 year-old Colonial brick home.
It is 4 miles from location 1 and 10 miles from 2.
It has 4 bedrooms, 1 bathroom, and 3 other rooms.
It also has: a fireplace, and a garage
Contact: Re-Max Realtors, Mike Foyer at 301-595-7942
TOUCH ANYWHERE ON SCREEN TO RETURN

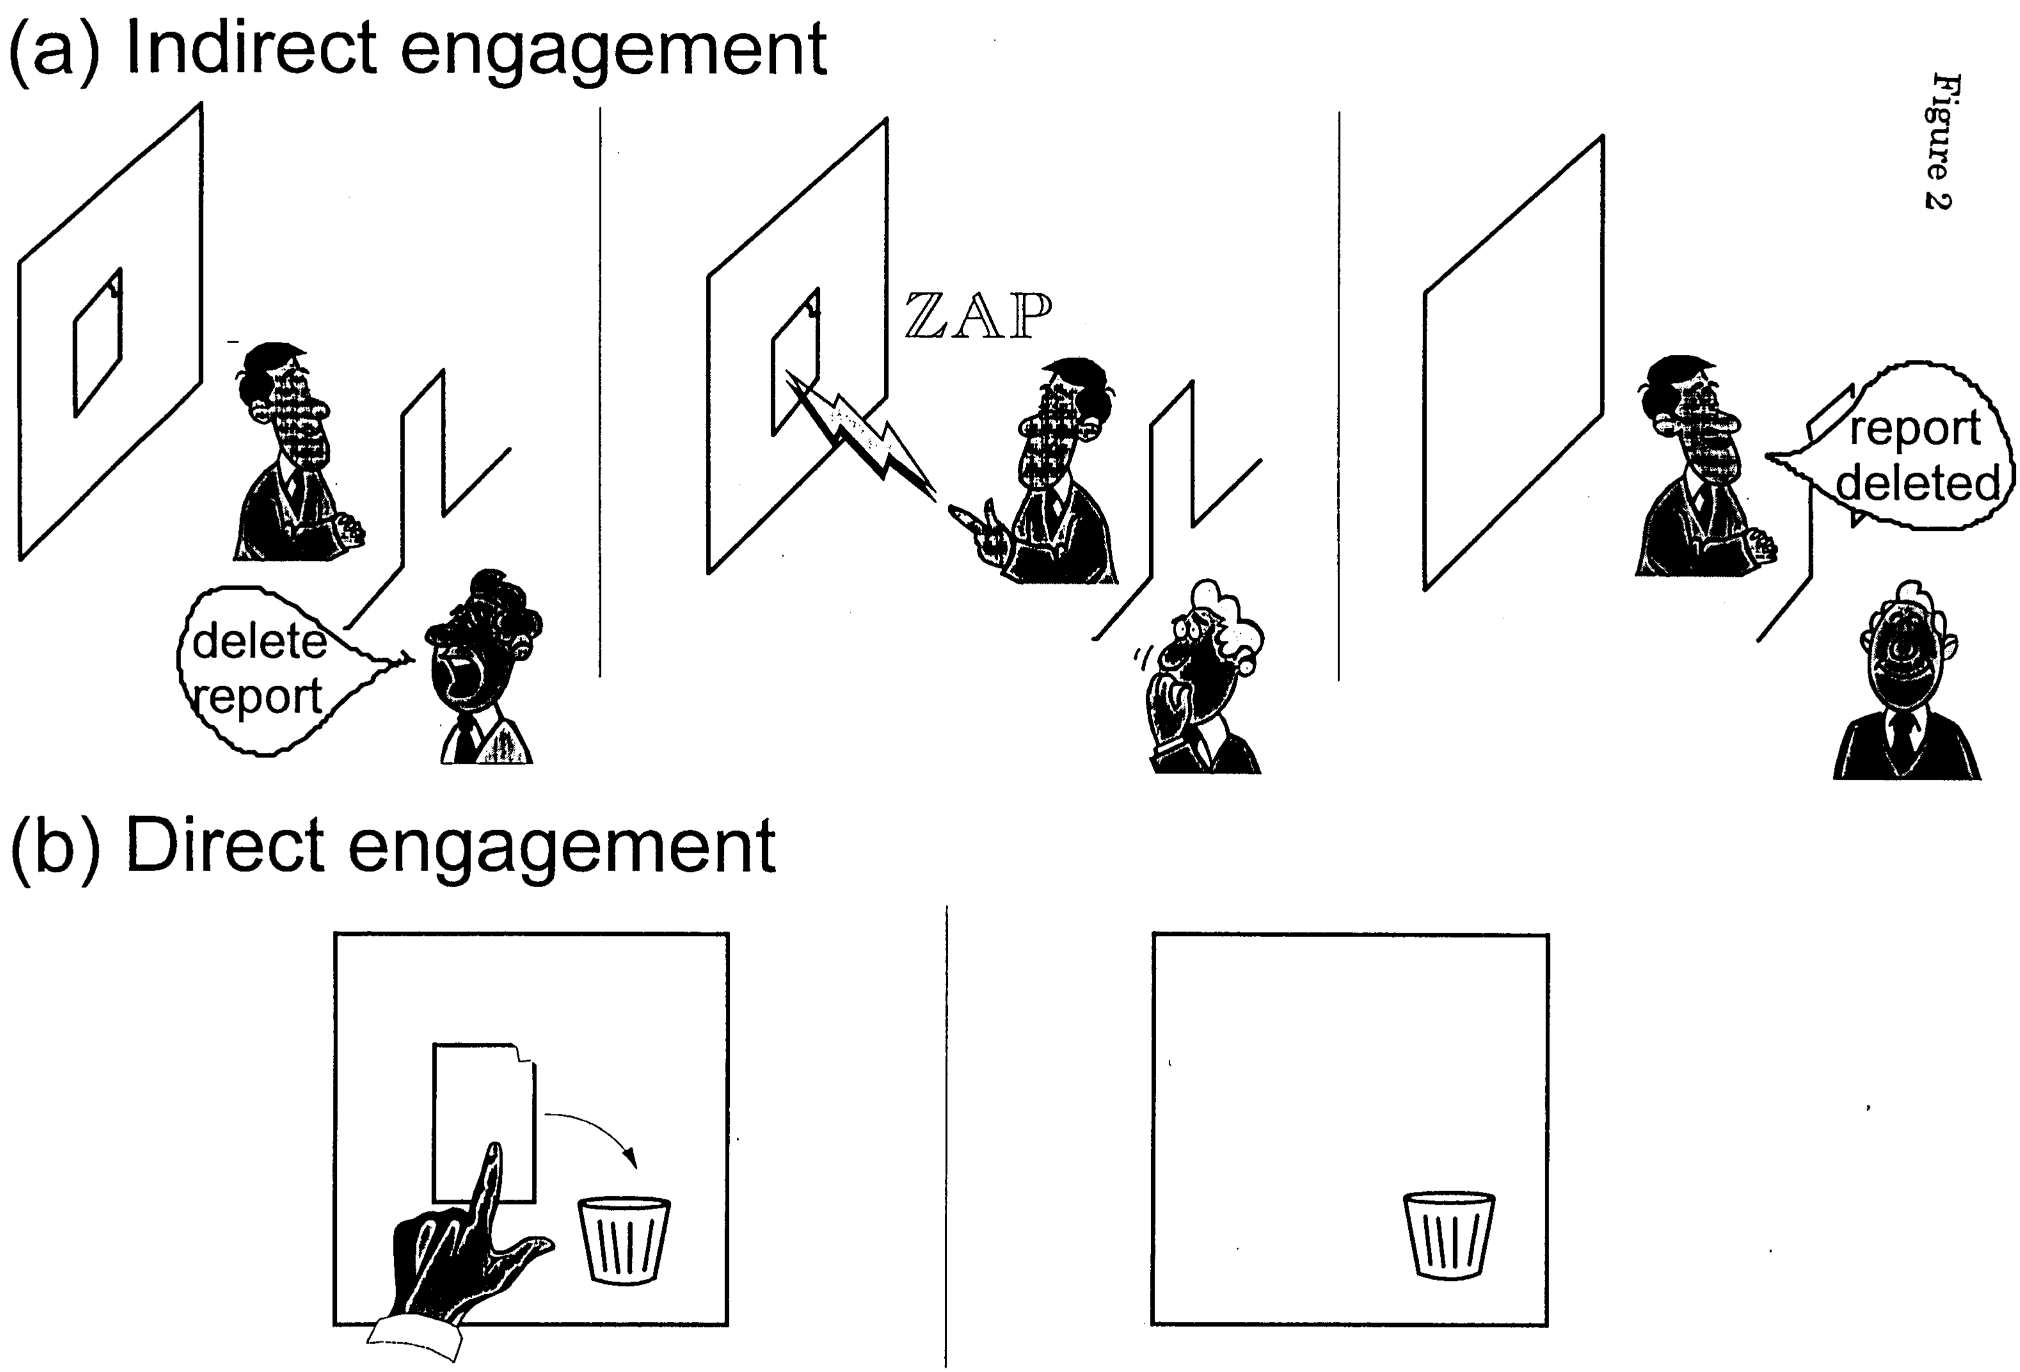
(a) Indirect engagement
delete report
ZAP
report deleted
(b) Direct engagement

Figure 2

Figure 3

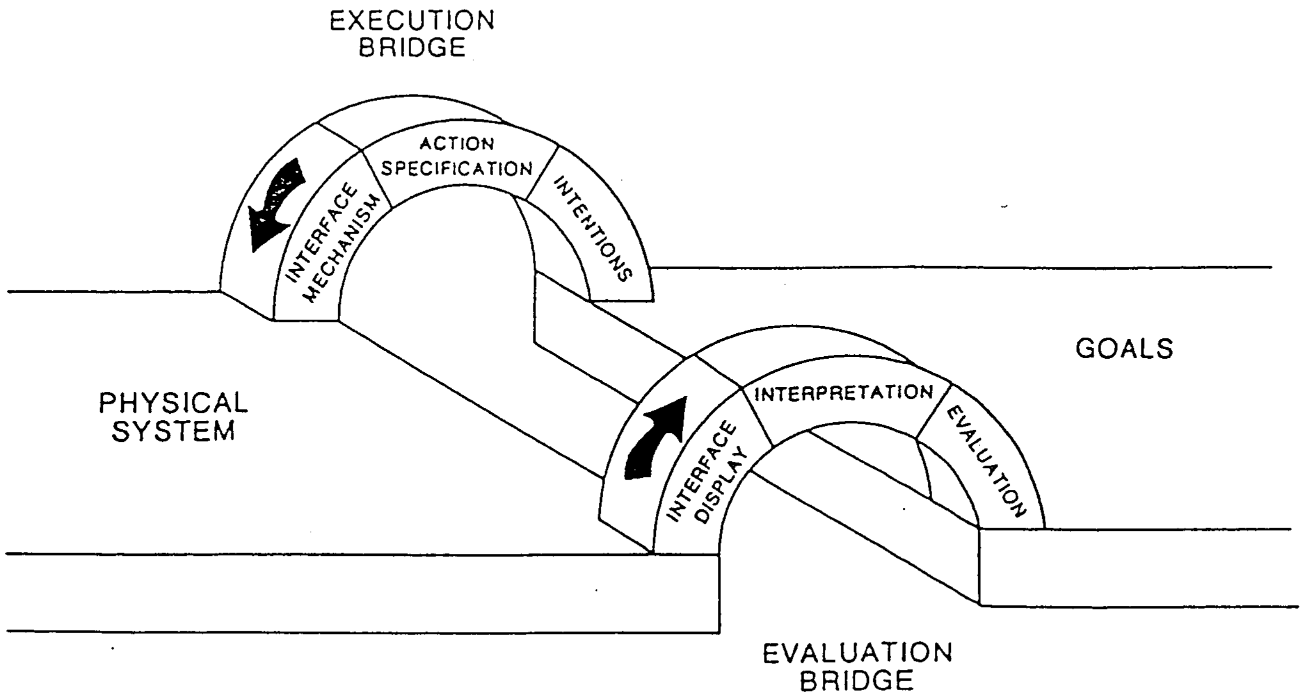

EXECUTION
BRIDGE
ACTION
SPECIFICATION
INTERFACE
MECHANISM
INTENTIONS
GOALS
PHYSICAL
SYSTEM
INTERPRETATION
INTERFACE
DISPLAY
EVALUATION
EVALUATION
BRIDGE

This analysis leads Hutchins et al to propose a space of interfaces varying in both engagement and distance as shown in Figure 4. **Direct manipulation** turns out to be only one of four possible types in which users treat the computer as a model world *and* find it easy to translate their goals into actions. By implication, they are the best of all interfaces. The possibility of poorly designed graphical systems is acknowledged in **low-level world** interfaces which are difficult or tedious to operate despite engaging users directly in a model world. Furthermore, the lack of a model world representation is not always shown to result in poorer interaction as evidenced in **high-level language** interfaces. These may support easy and effective interactions through the type of indirect engagement illustrated in Figure 2a by using commands which correspond closely to what users want to do. However, poor implementations of these systems result in **low-level languages**. These are the worst of all interfaces since they require users to work hard to achieve their goals through indirect engagement.

FIGURE 4 ABOUT HERE

Returning to direct manipulation interfaces, Hutchins et al go on to question some of Shneiderman's claims for them. Their *learnability* and *memorability* is said to depend upon the users prior knowledge and its semantic mapping to the system. The *need for error messages* is said to remain for task related problems, despite decreasing overall because of better feedback for the evaluation of actions. And *expert performance* is suspected to slow down rather than speed up through direct manipulation, for at least two reasons. First the loss of language involves the loss of descriptive expressions which allow the user to refer to classes of objects or actions in a single input. Second, the use of familiar real-world metaphors may lock the user into existing ways of doing things which are easy to learn and remember for the novice, but less effective than they could be for an expert in the new medium. They conclude with some force therefore that "much empirical work remains to be done to substantiate these claims" (Hutchins et al 1986, p123).

In the next section we turn to this work, and to a related body of research on mixed mode interfaces combining elements of language and action in the same system.

4. DATA AND DEVELOPMENTS

4.1 Tests

A large number of studies have tested the above claims in various ways. These tests break down into three categories:

(1) Uncritical comparative evaluations - measuring the usability of a direct manipulation interface against one or more alternative interfaces

(2) Critical comparative evaluations - measuring the usability of several different implementations of a direct manipulation interface, sometimes but not always against alternative interfaces, so as to separate different elements of the effect.

(3) Naturalistic choice studies - measuring interactional preferences for direct manipulation and alternative methods in a mixed mode interface

Figure 4

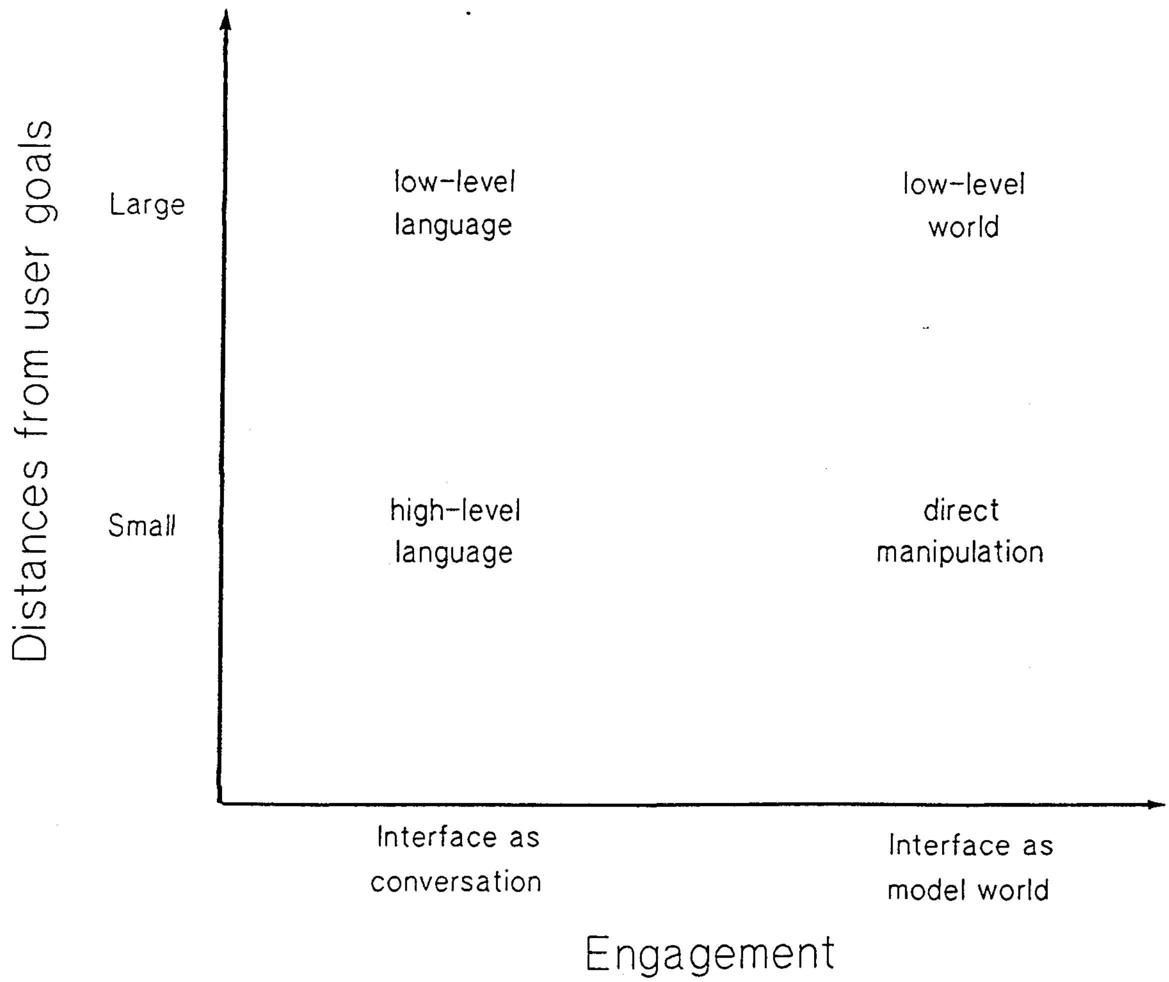
Distances from user goals
Large
Small
low-level
language
low-level
world
high-level
language
direct
manipulation
Interface as
conversation
Interface as
model world
Engagement

The first two types of tests employ a standard experimental paradigm in which independent groups of subjects perform the same task using one of several different interfaces. Dependent measures of task performance are then used to make group comparisons. The last type of test is based on a 'Combined' condition in which subjects are faced with a mixed mode interface. A more qualitative analysis of the choice of method used by subjects throughout the session is conducted to reveal the contexts in which direct manipulation is and is not favoured. The key insights from each of these tests are as follows.

4.1.1 Uncritical comparative evaluations

Most tests of the benefits of direct manipulation fall into this category. However their key insights are difficult to summarise because of differences in the way they implement direct manipulation and contrasting interfaces, in the tasks given to subjects, and in the way they measure usability. Consequently their findings are mixed; with some studies showing clear positive advantages for direct manipulation, others showing no differences, and yet others showing clear disadvantages.

Advantages for direct manipulation have been demonstrated on word processing tasks (Card, Moran & Newell 1983, Frese, Schulte-Gocking & Altmann 1987, Jordan 1992), file manipulation tasks (Karat 1987, Morgan, Morris & Gibbs 1991, Margono & Shneiderman 1987), database retrieval tasks (Egan, Remde, Landauer, Lochbaum & Gomez, 1989, Eberts & Bittianda 1993, Te'eni 1990, Ulich, Rauterberg, Mott, Greutmann & Strohmo 1991). The direct manipulation interfaces here were display editors, Macintosh-style operating systems, desktop filing systems, and an interactive graphics package. The alternative interfaces were all command line interfaces, except in Ulich et al (1991) where a menu interface was used. The studies measured different combinations of factors represented in Shneiderman's list of benefits and often found superiority of direct manipulation on only a subset of the measures used. For example, Eberts & Bittianda found benefits in speed and understanding but not accuracy; Morgan et al found benefits in accuracy and satisfaction but not speed; Karat found benefits in speed for all but one task; while Frese et al found performance benefits late but not early in practice.

No differences between direct manipulation and other interfaces have been found in studies of file managment (Whiteside, Jones, Levy & Wixon 1985), drawing (Eberleh, Korfmacher & Streitz 1992) and matching concepts and labels (Kacmar 1989). 'Iconic' interfaces were contrasted with command and/or menu interfaces by Whiteside et al and Kacmar, while Eberleh et al examined the use of commands and actions on a graphics task, with and without background distractions. There were no overall performance benefits in the use of icons and no savings in 'workload' experienced by subjects using the direct manipulation graphics interface. Furthermore, Whiteside et al and Eberleh et al both found that subjective preferences varied across individuals and could change with the task conditions.

Disadvantages for direct manipulation have been shown in studies of table manipulation (Tullis & Kodimer 1992), filing and retrieval (Jones & Dumais 1986), and browsing (Egido & Patterson 1988). Thus Tullis & Kodimer found two drag and drop methods to be slower than the use of radio buttons or a single data entry box for changing the order of fields in a table. In the other two experiments spatially arranged icons were contrasted with text labels and found to be less accurate for filing and retrieval (Jones & Dumais 1986) and slower for navigating through a catalogue (Egido & Patterson 1988).

The fact that the benefits of direct manipulation depend on implementation, task and measurement factors is perhaps the greatest insight provided by these tests. Taken together the tests suggest that direct manipulation interfaces do not provide a wide variety of usability benefits on any task. Rather, these interfaces seem to improve selected aspects of usability on a restricted set of tasks.

In general, the tests support Hutchins et al's view of a spectrum of manipulation-based and language-based interfaces varying in effectiveness (see again Figure 4). Although each study claims to have employed a direct manipulation interface, many may actually have used a low level world interface which doesn't map well to the goals users are trying to achieve with it. Depending on whether the alternative conversational interfaces are high or low level languages according to Figure 4, users may perform as well or even better with those. Hence the mixed results.

One problem here is that a surface description of a 'direct manipulation' interface doesn't ensure small 'distance' characteristics, and hence usability, because those features are dependent on the *task* the interface will be used for. This also means that the same manual interface might be classified as direct when used for one task and indirect (i.e. 'low level') when used for another. This is a problem for the direct manipulation philosophy since most interfaces support a variety of tasks or goals, some of which may be performed better through conversational rather than manual techniques[1]. A related problem is that even when manual techniques *are* appropriate to the current task, there is insufficient guidance in the original literature for designers to choose between competing implementations so as to achieve directness of manipulation in the resulting interface (Kunkel, Bannert & Fach 1995). It is to this issue that the next set of studies are addressed.

4.1.2 Critical comparative evaluations

These tests set out to examine the effectiveness of different implementations of direct manipulation interfaces against each other and against alternative interfaces. In doing so they help to tease apart the relative importance of different elements of direct manipulation interfaces and therefore to inform design choices within and between them.

The need for such tests is underlined by Kunkel et al (1995) who list 10 elements associated with direct manipulation interfaces in the literature. These should be recognisable from the descriptions in Sections 2 and 3:

1. object-orientation
2. permanent visualization
3. extensive use of metaphors
4. action-techniques using pointing devices
5. window-techniques
6. function-objects
7. integration of different input techniques

---

[1] A version of this multiple-task problem applies to any design philosophy which promotes one style of interaction over another. However in the case of direct manipulation it is compounded by the fact that the *definition* as well as the *success* of the promoted interaction method is relative to the tasks being performed.

8. use of icons
9. function activation by menu selection or dragging
10. user's adjustment of design features

Tests exist on the role of object-orientation (Kunkel et al 1995), permanent visualisation (Ahlberg, Williamson & Schneiderman 1992, Ballas, Constance, Heitmeyer & Perez 1992, Tabachneck & Simon 1994), metaphors (Anderson, Smyth, Knott, Bergan, Bergan & Alty 1994, Ankrah, Frohlich & Gilbert 1990), action-techniques (Ahlberg et al 1992, Ballas et al 1992) icons (Benbasat & Todd 1993, MacGregor 1992) and function activation (Benbasat & Todd 1993, Kunkel et al 1995).

Thus Kunkel et al (1995) themselves examined the relative benefits of interfaces varying in object orientation and function activation. Using the experimental design shown in Figure 5 they tested 4 independent groups of 16 subjects on one of four direct manipulation interfaces. Object orientation was varied by contrasting an object-function command order with a function-object command order. Function activation was varied by contrasting clicking with dragging selected objects. The actual system used was a graphics editor displaying triangle, circle and square objects together with icons for colour, pattern and line thickness functions. Users had to reproduce 10 abstract pictures on one of the interfaces. Command order made no significant difference to subjects performance on the task, whereas the mode of function activation did. Explicit function activation by clicking was found to be more accurate and efficient than implicit activation by dragging. There were no differences in subjects attitudes to the interfaces.

FIGURE 5 ABOUT HERE

In a related experiment, Benbasat and Todd (1993) varied the method of function activation in an office system interface by supporting either clicking or dragging operations. They refer to the clicking interface as 'Menu' based and the dragging interface as 'Direct manipulation' based. They also vary the iconic nature of each interface, by presenting actions and objects as icons or as textual labels. This leads to the four possible interfaces shown in Figure 6. Four independent groups of 12 subjects performed a meeting arrangement task on one of the interfaces. There were no significant differences between the Icon and Text conditions in speed or accuracy of performance, although dragging ('Direct manipulation') was significantly faster than clicking ('Menu'), at least on early trials.

FIGURE 6 ABOUT HERE

In a direct study of the effect of icons, MacGregor (1992) tested the relative usability of the three interfaces to a Videotex information system shown in Figure 7. The 'Label' interface simply presented a menu of textual items refering to different sections of information in the system. The 'Descriptor' interface presented the same labels with textual examples next to each item, while the 'Icon' interface presented the labels with examples shown as icons. Three independent groups of 10 subjects answered 36 questions on one of the interfaces. Use of the Descriptor and Icon interfaces resulted in significantly more accurate performance than that on the Label interface, but were themselves equivalent. This implies that there is no general advantage to using icons rather than textual menus, but that in particular cases icons may be better because of the additional information they carry. However, other studies have shown that poorly designed icons can actually be worse than labels because they carry less information (e.g. Franzke 1995).

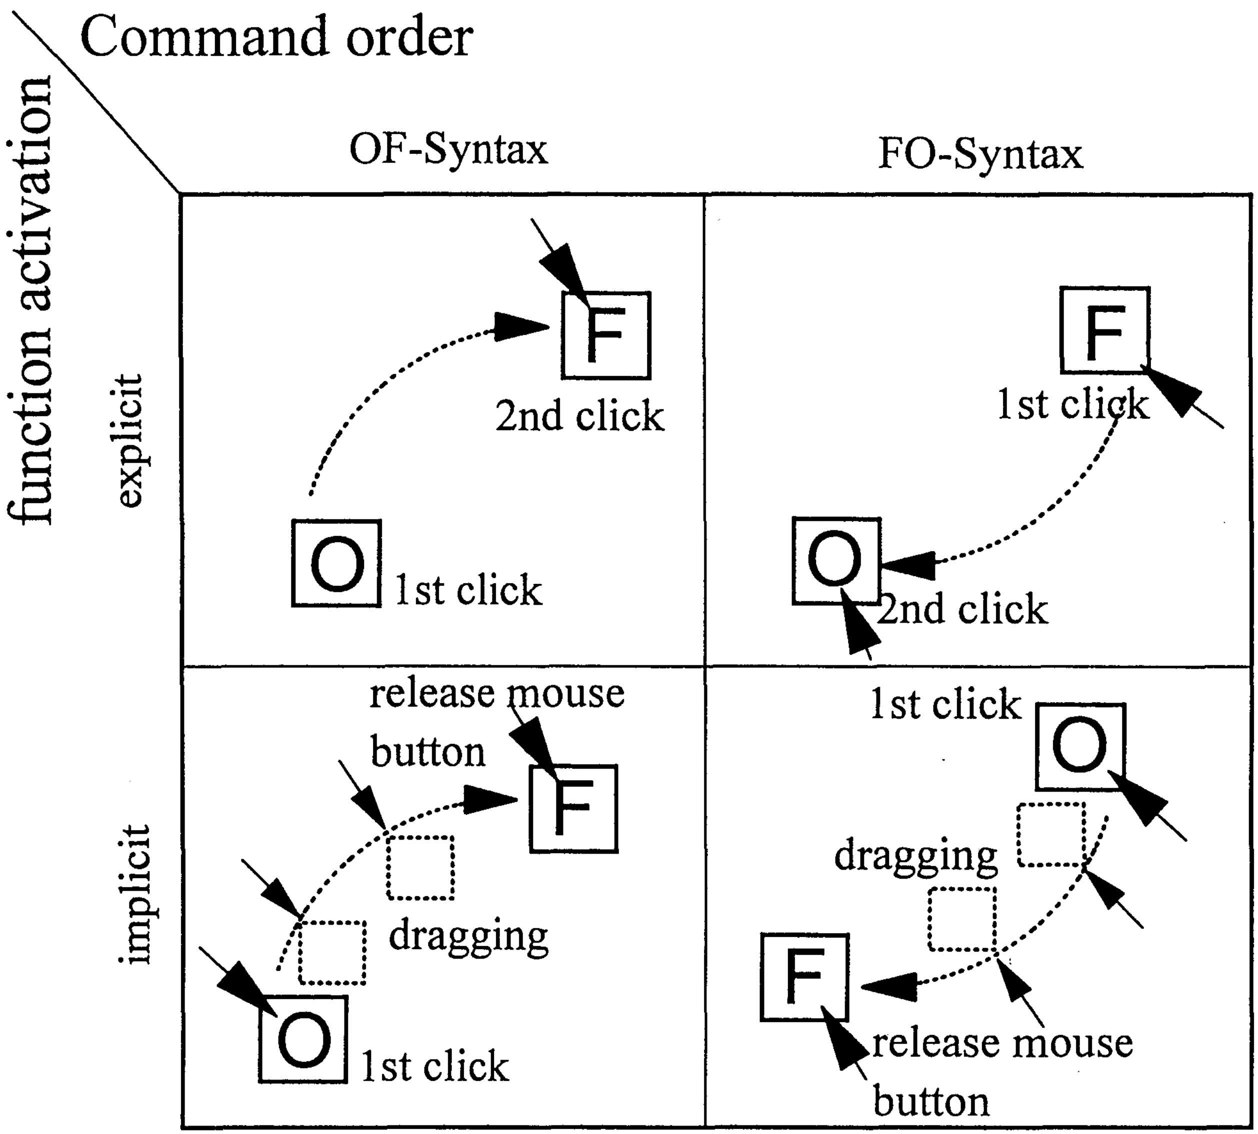

Command order
function activation
OF-Syntax
FO-Syntax
explicit
implicit
F
2nd click
O
1st click
F
1st click
O
2nd click
release mouse
button
F
dragging
O
1st click
1st click
O
dragging
F
release mouse
button

Figure 5

Figure 6

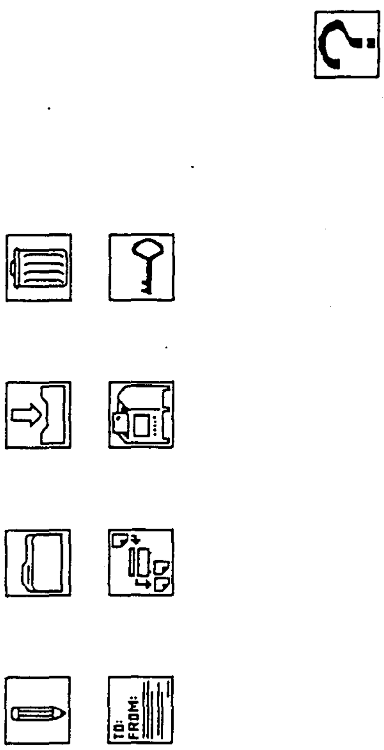

Direct manipulation—icon-based interface.

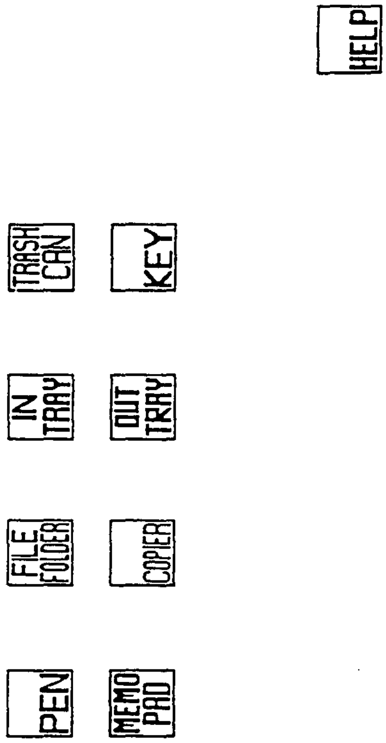

Direct manipulation—text-based interface.

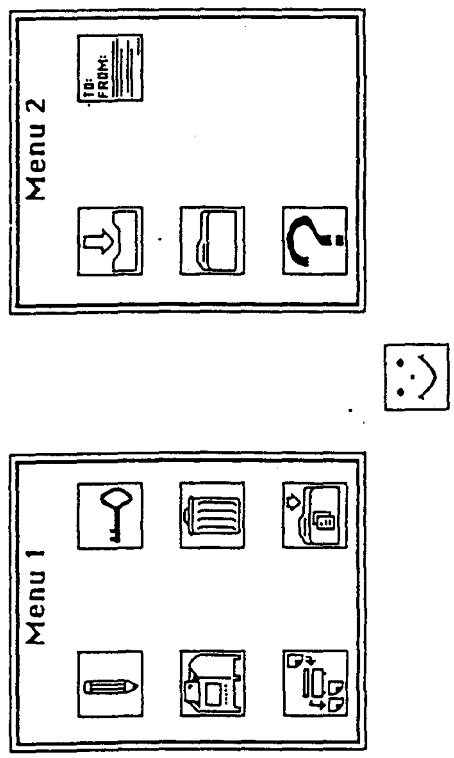

Menu—Icon-based interface.

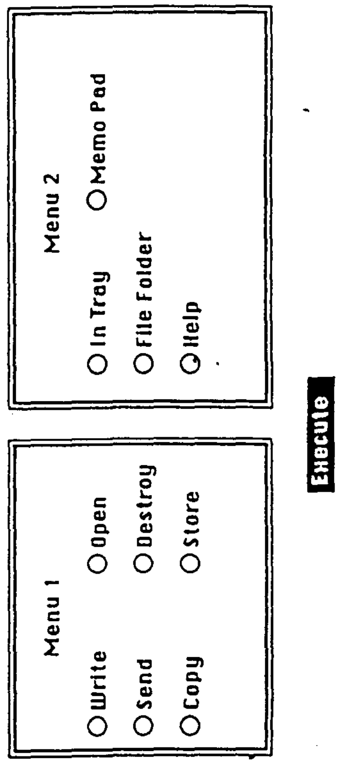

Menu—text-based interface.

FIGURE 7 ABOUT HERE

A further aspect of iconic representation is the fact that icons can convey real-world metaphors to the user. In an attempt to separate the effects of metaphor from 'directness' at the interface, Ankrah et al (1990) chose to implement a version of Norman's (1986) example task of filling a bath to a required level and temperature (see Figure 8). In a 'With metaphor' condition users could see an iconic bathtub fill with water and watch the temperature rise in an iconic thermometer. In a 'Without metaphor' condition subjects saw only the movement of lines within two similarly shaped boxes. Within each of these interfaces directness was varied by modifying the type of tap control system embodied in two screen-displayed sliders. Following Norman's example, 'Direct' controls were supported by the use of sliders corresponding to rate of flow and temperature, while 'Indirect' controls were supported by the use of sliders for hot and cold water. Four independent groups of 10 subjects performed the process control task in one of the four resulting interfaces. Subjects performed better overall on the Direct interfaces but not on the With Metaphor interfaces. In fact the most usable interfaces were the Direct-Without Metaphor and Indirect-With Metaphor ones as revealed in a significant interaction between the factors. This effect appears to be related to the fact that the bath metaphor conveys a message not only about the task but also about the control system involved. British subjects can be expected to assume hot and cold tap controls in both the Indirect *and* Direct With Metaphor conditions, which would favour performance on the Indirect With Metaphor system only. The absence of this assumption also favoured performance on the Direct Without Metaphor interface.

FIGURE 8 ABOUT HERE

In a further study of the effect of metaphor on human computer interaction Anderson et al (1984) varied instructions in the use of a desktop videoconferencing system. Three sets of instructions were use to communicate Office Door, Dogs or Traffic Lights metaphors for system operation. Three independent groups of 6 subjects were then asked to make and receive the same connections in a communication task. A post-session questionnaire revealed that subjects made different types of errors in their reasoning about the system depending on the metaphor to which they had been exposed. Thus the Doors metaphor led subjects to believe that the system could do more than it could more often than did the Traffic Lights metaphor. Furthermore within each metaphor condition there were significant differences in the classes of reasoning mistakes made. The findings were generally consistent with a proposed model of metaphor, reproduced in Figure 9, which distinguishes between various overlaps in the supported (+) or unsupported (-) functionality of any real-world 'vehicle' (V) and computer system domain (S).

FIGURE 9 ABOUT HERE

Three studies examine the utility of data visualisation at the interface. In the first study Ballas et al (1992) attempted to vary distance and engagement in a naval command and control interface. Distance approximated to an output visualisation factor which was varied by displaying radar contact information either as a map of the spanned area (Low Distance) or as a table of contacts and properties (High Distance). Engagement approximated to an input visualisation factor which was varied by use of touchscreen (Direct Engagement) or keypad (Indirect Engagement) input devices. This resulted in the four interfaces shown in Figure 10.

Figure 7

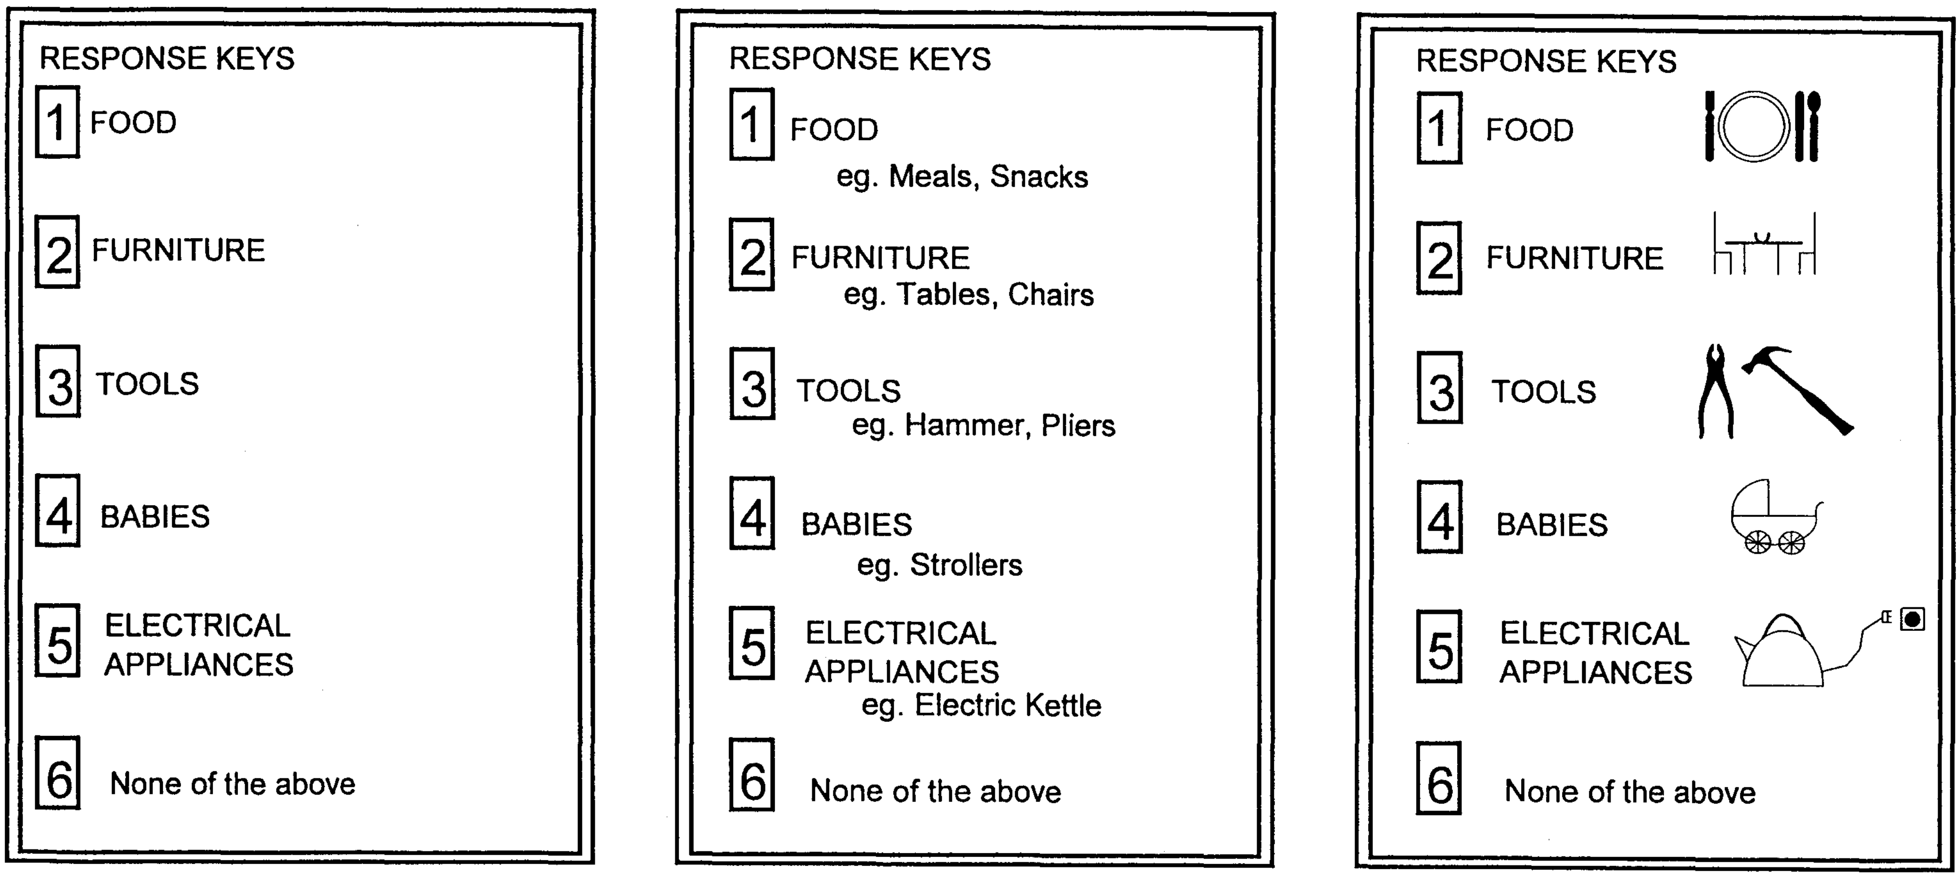

(a)Label (b)Description (c)Icon

Figure 8

(a) With Metaphor

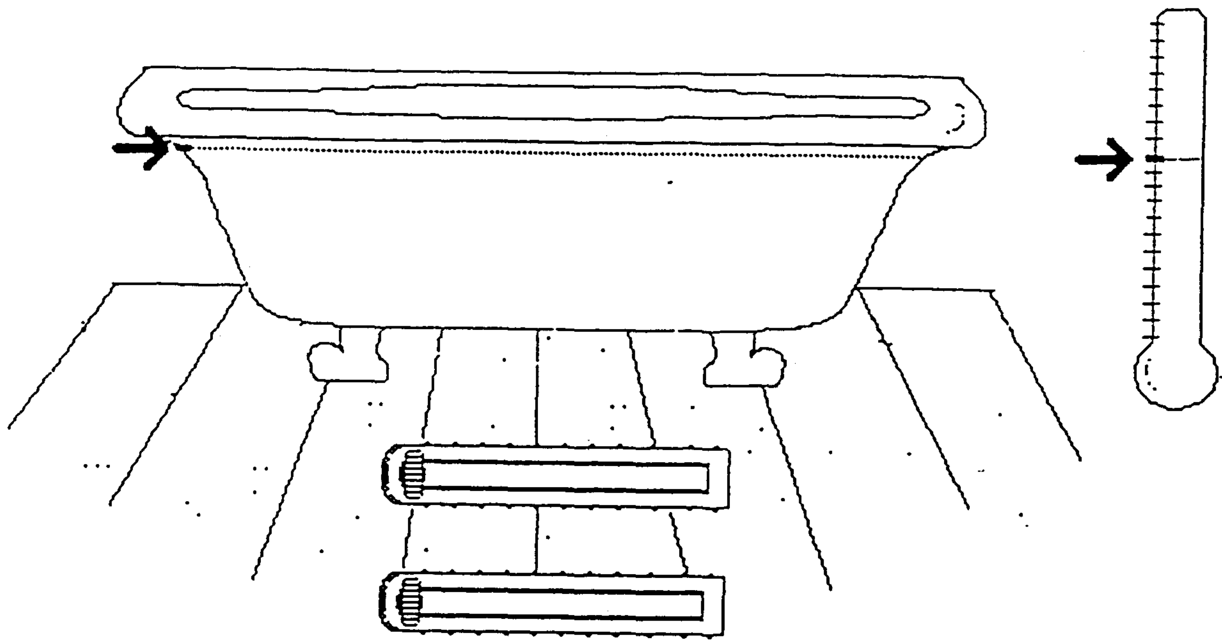

(b) Without Metaphor

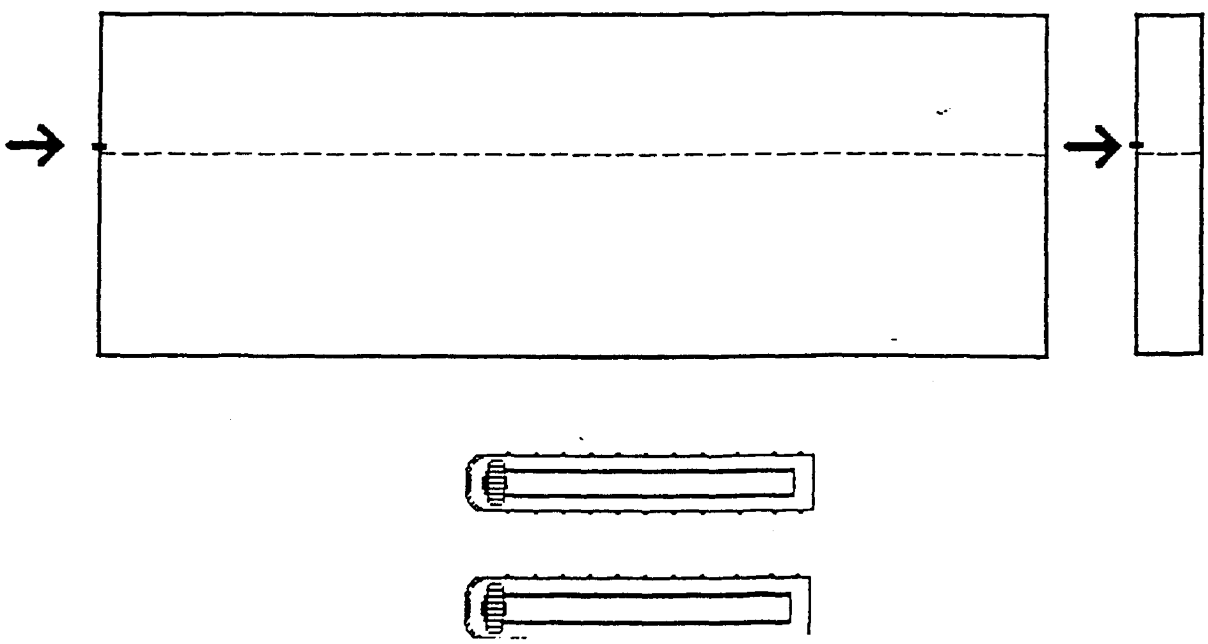

Figure 9

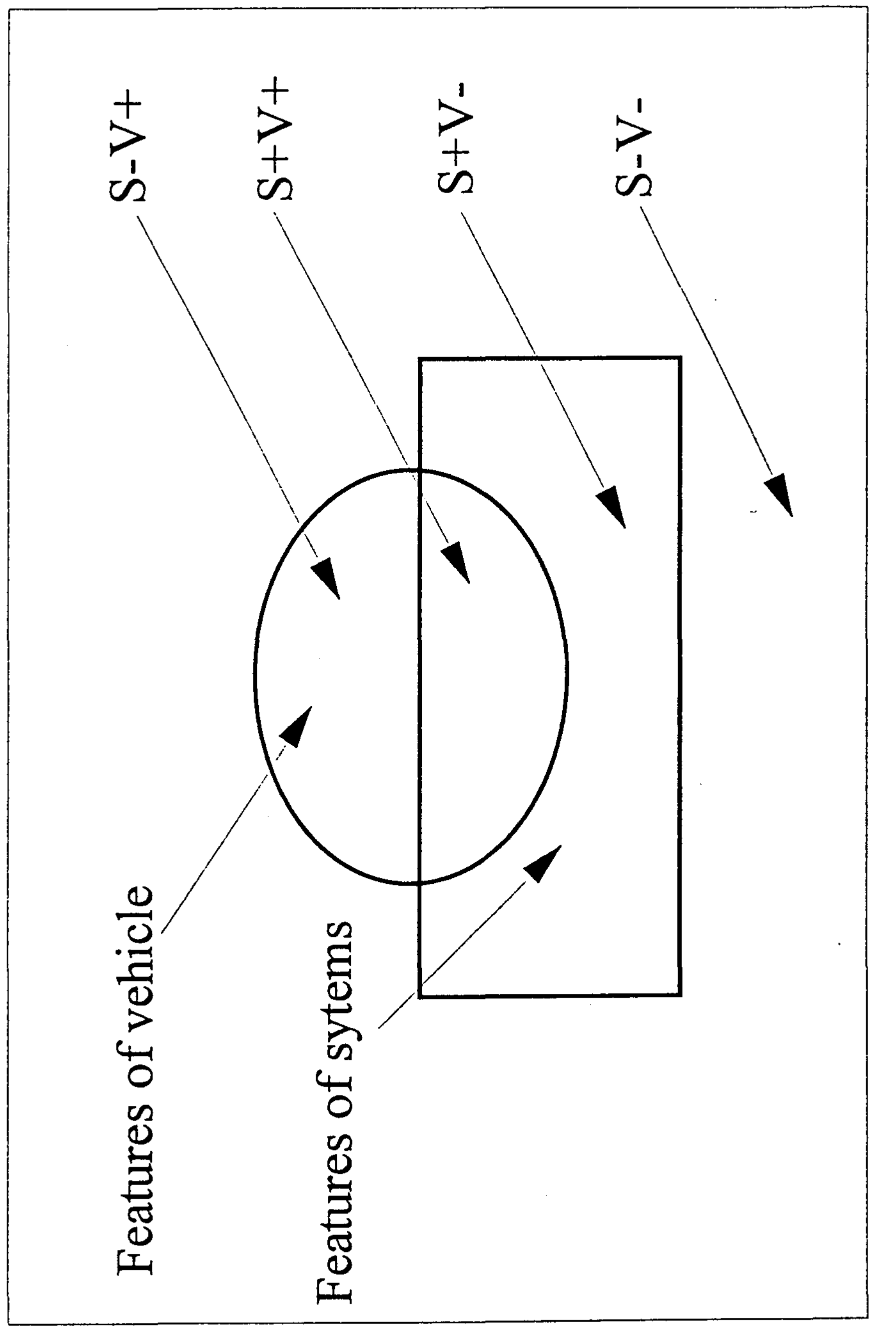

Four independent groups of 5 subjects performed tracking and classification tasks on the interfaces. Although there were no significant differences in the speed and accuracy of task performance across the interfaces, subjects using the graphical map displays felt themselves better able to monitor and predict changes in the tactical situation. This was confirmed by the greater ability of subjects on the map display to resume control of a previously automated activity (at least for confirmation decisions).

FIGURE 10 ABOUT HERE

A similar study was carried out by Ahlberg et al (1992) in a different domain. They varied input and output visualisation in an interface to the Periodic Table. Output varied between a tabular arrangement of highlighted element symbols (G=Graphical output) and a list of selected symbols (T=Textual output). Input varied between use of screen-displayed sliders for atomic properties (S=Slider input) and keyboard input of property values (F=Form fill-in). Although this would theoretically generate four interface designs, only the three shown in Figure 11 were evaluated in the study, with the SG combination refered to as the 'Dynamic Query' interface. An ST interface combining sliders with textual output is missing from the study. Three independent groups of 6 subjects answered questions about the Periodic Table each using one of the interfaces shown. Subjects using the Dynamic Query interface answered questions faster than those using the other two interfaces, and more accurately than subjects using the FT interface (with Textual output). This suggests that removal of *output* visualisation may have a more detrimental effect on (the accuracy of) performance than removal of *input* visualisation.

FIGURE 11 ABOUT HERE

Finally, Tabachneck & Simon (1994) found some detrimental effects of output visualisation in a computer assisted learning task. Three independent groups of 4 economics students were given questions to answer from data displayed either in graphs, tables or algebraic expressions. Subjects using graphs were better at answering quantitative questions than subjects using equations or tables, but not at answering explanatory questions requiring underlying reasons. Two subsequent experiments appear to confirm a hypothesis that visualisation appears to induce superficial visual reasoning from the display which may compromise a deeper understanding of actual causal factors.

Taken together, these studies begin to break down the rather nebulous notion of direct manipulation into more concrete pieces. They also challange some of the unwritten assumptions about what features of graphical interfaces work best. I have described them in some detail here because they show that there are many dimensions to the concept which vary in importance and interact with each other in complex ways. From the available data we can see the following themes emerge:

- Order of command activation may not matter
- Clicking may be slower but more accurate than dragging
- Icons may be equivalent to menus with extra information
- Metaphors should be selected with a view to the device as well as the task characteristics they convey
- Choice between alternative metaphors may be informed by mapping out their applicability and inapplicability to the target computer domain
- Visualisation of output data seems critical to the benefits of direct manipulation

Figure 10

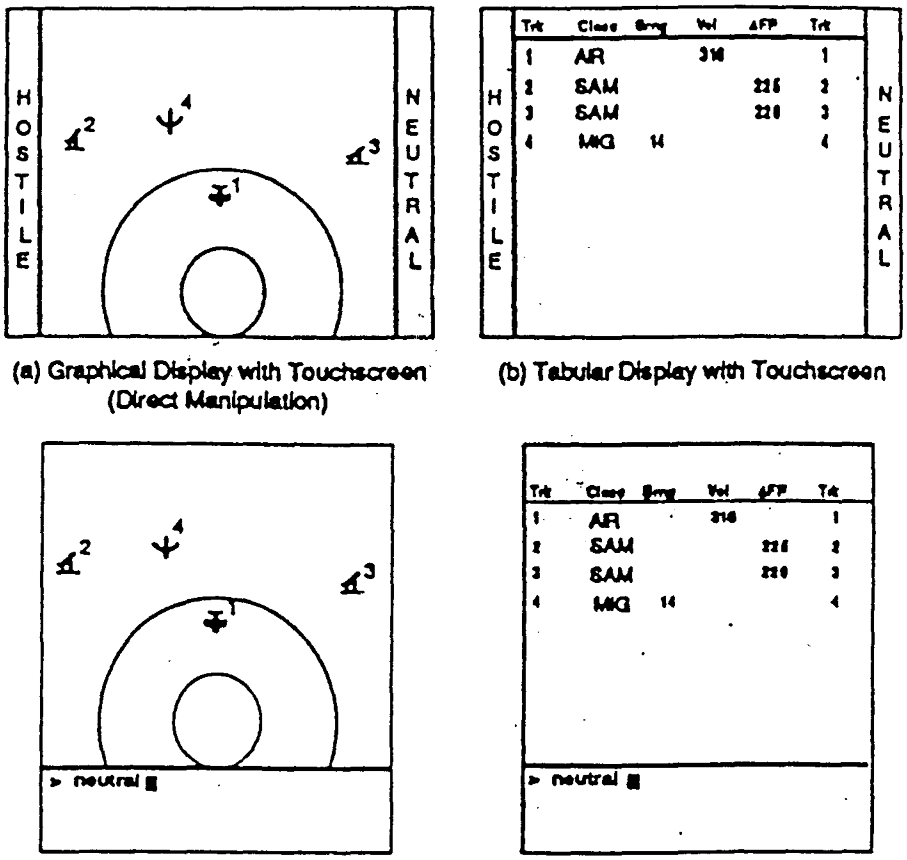

(a) Graphical Display with Touchscreen (Direct Manipulation)

(b) Tabular Display with Touchscreen

(c) Graphical Display with Keypad Input

(d) Tabular Display with Keypad Input (Command Language)

Figure 11

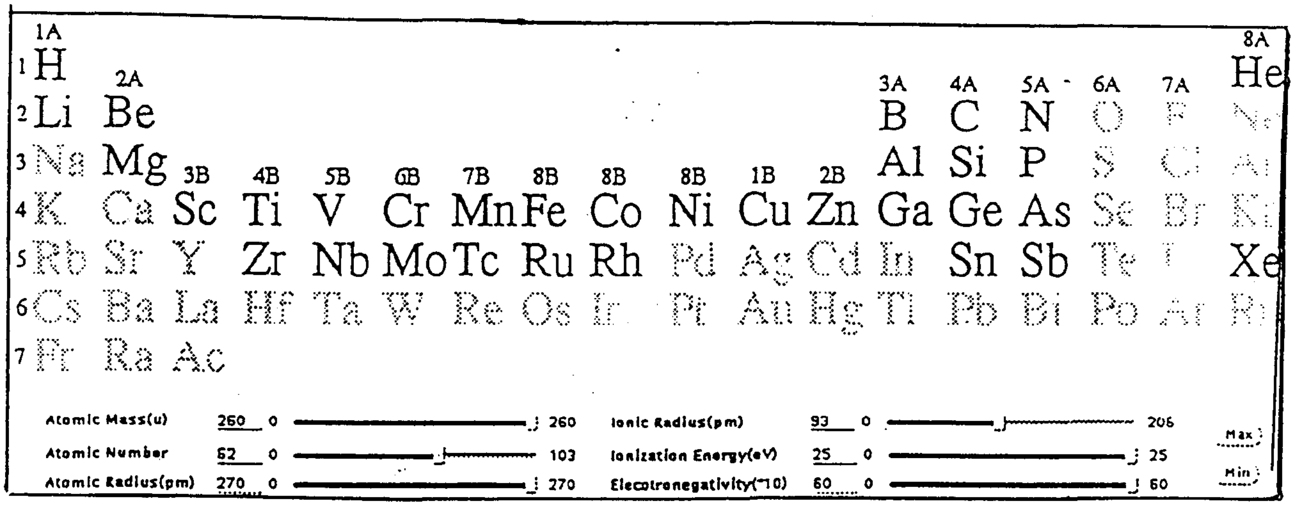

(a) SG

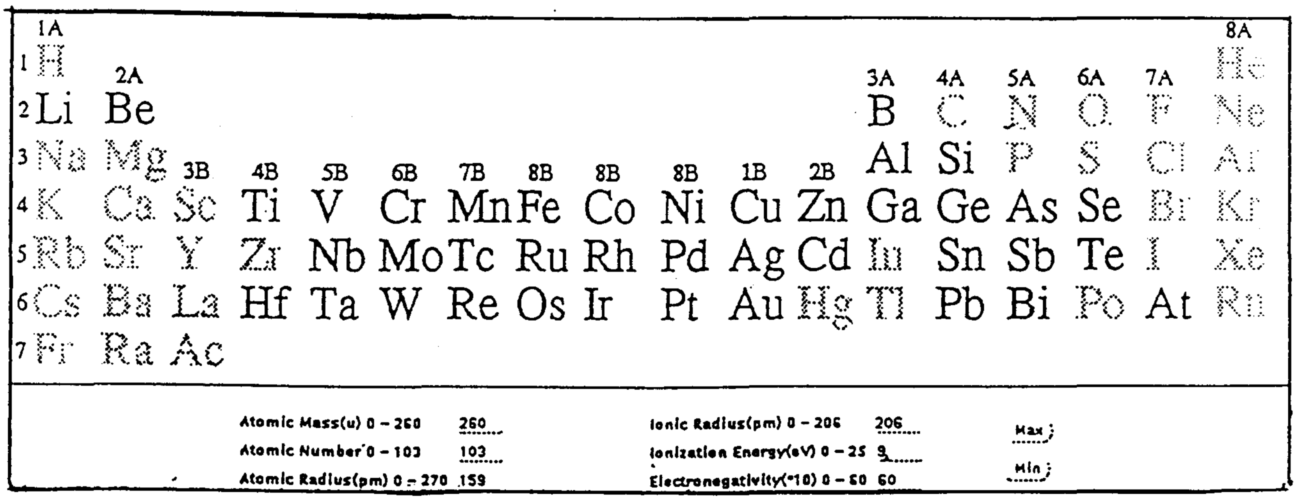

(b) FG

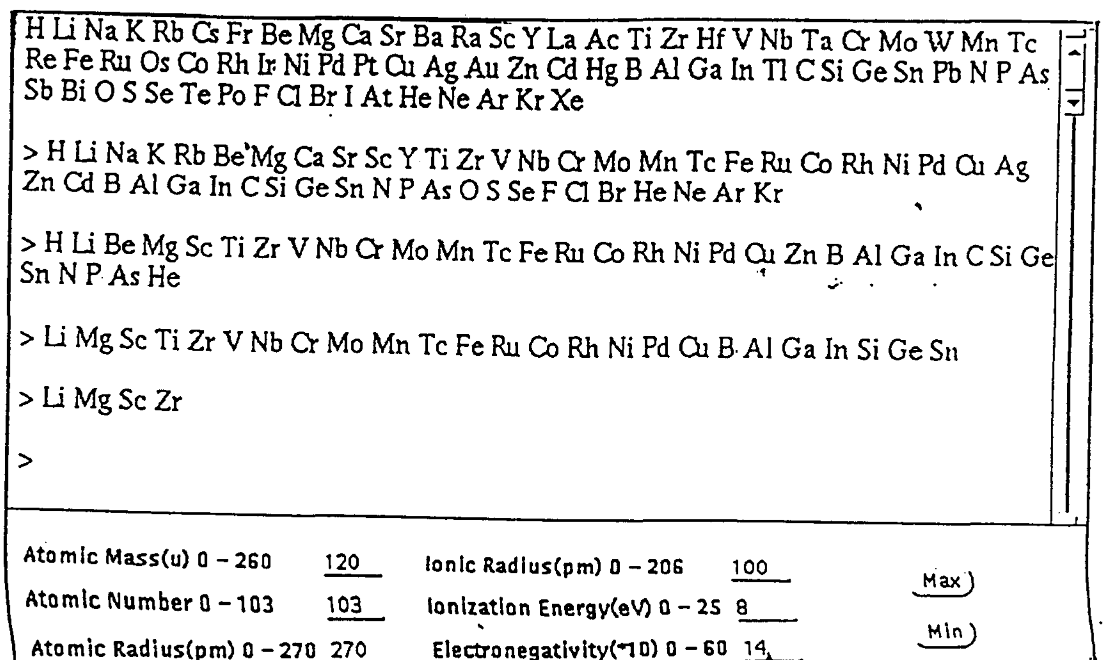

(c) FT

- Visualisation of input data may be less important

Further critical studies are needed to confirm these impressions and to address the importance of other associated features of direct manipulation such as window-techniques, function-objects, integration of input techniques and user's adjustment of design features.

4.1.3 Naturalistic choice studies

A final approach to testing the benefits of direct manipulation interfaces is to present users with a mixed mode interface incorporating elements of both manual and linguistic interaction. 'Naturalistic' choices made by users during the course of interacting with the interface would reveal not only *whether* direct manipulation methods were used more often than alternative methods but also *when* and *how*.

Unfortunately the number of studies of this kind is very limited. As we have just seen, most tests of direct manipulation interfaces vary interface type over independent groups of subjects, so that it is rare for the same subjects to experience more than one type. An exception to this was the study by Eberts & Bittianda (1993) mentioned in Section 4.1.1. They included a 'Combined' direct manipulation and command interface in two experiments testing preference for direct manipulation. However in the first experiment subjects were asked to alternate between each method in a training regime and hence had no choice open to them. In the second experiment subjects were asked to complete written tests of reasoning about the system. They found that subjects trained on the Combined interface used concepts involving direct manipulation operators more frequently than command line operators, although no analysis was conducted on *when* each type was used (personal communication).

More direct evidence regarding naturalisitic choice of interaction style comes from evaluations of a multimodal file management system called EDWARD (Bos, Huls & Claassen 1994). EDWARD is essentially a graphical 'desktop' interface with a natural language dialogue window (see Figure 12). It supports a mixture of menu, command, direct manipulation and natural language interactions, down to the individual interface expressions (such as move this file → to this directory →, where → denotes a pointing gesture). The use of EDWARD in a series of experiments is described in Huls (1995). Early feedback from 10 novice users performing file manipulation tasks on the unrestricted system revealed that

(1) All users employed more than one interaction style
(2) All users utilized natural language input
(3) Individual users switched frequently between methods
(4) Users said they liked the freedom of style choice

FIGURE 12

In a more formal experiment, 22 subjects used this 'FREE choice' condition, in contrast to 66 subjects in three RESTRICTED groups who were forced to use natural language ('Language'), direct manipulation ('Action') or both together ('Multimodal'). Again there was a great diversity in the choice of interaction style in the FREE condition. This is shown in Figure 13 which indicates massive individual differences in preferred mode. On the whole unimodal interface expressions were preferred to multimodal expressions, although direct manipulation and natural language expressions were used selectively in the proportions shown. Analysis of when users chose these methods suggested that direct manipulation was often used to

Figure 12

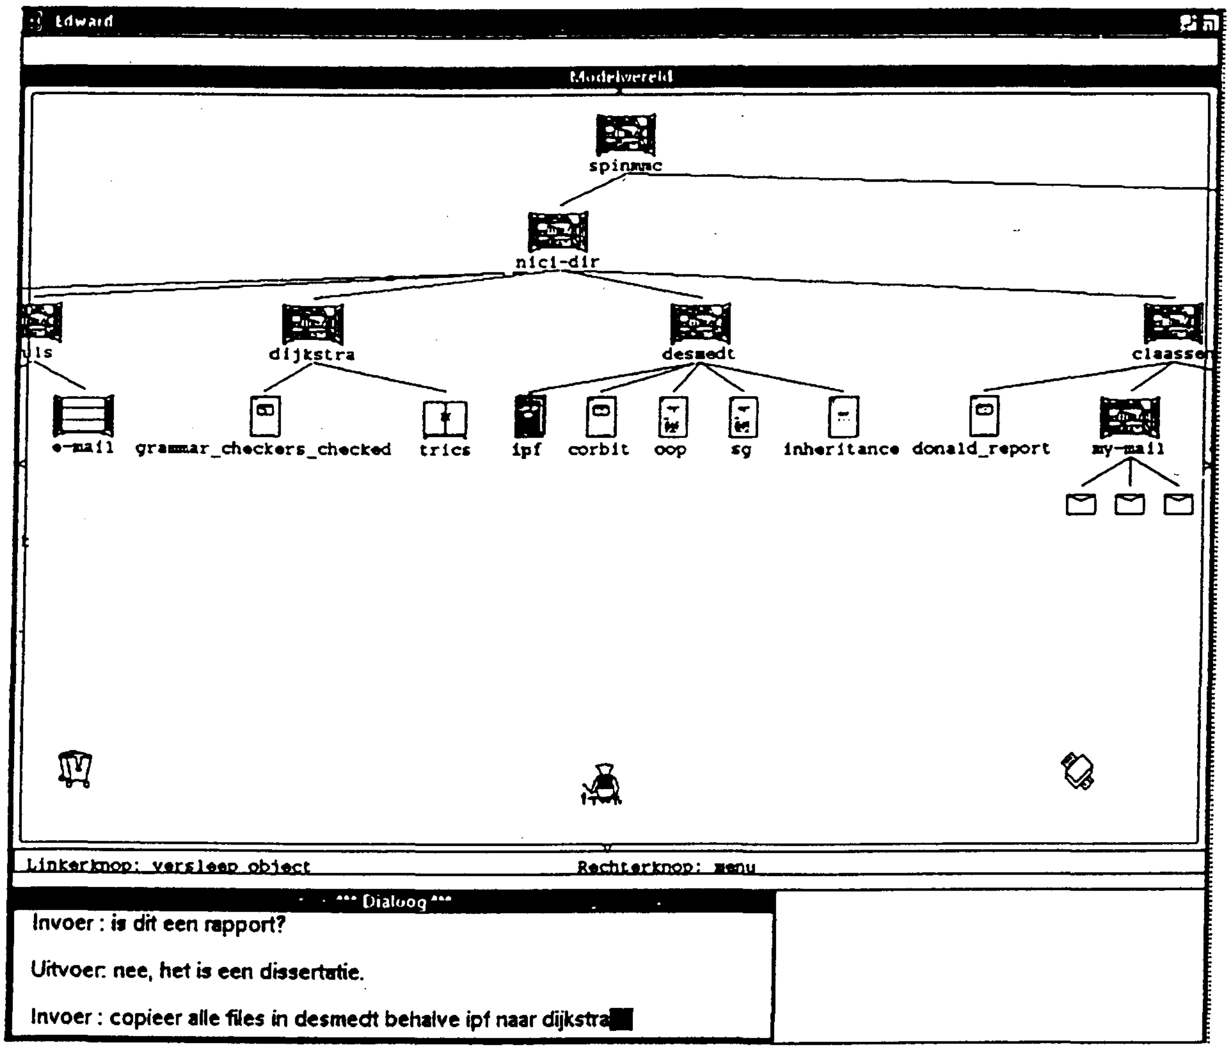

Edward
Modelwereld
spinmac
nici-dir
dijkstra
desmedt
e-mail
grammar_checkers_checked
trics
ipf
corbit
oop
sg
inheritance
donald_report
my-mail
Linkerknop: versleep object
Rechterknop: menu
*** Dialoog ***
Invoer : is dit een rapport?
Uitvoer: nee, het is een dissertatie.
Invoer : copieer alle files in desmedt behalve ipf naar dijkstra

avoid typing long object names in the dialogue box, and natural language was used to refer to objects which were not visible on the screen.

FIGURE 13

Such evaluations of mixed mode interfaces are rare, perhaps because most research papers in this area concentrate on implementation rather than usage issues, while few commercial system evaluations are ever published. This is an important gap in the literature which could help to identify the limitations as well as the benefits of direct manipulation and thereby prescribe effective combinations of manual and conversational forms of interaction.

4.1.4 Summary

In short, the above tests confirm some but not all of the proposed benefits of direct manipulation and show them to depend on implementation, task and measurement factors. A key problem with the concept is its operationalisation in specific designs. At present, a large number of graphical features are associated with direct manipulation and a new class of tests are beginning to identify which combination of features are critical to its success. The possibility of limitations and well as benefits for direct manipulation is suggested by a small number of traditional tests which show disadvantages to its use on certain tasks, and by an even smaller number of studies illustrating naturalistic choice of alternative methods in mixed mode interfaces.

4.2 Mixed mode interfaces

The design of mixed mode interfaces is a significant development for the philosophy of direct manipulation, mainly because it cannot easily accomodate them. The original insights by Shneiderman promote manual over conversational forms of interaction and do not consider how best to mix these forms in hybrid interface designs. This omission is also represented in the paper by Hutchins et al who fail to represent mixed conversational and model-world systems in their space of interfaces (shown in Figure 4). In this section I use an alternative framework to indicate the variety of ways in which manual and conversational forms of interaction are being mixed in both prototype and commercial systems. These developments serve to shift the locus of the direct manipulation debate from *whether* direct manipulation is better than other forms of interaction to *when and how* its benefits should be combined with other forms.

Figure 14 shows an adpated version of my own 'space of interfaces' diagram, published in the context of a discussion of multimedia systems (Frohlich 1991). This distinguishes between the input and output interfaces to any computer system and is organised around two central *modes* of interaction: language and action. These modes correspond to Hutchins et al's conversational and model-world interfaces, but are shown to co-exist for the same input or output interface. This allows for separate mixed mode (language/action) combinations *within* the input and output interfaces, as well as for cross-modal combinations *between* them (i.e. language in/action out or action in/language out). Furthermore, these modes are said to be subjective properties of the interface which depend on the way users intend their own inputs to be interpreted by the computer, and how they perceive the outputs of the computer. The critical distinction is whether or not they intend or interpret them as social or physical actions.

Figure 13

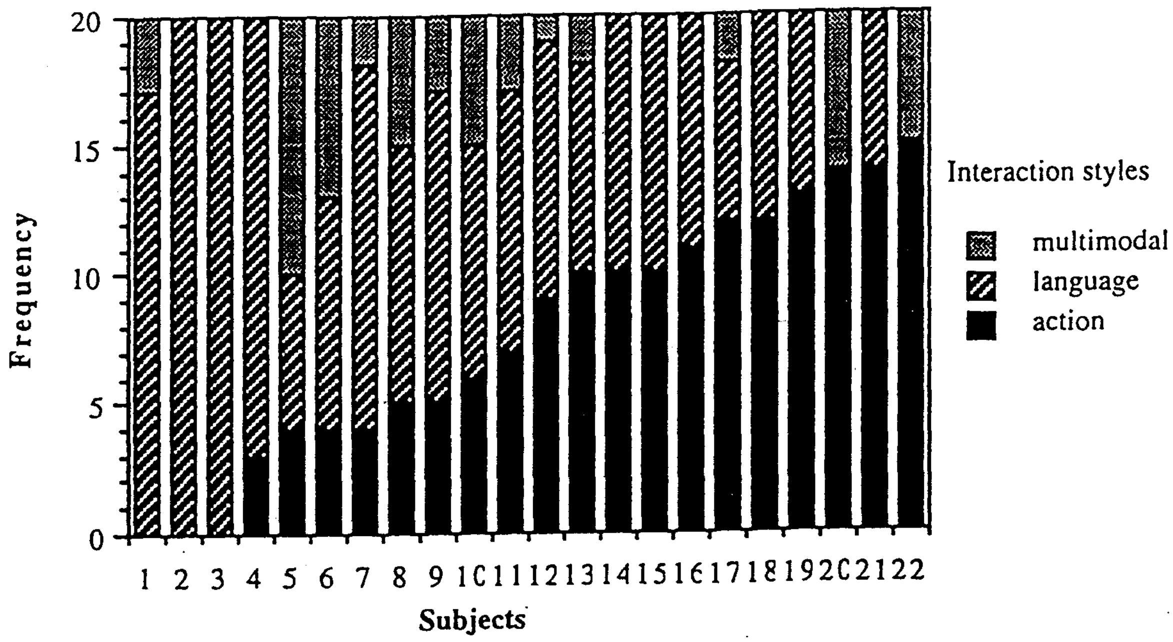

20
15
10
5
0
Frequency
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22
Subjects
Interaction styles
multimodal
language
action

Of course the expectations about whether to adopt a social or physical relationship with the machine are set by the *medium* and *style* of input allowed by the computer and the computer's own medium and style of response. In practice then, the social and physical tone of the interaction is given by the combination of interface styles employed, but differences in the balance of combinations can lead to subtle differences in the dominance of one mode over the other.

FIGURE 14 ABOUT HERE

One thing this framework makes clear is that many interfaces which have traditionally been thought of as direct manipulation interfaces are themselves mixed mode. For example the Macintosh interface incorporates many elements of language as well as action through the user's use of menus, field filling and quick command gestures and the system's use of natural language dialogue boxes containing field descriptors and menu items (see Figures 3 and 4 in Frohlich 1991 for a full characterisation using the framework). Pure direct manipulation interfaces according to the framework would be model-world interfaces based on an Action in/Action out modality involving only the media of sound, graphics and motion. Interestingly, these are exactly the properties of information visualisation interfaces such as the Dynamic Homefinder interface described in Section 2.2.

Other mixed mode interfaces are those which explicitly mix model-world and agent metaphors for interaction. Hewlett Packard's own desktop office interface called NewWave does this by incorporating an 'Agent' icon on the desktop, with functionality to automate routine or background tasks such as file backup or meeting reminders (Stearns 1989). In practice the user clicks on pre-recorded Agent 'Tasks' or can record new tasks by demonstration (see Myers 1992 for a discussion of demonstrational interfaces). Metaphorically, the user is intended to relate to the Agent as a personal assistant in the office world of interest. This impression is encouraged through animation of Agent actions on documents, folders and other office icons so that users can see the Agent acting on the world in much the same way as they do. The use of such interface agents as 'personal digital assitants' seems set to rise in the context of computer telephony and networking where users cannot be expected to deal manually with the vast amount of information that will become available to them (c.f. Maes 1994).

The same kind of mixing of agent and model world attributes happens in a different way in computer mediated communication tools. In these cases the interfaces agents are other people who may act on a shared workspace or document whilst talking to you (c.f. Whittaker, Geelhoed & Robinson 1993).

The natural interleaving of linguistic and physical actions in human-human conversation appears to be the inspiration for a number of attempts to mix interface styles at a fine grained level in human-computer interaction (Neilson & Lee 1994). For example, the ACORD system is a transportation monitoring application which allows users to update a knowledge base of truck positions either by dragging truck icons around a geographical map or by informing the system of truck movements in natural language (Lee & Zeevat 1990). Both interaction styles can be combined in individual input expressions: such as typing 'This goes to Munich' after clicking on the desired truck icon. Haddock (1992) shows how these kind of referring expressions can be used to good effect in multimodal database queries in the SAMi system. Natural language follow-up questions are especially efficient when scoped by

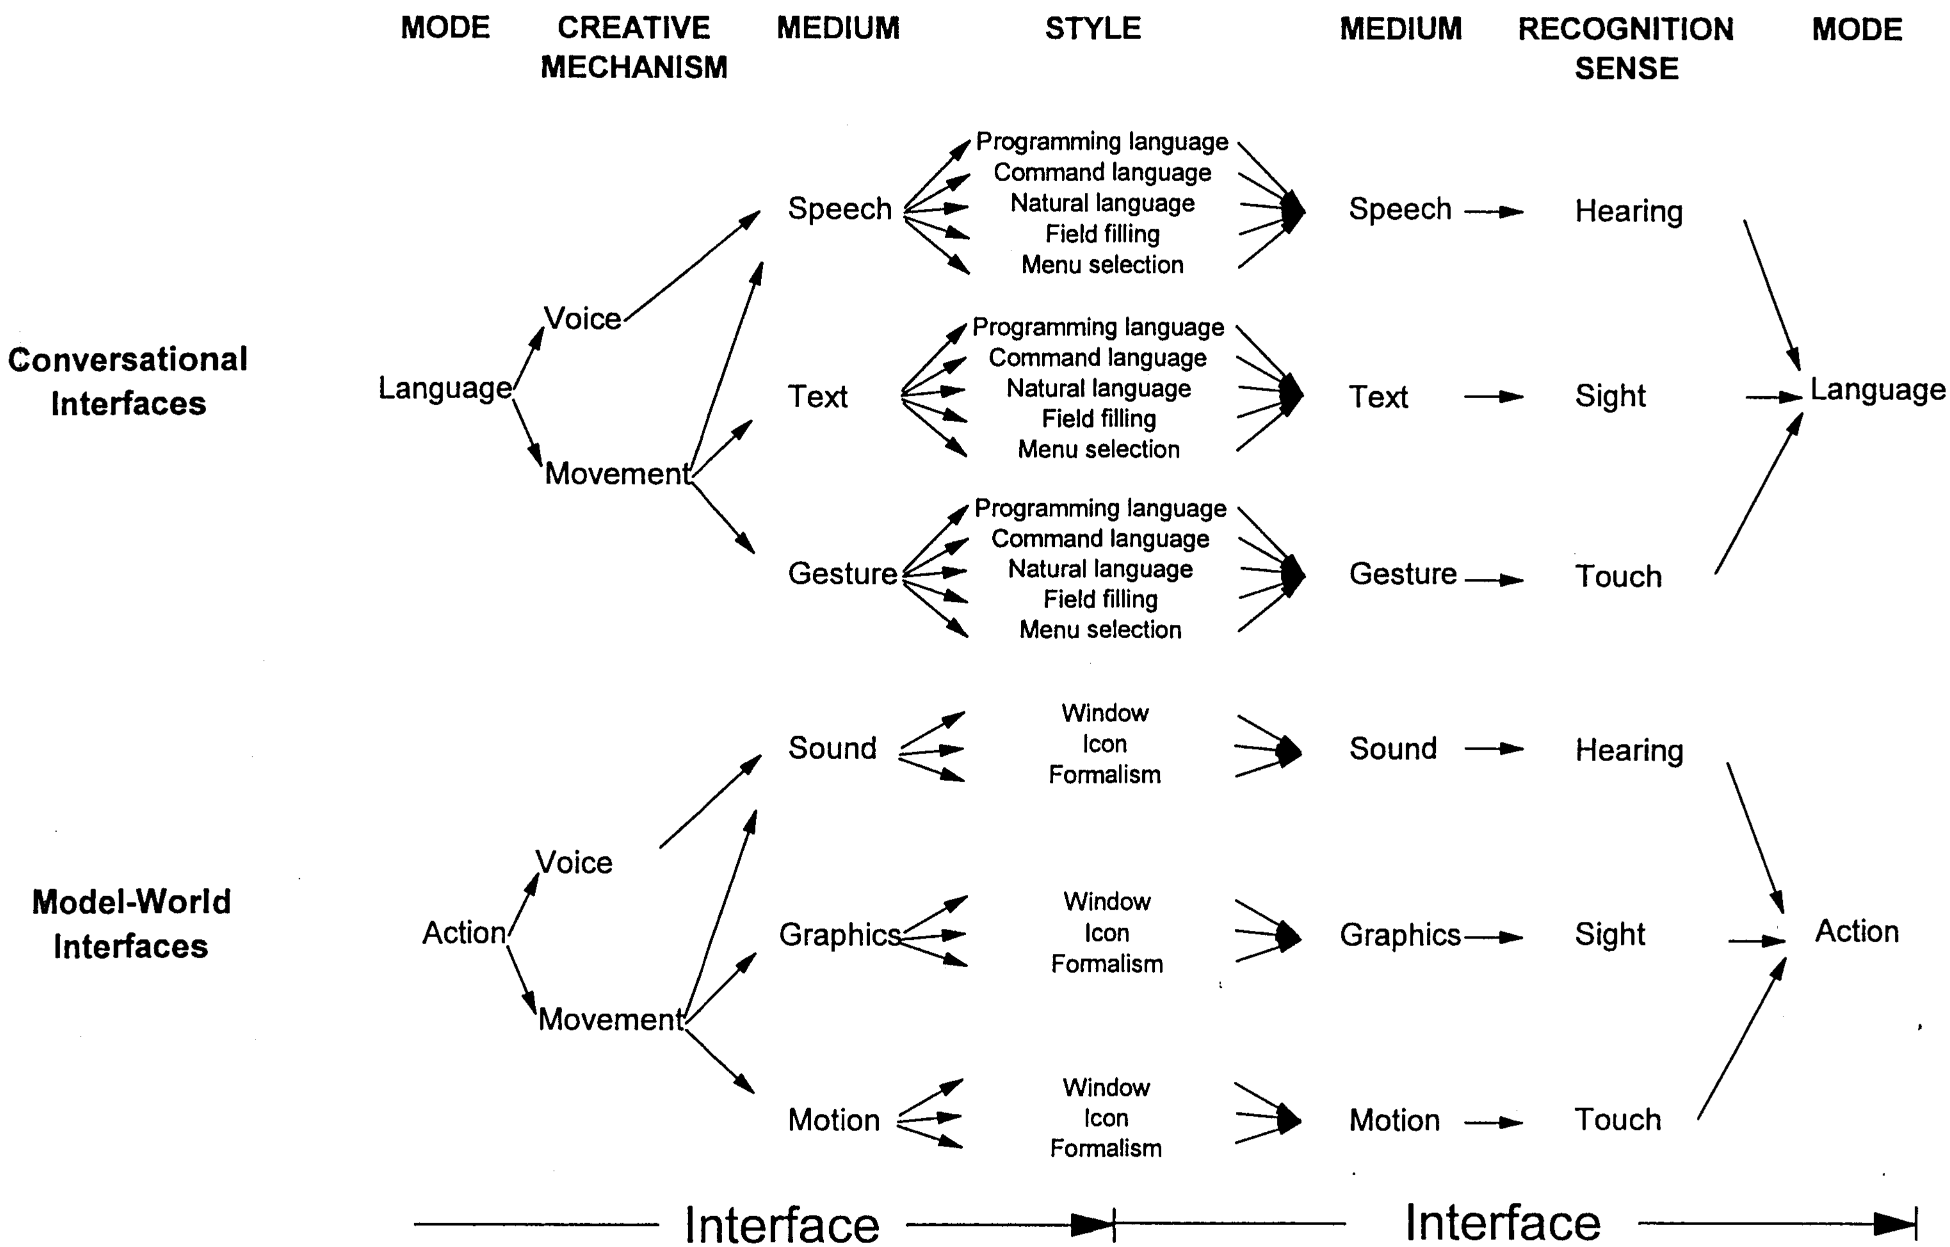
MODE
CREATIVE MECHANISM
MEDIUM
STYLE
MEDIUM
RECOGNITION SENSE
MODE
Conversational Interfaces
Language
Voice
Movement
Speech
Text
Gesture
Programming language
Command language
Natural language
Field filling
Menu selection
Speech
Text
Gesture
Hearing
Sight
Touch
Language
Model-World Interfaces
Action
Voice
Movement
Sound
Graphics
Motion
Window
Icon
Formalism
Sound
Graphics
Motion
Hearing
Sight
Touch
Action
Interface
Interface

Figure 14

pointing to elements of the displayed output from a previous query: as in 'Which of these sales is for Walkers?', or 'Show these as a bar chart'. Cohen, Dalrymple, Moran, Periera, Sullivan, Gargan, Schlossberg & Tyler (1989) take these techniques further in the SHOPTALK system which utilizes 'natural language forms' and 'discourse tree structures' in the context of an iconic representation of a circuit-board factory.

In these last examples, conversational and manual interface styles are being mixed mainly within the input interface. Cross-modal mixing across the input and output interfaces also occurs in systems such as GeoSpace (Lokuge & Ishizaki 1995). This is like a Dynamic Homefinder system you can talk to, since it combines natural language input with a map formalism output to give users access to a database of area and property information. Users ask for different graphical views of the information in a Language in/Action out interaction. The opposite Action in/Language out arrangement is found in certain intelligent tutoring systems which critique the user's manipulation of some world of interest in natural language output.

All these forms of mixed mode interaction show how the character of model world or conversational interfaces can be altered by introducing elements of the complementary mode into the input or output interface. The range, effectiveness and spread of such systems suggests that there is great value in choosing the middle way.

4.3 Theory

Further discussion of the concept of direct manipulation has continued in parallel to the studies and systems mentioned above. The same lessons are reinforced through attempts to explain the value of mixing manual and conversational forms of interaction and to further understand what makes manipulation direct. Significant contributions to each issue are reviewed in the next two sections.

4.3.1 The value of mixed mode interaction

In developments of their original positions, both Hutchins and Laurel promote interface agents as additions to direct manipulation interfaces.

Hutchins (1989) discusses four possible relationships a user can have with the computer which he calls 'mode of interaction' metaphors:

1. *The conversation metaphor* - where expressions have the character of utterances in a conversation about the task at hand
2. *The declaration metaphor* - where expressions take on the character of speech acts which magically cause things to happen in the world of interest
3. *The model-world metaphor* - where expressions have the character of actions taken in the world of interest
4. *The collaborative manipulation metaphor* - where expressions have the character of carrying out a task with someone elses help

Pointing out the limitations of the first three modes, Hutchins notes that conversation is limited because users have to learn a new language, maintain a mental model of the world being talked about, and converse in a manner very different from that of ordinary conversation. Declaration is limited because it depends on a practice effect in using a conversational interface which is always liable to dissipate if the user declares something that cannot be done. Model-world interaction is limited precisely because of its directness in collapsing abstract descriptions into concrete actions. Users can no longer refer to *classes* of objects but must identify each instance of a class manually. These limitations in each individual mode leads Hutchins to promote their combination in a form of collaborative manipulation in which users act within a model-world alongside an intermediary who can also act on their behalf (see Figure 15). Provision of a linguistic interface to the NewWave Agent mentioned in Section 4.2 would create this arrangement within a desktop office system.

FIGURE 15 ABOUT HERE

Laurel (1990) makes a similar argument for the re-introduction of a conversational metaphor within a model-world context. Whereas before she had argued against the *hidden intermediary* as a primary vehicle for interaction (Laurel 1986) she now points out the value of a *visible intermediary* (my term) within a model-world. This she refers to as an interface agent, defined as "a character enacted by a computer, who acts on behalf of the user in a virtual (computer-based) environment" (p356). Like Hutchins, Laurel begins to note a number of computer-based tasks which are difficult to perform through action and might be more easily achieved by delegation through language to an interface agent. These include retrieving, sorting, organising, programming and scheduling. In addition, Laurel promotes the interface agent as a convenient new metaphor for representing the more (pro-)active and autonomous components of modern computer systems. These include various forms of agency, such as facilities to filter information, remind, help, tutor, advise, perform or play.

The same conclusions are reached via a different route by Whittaker (1990), Stenton (1990) and Walker (1989). All three authors discuss the limitations of direct manipulation in the context of desktop office systems such as the Xerox Star, Apple Macintosh, or HP NewWave systems.

Whittaker (1990) examines each of Shneiderman's three principles of direct manipulation as they apply to a desktop model world. He describes particular limitations with Principles 1 and 3:

> Principle (1) Continuous representation of the object of interest - This leads to the problems of locating objects that are not visible in a large or distributed system. Manual search may be inefficient even when users know the location of target documents, and completely impracticable when they don't.
> Principle (3) Rapid incremental reversible operations ... - The need for incremental operations leads to inefficiencies in the execution of common compound actions which might be done faster through a single command. It also prevents automatic application of repetitive actions to a set of objects.
> Principle (3) ...whose impact is immediately visible - This does not allow for processes which users may want to run in the future or while they are absent, such as reminders or mail forwarding.

These problems are said to relate to the requirement that change in the model world be a direct result of a corresponding user action. Relaxing this requirement leads to a form of agency in which the system itself takes responsibility for executing certain actions.

Figure 15

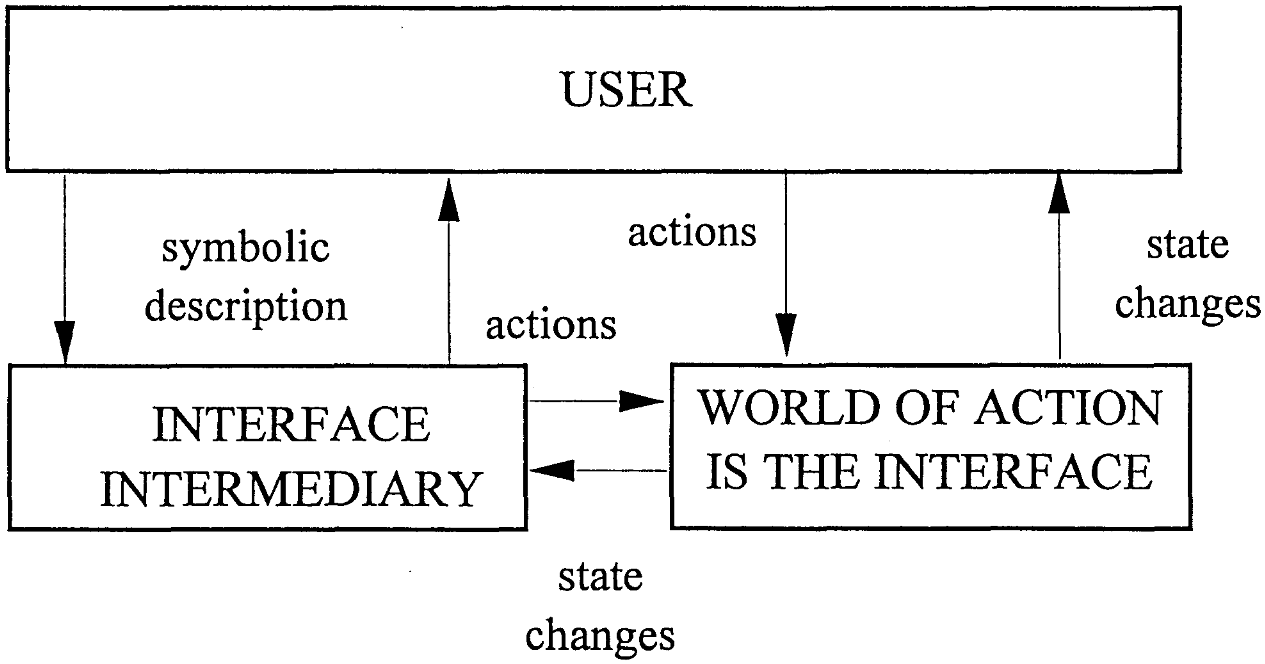

USER
symbolic
description
actions
actions
state
changes
INTERFACE
INTERMEDIARY
WORLD OF ACTION
IS THE INTERFACE
state
changes

Stenton (1990) elaborates on this view by suggesting that the interaction between a desktop system and its user can be improved by adding an interface agent who can take the initiative from time to time. Such initiative is required to provide feedback on the kinds of delayed operations discussed by Whittaker, and could be used more generally to provide cooperative suggestions or responses for the task at hand. Without the metaphor of an interface agent to perform these actions Stenton points out that users would have to believe in a talking or listening desktop.

Walker (1989) lists a set of desktop interaction tasks that may be more suited to control through natural language than through direct manipulation. Indeed, a number of these tasks would be impossible to perform manually because they involve actions out of the 'here and now'; referring to unseen objects or to events forwards or backwards in time. The tasks are:
DEFINITE DESCRIPTION - referring to objects by their attributes
e.g. Get me the messages FROM ERIK.
DISCOURSE REFERENCE - referring to previously evoked objects with pronouns
e.g. Get me the messages from Erik. Are any of THEM about today's meeting?
TEMPORAL SPECIFICATION - specifying future times or intervals for an event
e.g. Send a copy of the proposal to Susan EVERY WEEK.
QUANTIFICATION - referring to object quantities
e.g. HOW MANY messages do I have today? Are there ANY from Erik?
COORDINATION - combining operations in compound phrases
e.g. print the report AND send a copy to Steve
NEGATION - refering to negative attributes
e.g. Are there any messages I HAVEN'T answered?
COMPARATIVES - comparing objects by attributes
e.g. Are there any messages OLDER than two months?
SORTING EXPRESSIONS - referring to the attributes by which objects should be displayed
e.g. Show me my messages BY topic
Walker is keen to recommend empirical tests which would identify which of these tasks are *actually* performed better via natural language and are valuable enough to develop as additional features of a desktop system. Her own work of this kind was described in Section 4.1.3, and suggested value in natural language DISCOURSE REFERENCE, COORDINATION, and SORTING EXPRESSIONS.

A very similar argument is made by Cohen (1992) in discussing the symmetrical advantages and disadvantages of natural language and direct manipulation. Thus he identifies the following tasks as being difficult to do through direct manipulation but easy to do through natural language:

1. Description including quantification, negation and temporal information
2. Anaphora (e.g. pronouns)
3. Operations on large sets of objects
4. Delayed actions

Cohen goes on to describe features of the SHOPTALK system discussed in Section 4.2 which combines elements of natural language with direct manipulation. Unfortunately he fails to provide data on the *use* of SHOPTALK which would help to substantiate his own claims and the earlier ones of Walker (1989).

4.3.2 Further insights on directness

Along with the critical evaluations of direct manipulation described in Section 4.1.2 go a number of theoretical discussions of directness. These begin with those relating to the essence of good manual interaction and end with a number of attempts to apply the concept to conversational forms of interaction.

Regarding the directness of manual interaction, Johnson, Roberts, Verplank, Smith, Irby, Beard, Mackay (1989) look back at the design of the Xerox Star system to consider what made it so successful in the industry. They attribute a large part of this success to the *desktop metaphor* and everything that entailed for the direct manipulation of graphical objects on the screen, including:

- a high resolution bitmapped display
- a pointing device
- an object-oriented iconic interface
- generic commands (such as move, open, copy, delete, show properties, same)

A particular guiding principle behind the Star's design was to encourage 'seeing and pointing over remembering and typing' through the graphical interface. However, they note that one of the lessons of the Star experience is 'Don't be dogmatic about the desktop metaphor and direct manipulation' because there are things that remembering and typing are better for.

Johnson (1987) elaborates on this point in another article in which he argues that the faithfulness of the desktop metaphor depends on the individual office functions represented. Some functions are best implemented in a new way. Thus he lists a number of facilities in the Star interface which deviate from the desktop metaphor to good effect, including the sorting of files by different attributes, the nesting of folders to any depth, the shrinking of documents to icons when not in use and the tiling of windows. He also notes the absence of additional office objects such as staplers which are simply not needed on the electronic desktop where pages of documents do not naturally come apart. Johnson suggests that designers must trade off faithfulness to the metaphor with the utility of other methods, and users must learn to adapt accordingly (see also Halasz & Moran 1981).

The role of icons in representing metaphors is explored in recent work by Familant & Detweiler (1993). An icon is defined as 'a sign that shares characteristics with the objects to which it refers' (p706), and is distinguished from other signs such as an index and a symbol. The authors show how many of the so called icons used in modern graphical interfaces have lost this property of shared characteristics with referent objects, largely because they refer to highly abstract computer facilities which have no real world counterparts (e.g. undo, find, insert frame). These do not employ metaphors and should really be called symbols. Furthermore, true icons invoking some real world metaphor can incorporate direct or indirect reference depending on various relationships between the thing depicted, the thing to which it refers, and the thing ultimately denoted or pointed to. The implications for directness of manipulation seem to be that icons will only really improve the 'fit' between user and system in cases where they depict a real world object with clear connections to the referenced computer object.

On a different tack, Desain (1988) equates distance with directness (as defined by Hutchins et al) and applies it as a measure of the quality of both manual and conversational interfaces. In this simple move he shows how both kinds of interfaces can turn out to be direct to use. He suggests that the most direct manual interfaces are those which use a well-known graphical formalism for the application domain and support the corresponding real-world actions (such as grasping). The most direct conversational interfaces are said to be those employing natural language syntax already known to the user with the 'jargon' vocabulary of the application domain. Various strategies for minimizing semantic and syntactic distance are discussed for both types of systems, including issues of representation, abstraction, reference, error management, and precision.

Both Claassen, Bos & Huls, (1990) and Brennan (1990) adopt the same position to attack the view that conversational interfaces are necessarily indirect to use. They highlight some of the more valuable attributes of converational interfaces including the handling of error and repair, the ability to ask questions, and the range of expressions for refering to abstract properties and events. As further evidence for the directness of conversational interaction in general Brennan also cites the extremely efficient and cryptic quality of interactions between people who know each other well.

Finally in my own previous review of the area I took up this last point in a further discussion of what makes conversation direct (Frohlich 1993). I suggested that good 'rapport' between partners is the quality to aim for; comparable to good 'engagement' within manual interactions and characterised by the kind of brief, cryptic exchanges referred to by Brennan (1990). These exchanges are a puzzle for the philosophy of direct manipulation because in one sense they are highly indirect. They increase the cognitive load on conversational partners whilst decreasing the interactional work that is done between them. Thus by making more inferences about what the other partner is saying it is possible for both partners to conduct a shorter and more efficient conversation. This observation led me to distinguish between two types of directness in human computer interaction:
(i) cognitive directness already defined by Hutchins et al as 'least cognitive effort', and
(ii) social directness defined as 'least collaborative effort'.
Within this view, cryptic conversations can be seen as increasing social directness at the expense of cognitive directness. To avoid too much confusion between the terms cognitive and social directness I suggested referring to social directness as gracefulness. Hence graceful interaction would be characterised by a feeling of pace and efficiency in the flow of exchanges between user and computer, whether they be through manual actions or linguisitc utterances.

4.3.3 Summary

In short, recent thinking on direct manipulation reinforces many of the lessons of the empirical tests and mixed mode interface designs reviewed in Sections 4.1 and 4.2. They either cast a critical eye over what accounts for the directness effect or they acknowledge that manual interaction is limited and should be combined with conversational interaction. The first line of argument results in further insights on good features of both manual and conversational interaction; such as 'seeing and pointing over remembering and typing', a clear application metaphor which is creatively implemented, icons which make sense in terms of the metaphor, words which include the 'jargon' of the domain, and sense of gracefulness and efficiency in the flow of the interaction. The second line of argument results in discussion of the value of a 'collaborative manipulation' paradigm in which users work together with a virtual partner in the model world of interest. This allows them to use conversation with the partner to perform all those interaction tasks that are difficult or impossible to perform by manipulation alone.

5. TWO PHILOSOPHIES, TWO DEBATES

5.1 Separating directness and manipulation

Having updated the debate about direct manipulation interfaces what can now be said about the direct manipulation philosophy?

Historically we can see that direct manipulation interfaces certainly served to break the mold of programming and command line interfaces which required users to remember and type complex phrases in an artificial language and then to hold in their heads an impression of the effects of these commands on the application. When compared to this baseline, manual interaction looked like an extremely attractive alternative which deserved much of the praise and success it received in the industry. However the industry has moved on and with it the baseline. Menu, form and natural language interfaces now constitute a fairer comparison with manual/graphical interfaces, and there are a growing number of mixed mode interfaces which merge manual and conversational intearction to good effect. In line with these developments, straightforward comparisons of direct manipulation interfaces with other kinds have come up with conflicting results, while more critical tests have begun to show ways in which manual interfaces themselves can vary in design. In this modern context the old formulation of the philosophy begins to look decided shakey.

A new formulation is suggested by the progress that has been made in understanding *how* and *when* to utilize manual forms of interaction at the interface. These two questions run through the direct manipulation debate and ultimately serve to split it apart. For example, questions of how to use manual interaction are addressed by critical tests comparing different implementations of manual interfaces with each other, and by discussions of 'directness' in general (described in Sections 4.1.2 and 4.3.2). They show that some features of manual *and* conversational interaction are more effective than others. Questions of when to use manual interaction are addressed by uncritical tests comparing manual with conversational interfaces for the same task, and by the work on mixed mode interfaces (described in Sections 4.1.1, 4.1.3, 4.2 and 4.3.1). They show that manual interfaces are not always better than conversational ones, and that combined interfaces can leave the choice very effectively in the hands of users.

Thus we have ended up with two distinct strands to the current direct manipulation debate relating to questions of directness and questions of manipulation. Directness is now equated with distance and manipulation with engagement. That these components are separable was first suggested by Hutchins et al's space of interfaces (shown in Figure 4). However, recent research has been steadily populating the space and discovering the truth of its comparisons. In addition, it has added a class of mixed mode interfaces which themselves vary in directness. This leads to a reconceptualisation of the space of interfaces shown in Figure 16 (adapted from Figure 2 in Frohlich 1993).

FIGURE 16 ABOUT HERE

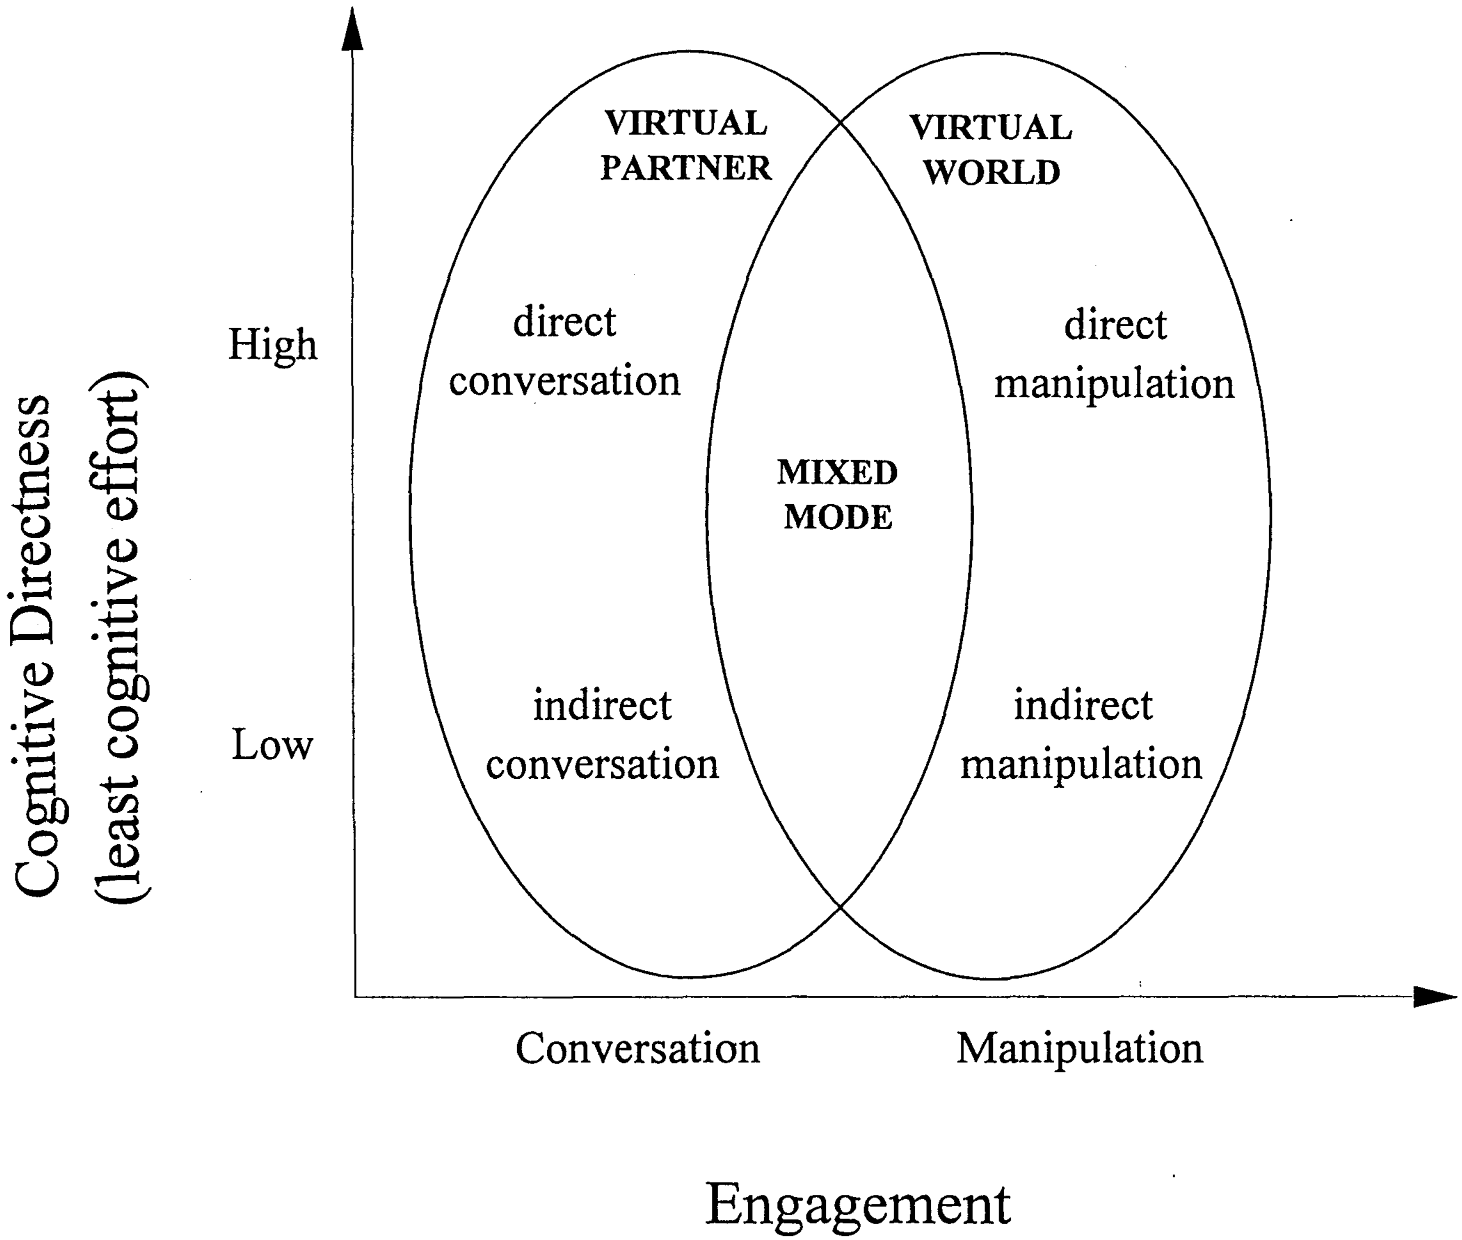
VIRTUAL PARTNER
VIRTUAL WORLD
High
direct conversation
direct manipulation
MIXED MODE
Low
indirect conversation
indirect manipulation
Cognitive Directness
(least cognitive effort)
Conversation
Manipulation
Engagement

Figure 16

Work on each dimension of this new interface space is leading to two emergent philosophies of interaction; a *directness philosophy* relating to what makes an interface easy to use, and a *manipulation philosophy* relating to why manual forms of interaction are preferable to conversational forms. In the next two sections I will suggest a characterisation of each philosophy based on the research reviewed. Each characterisation is done in terms of a set of design recommendations which can be used both to inform design and influence future research.

5.2 A new philosophy of directness

Table 1 shows a set of principles for the design of good manual and conversational interaction. The principles for manual interaction are based on a revision of Shneiderman's original principles of direct manipulation in the light of the findings about *how* to use manipulation. They are partitioned by whether or not they relate to cognitive directness in the way the user thinks about the system, or to gracefulness in the form of interaction between them. The main change is the addition of a 'Real-world metaphor' principle which seems to be associated with the most successful direct manipulation interfaces. 'Natural actions' is a development of Shneiderman's 'Physical actions' (Principle 2) which requires that the manipulations used should posess the same control dynamics as those used in the corresponding real-world task wherever possible. This serves to extend the concept of metaphor to the device as well as the task. 'Continuous representation' is Shneiderman's Principle 1, while 'Responsive visualisation' corresponds to his Principle 3, with emphasis on fast analogue change in the output visualisation as a property of the interaction.

	MANIPULATION	CONVERSATION
COGNITIVE DIRECTNESS	• Coherent real-world metaphor • Natural actions • Continuous representation	• Familiar terminology • Natural language • Personal relevance
INTERACTIONAL GRACEFULNESS	• Responsive visualisation	• Short rapid turns • Mixed initiative • Explicit repair

Table 1. Principles of good manual and conversational interaction

The principles proposed for conversational interaction are derived partly from the literature reviewed, and partly from related work on the design of human-computer dialogues with conversation-like properties (e.g. Frohlich & Luff 1990, Hayes & Reddy 1983, Luff & Frohlich 1991, Nickerson 1981, Yankelovich, Levow & Marx 1995). They are included here in response to the insight that directness (and gracefulness) can be applied to conversational as well as manual interface design, and in the hope that they might stimulate the same sort of debate as Shneiderman's original principles for manipulation. Thus, 'Familiar terminology' refers to use of application 'jargon', 'Natural language' refers to the use of a native syntax for combining and interpreting interface expressions, 'Personal relevance' refers to the use of a relationship history by the computer, 'Short rapid turns' refers to the highly interactive exchange of turns, 'Mixed initiative' refers to a switching of who is leading the conversation, while 'Explicit repair' refers to the incorporation of procedures for handling user or computer trouble in understanding.

Further research is needed to test and revise both sets of principles. Continuation of the critical tests reviewed in Section 4.1.2 for manual interfaces would help to assess the revision of Shneiderman's principles, while their extension to language-based interfaces would help to test the proposed conversational principles. Whether or not a third set of principles is required for mixed mode interaction also remains to be determined.

5.3 A new philosophy of manipulation

Table 2 presents a rough guide to the selection of interface modality. It is based on findings from the review regarding *when* to use manual and conversational forms of interaction. The essence of the new philosophy is therefore to utilize manipulation selectively! Various features of the target application can be checked off in the table. These range from high level features such as the overall activity being supported and the interaction tasks involved, to low level features such as the interaction media and technology available. An additional category of contextual factors has also been included to take into account other activities the user might be expected to perform concurrently. The designer can work through the table literally top-down or bottom-up to arrive at a set of indicators for manual or conversational interaction. If ticks appear in only one of the columns then the modality is clear. If ticks appear in both columns then a mixed mode interface is indicated, unless the designer is willing to trade-off some inefficiencies for a preferred mode.

Selective application of the manipulation modality should itself ensure a greater directness and gracefulness of manipulation when it *is* used. Whether or not these are the right features on which to base such selection remains to be tested. This will require a continuation of the uncritical tests reviewed in Section 4.1.1 together with many more evaluations of the naturalistic use of mixed mode interfaces as described in Section 4.1.3. In fact, the development of mixed mode interfaces itself constitutes the biggest challenge for advocates of manual interaction. This is because, in the long term, the use of language in a manual interface and the use of manipulation in a conversational one promises to extend the reach of manipulation techniques beyond their natural limits. We will only begin to see this if future tests of Manual versus Conversational interaction begin to routinely incorporate a 'Combined' condition. For many tasks this might often turn out to be the best.

(TABLE 2 OVERLEAF)

Acknowledgements

Thanks to Ben Shneiderman and two anonymous referees for their helpful comments on an earlier draft of the chapter; to my colleagues in Personal Systems Lab for numerous discussions and feedback on ideas; and to Charlotte Swain for her expert help in tracking down references.

APPLICATION FEATURES	MANUAL MODE	CONVERSATIONAL MODE
ACTIVITY METAPHORS	☐ Looking ☐ Browsing ☐ Exploring ☐ Navigating ☐ Controlling ☐ Monitoring ☐ Constructing ☐ Creating	☐ Informing ☐ Requesting ☐ Asking ☐ Advice seeking ☐ Understanding ☐ Negotiating ☐ Delegating ☐ Problem solving
INTERACTION TASKS	☐ Selecting seen objects ☐ Executing actions with immediate consequences ☐ Responding immediately to feedback ☐ Identifying relationships between objects	☐ Selecting unseen objects ☐ Identifying sets of objects ☐ Referring back ☐ Scheduling forwards ☐ Repeating actions ☐ Combining actions ☐ Specifying exact values
INTERACTION MEDIA	☐ Sound ☐ Graphics ☐ Motion	☐ Speech ☐ Text ☐ Gesture
INTERACTION TECHNOLOGY	☐ Large displays ☐ Sound effects ☐ Motion sensing	☐ Small displays ☐ Keyboard only ☐ Voice activation
INTERACTION CONTEXT	☐ Hands and eyes free ☐ Ears busy	☐ Hands or eyes busy ☐ Ears free

Table 2. Modality selection checklist

Figure captions

Figure 1.	The Dynamic Homefinder interface. Figure 2 in Shneiderman, Williamson & Ahlberg (1992).
Figure 2.	Two types of engagement.
Figure 3.	Two types of distance. Figure 3.2 in Norman (1986).
Figure 4.	A space of interfaces. Figure 5.8 in Hutchins, Hollan & Norman (1986).
Figure 5.	Four direct manipulation interface designs. Figure 1 in Kunkel, Bannert & Fach (1995).
Figure 6.	Two direct manipulation and two menu interfaces. Figures 1-4 in Benbasat & Todd (1993).
Figure 7.	Three types of menu page interfaces. Figure 1 in MacGregor (1992).
Figure 8.	Two interfaces to a bath-filling task with and without the metaphor. Figures 1 and 2 in Ankrah, Frohlich & Gilbert (1990).
Figure 9.	A model of metaphor. Figure 2 in Anderson, Smyth, Knott, Bergan & Alty 1994.
Figure 10.	Four interfaces varying in distance and engagement. Figure 2 in Ballas, Constance, Heitmeyer & Perez (1992).
Figure 11.	Two graphical and one textual interface to the Periodic Table. Figures 2-4 in Ahlberg, Williamson & Shneiderman (1992). S=Slider input, F=Form fill-in, G=Graphical output, T=Text output.
Figure 12.	The EDWARD interface. Figure 1 in Bos et al (1994).
Figure 13.	Free choice of alternative interaction styles in EDWARD. Figure 2.3 in Huls (1995).
Figure 14.	The design space of interfaces. Adapted from Figures 1 and 2 in Frohlich (1991).
Figure 15.	The collaborative manipulation interface. Figure 2.6 in Hutchins (1989).
Figure 16.	Interface categories partitioned by Directness and Engagement. Adapted from Figure 2 in Frohlich (1993).

REFERENCES

Ahlberg, C., & Shneiderman, B. (1994a). The Alphaslider: A compact and rapid selector. *Proceedings of CHI '94: Human Factors in Computing Systems,* 365-371. New York: ACM.

Ahlberg, C., & Shneiderman, B. (1994b). Visual information seeking: Tight coupling of dynamic query filters with starfield displays. *Proceedings of CHI '94: Human Factors in Computing Systems*, 313-317, New York: ACM

Ahlberg, C., Williamson, C., & Schneiderman, B. (1992). Dynamic queries for information exploration: An implementation and evaluation. *Proceedings of CHI '92: Human Factors in Computing Systems,* 619-626, New York: ACM.

Ankrah, A., Frohlich, D. M., & Gilbert, G. N. (1990). Two ways to fill a bath with and without knowing it. *Proceedings of INTERACT '90,* 73-78. North Holland: Amsterdam.

Anderson, B., Smyth, M., Knott, R. P., Bergan, M., Bergan, J., & Alty, J. L. (1994). Minimising conceptual baggage: Making choices about metaphor. In G. Cockton, S. W. Draper, & G. R. S. Weir (Eds.), *People and Computers IX*, Cambridge, UK: Cambridge University Press.

Ballas, J. A., C. L., Heitmeyer, & Perez, M. A. (1992). Evaluating two aspects of direct manipulation in advanced cockpits. *Proceedings of CHI '92: Human Factors in Computing Systems*, 127-137, New York: ACM.

Benbasat, I., & Todd, P. (1993). An experimental investigation of interface design alternatives: Icon vs. text and direct manipulation vs. menus. *International Journal of Man-Machine Studies,* **38**, 369-402.

Boss, E., Huls, C., & Claassen W. (1994) EDWARD: full integration of language and action in a multimodal user interface. *International Journal of Human-Computer Studies*, **40**, 473-495.

Brennan, S. E. (1990). Conversation as direct manipulation. In B. Laurel (Ed.), *The Art of Human-Computer Interface Design*, 393-404, CA: Addison Wesley.

Buxton, B. (1993). HCI and the inadequacies of direct manipulation systems. *SIGCHI Bulletin* **25**, 21-22, New York: ACM

Canfield-Smith, D., Irby, C., Kimball, R., & Verplank, B. (1992). Designing the star user interface. *Byte Magazine*, April, 210-223.

Card, S. K. (1995). The Anti-Mac: Violating the macintosh human interface guidelines. *Proceedings of CHI '95: Human Factors in Computing Systems*, 183, New York: ACM

Card, S.K., Moran, T. & Newell, A. (1983). *The Psychology of Human-Computer Interaction.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Card, S. K., Robertson, G. G., & Mackinlay, J.D., (1991). The information visualizer: An information workspace. *Companion Proceedings of CHI '91: Human Factors in Computing Systems*, 181, New York: ACM.

Claassen, W., Bos, E., & Huls, C. (1990). *The Pooh way in human-computer interaction: Towards multimodal interfaces.* MMC Research Report No. 5. Nijmegen Institute for Cognition Research and Information Technology (NICI), Nijmegen, Netherlands.

Cohen, P. R. (1992). The role of natural language in a multimodal interface. *Proceedings of the 5th Annual Symposium on User Interface Software Technology*, 143-149. New York: ACM

Cohen, P. R., Dalrymple, M., Moran, D. B., Periera, F. C. N., Sullivan, J.W., Gargan, J.R., Schlossberg, J. L., & Tyler, S. W. (1989). Synergistic use of direct manipulation and

natural language. *Proceedings of CHI '89: Human Factors in Computing Systems*, 227-233, New York: ACM.
Desain, P. (1988). Direct manipulation and the design of user interfaces. *Journal for the Integrated Study of AI, Cognitive Science, and Applied Epistemology,* **5**, 225-246.
Eberleh, E., Korfmacher, W., & Streitz, N. (1992). Thinking of acting? Mental workload and subjective preferences for a command code and a manipulation interaction style. *International Journal of Human-Computer Interaction,* **4**, 105-122.
Eberts, R. E., & Bittianda, K. P. (1992). Preferred mental models for direct manipulation and command-based interfaces. *International Journal of Man-Machine Studies,* **38**, 769-785.
Egan, D. E., Remde, J. R., Landauer, J. K., Lochbaum, C. C., & Gomez, L. M. (1989). Behavioural evaluation and analysis of a hypertext browser. *Proceedings of CHI '89: Human Factors in Computing Systems*, 205-210, New York: ACM.
Egido, C., Patterson, J. (1988). Pictures and category labels as navigational aids for catalogue browsing. *Proceedings of CHI '88: Human Factors in Computing Systems*, 127-132, New York: ACM.
Familant, M. E., & Detweiler, M. C. (1993). Iconic reference: evolving perspectives and an organising framework. *International Journal of Man-Machine Studies,* **39**, 705-728.
Foley, J. D. (1987). Interfaces for advanced computing. *Scientific American,* October, 83-90.
Franzke, M. (1995) Turning research into practice: Characteristics of display-based interaction. *Proceedings of CHI '95: Human Factors in Computing Systems*, 421-428, New York: ACM.
Frese, M., Schulte-Gocking, H., & Altmann, A. (1987). Direkte manipulation vs. konventionelle interaktion. Cited in Ziegler and Fahnrich (1988).
Frohlich, D. M. (1991). The design space of interfaces. In L. Kielldahl (Ed.) *Multimedia-Principles, Systems and applications,* 53-69. Berlin: Springer-Berlag.
Frohlich, D. M. (1993). The history and future of direct manipulation. *Behaviour and Information Technology,* **12,** 315-329.
Frohlich, D. M., & Luff, P. (1990). Applying the technology of conversation to the technology for conversation. In P. Luff, G.N. Gilbert, & D. M. Frohlich (Eds.) *Computers and conversation analysis*, 187-220. London: Academic Press.
Gittens, D. (1986). Icon based human computer interaction. *International Journal of Man-Machine Studies,* **24**, 519-543.
Haddock, N. J. (1992). Multimodal database query. *Proceedings of COLING.* 1274-1278.
Halasz, F. & Moran, T. P. (1981). Analogy considered harmful. *Proceedings of CHI '81: Human Factors in Computing Systems,* 383-386. Gaithesburg, Maryland.
Harrison, B. L., Ishii, H., Vicente, K. J., & Buxton, W. A. S. (1995). Transparent layered user interfaces: An evaluation of a display design to enhance focussed and divided attention. *Proceedings of CHI '95: Human Factors in Computing Systems*, 317-324. New York, ACM.
Hayes, P.J., Reddy, D.R. (1983). Steps towards graceful interaction in spoken and written man-machine communication. *International Journal of Man-Machine Studies*, **19**, 231-284.
Huls, C. (1995) *Fitting language into the picture.* Published PhD thesis. Nijmegen Institute for Cognition and Information, Technical Report No. TR 95-01.
Hutchins, E. L. (1989). Metaphors for interface design. In M. M. Taylor, F. Neal, & D. G. Bouwhuis (Eds.) *The Structure of Multimodal Dialogue,* 11-28. Amsterdam: Elsevier Science.

Hutchins, E. L., Hollan, J. D., & Norman, D. A. (1986). Direct manipulation interfaces. In D. A. Norman & S.W. Draper (Eds.) *User centered system design,* 87-124. Hillside, NJ: Lawrence Erlbaum Associates, Inc.

Iseki, O. , & Shneiderman, B. (1986). Applying direct manipulation concepts: Direct Manipulation Disk Operating System (DMDOS). *ACM SIGSOFT Software Engineering Notes,* **11**, (April 1986), 22-26.

Jacob, R. J. K. (1991). The use of eye movements in human computer interaction techniques: What you look at is what you get. *Transactions on Information Systems,* **9**, 152-169. New York: ACM

Johnson, J. (1987). How faithfully should the electronic office simulate the real one? *SIGCHI Bulletin*, October, 21-25. New York: ACM

Johnson, J., Roberts,T. L., Verplank, W., Smith, D. C., Irby, C. H., Beard, M., & Mackay, K. (1989). The Xerox Star: A retrospective. *IEEE Computer*, September, 11-29.

Jones, W. P., & Dumais, S.T. (1986). The spatial metaphor for user interfaces: Experimental tests of reference by location versus name. *Transactions on Office Information Systems,* **4**, 42-63.

Jordan, P.W. (1992). Claims for direct manipulation interfaces investigated. *Industrial Management and Data Systems*, 3-6.

Kacmar, C. J. (1989). *An experimental comparison of text and icon menu item formats.* Working paper, Texas A&M University Computer Science Department.

Karat, J. (1987). Evaluating user interface complexity. *Proceedings of the Human Factors Society 31st Annual Meeting,* 566-570. Santa Monica, CA: Human Factors Society.

Kunkel, K., Bannert, M., & Fach, P. W. (1995). The influence of design decisions on the usability of direct manipulation user interfaces. *Behaviour and Information Technology,* **14**, 93-106.

Kurtenbach, G. P., Sellen, A. J., & Buxton, W.A.S. (1993). An empirical evaluation of some articulatory and cognitive aspects of making menus. *International Journal of Human-Computer Interaction,* **8**, 1-24.

Laurel, B. (1986). Interface as memesis. In D. A. Norman and S. Draper (Eds.) *User Centered System Design: New Perspectives on Human Computer Interaction.* New York: Lawrence Erlbaum Associates, Inc.

Laurel, B. (1990). Interface agents: Metaphors with character. In B. Laurel (Ed.) *The Art of Computer-Interface Design.,* 355-366. CA: Addison-Wesley.

Lee, J., & Zeevat, H. (1990). Intergrating natural language and graphics in dialogue. *Proceedings of INTERACT,* 479-484. Amsterdam: Elsevier Science.

Lokuge, I., & Ishizaki, S. (1995). Geospace: An interactive visualization system for exploring complex information spaces. *Proceedings of CHI '95: Human Factors in Computing Systems,* 409-414. New York: ACM.

Luff, P., Frohlich, D. M. (1991). Mixed initiative interaction. In T. Bach-Capon (Ed.) *Knowledge Based Systems and Legal Applications,* 265-294. London: Academic Press.

Macgregor, J. N. (1991). A comparison of the effects of icons and descriptors in videotex menu retrieval. *International Journal of Man-Machine Studies,* **37**, 767-777.

Maes, P. (1994). Agents that reduce work and information overload. *Communications of the ACM*, **37**, 31-40.

Morgan, K., Morris, R. L., & Gibbs, S. (1991). When does a mouse become a rat? or ... Comparing the performances and preferences in direct manipulation and command line environment. *Computer Journal,* **34**, 265-271.

Margono, S., & Shneiderman, B. (1987). A study of file manipulation by novices using commands vs. direct manipulation. *Proceedings of 26th Annual Technical Symposium of the Washington, D.C.Chapter of the ACM,* 57-62, Gathersburg, MD: National Bureau of Standards.

Myers, B. A. (1992). Demonstrational interfaces: A step beyond direct manipulation. *Computer*, August, 61-73.

Nass, C., Steuer, J., Henriksen, L., & Dryer, C. (1994). Machines, social attributions and ethopoeia: Performance assessments of computers subsequent to self and other evaluations. *International Journal of Human-Computer Studies,* **40**, 543-559.

Nass, C., Steuer,J., Tauber,E.R. (1994). Computers as social actors. *Proceedings of CHI '94: Human Factors in Computing Systems*, 72-78. New York: ACM.

Nickerson, R. S. (1991). Some characteristics of conversations. In B. Shackel (Ed.) *Man Computer Interaction: Human Factors Aspects of Computers and People,* 53-64. The Netherlands: Sijthoff and Noordhoff.

Neilson, I., & Lee, J. (1994). Conversations with graphics: Implications for the design of natural language/graphic interfaces. *International Journal of Human-Computer Studies,* **40**, 509-541.

Norman, D. A. (1986). Cognitive engineering. In D. A. Norman & S. W. Draper (Eds.) *User Centered System Design.* Hillside, NJ: Lawrence Erlbaum Associates, Inc.

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

Shneiderman, B. (1983). Direct manipulation: A step beyond programming languages, *IEEE Computer,* **16**, 57-69.

Shneiderman, B. (1992). *Designing the user interface: Strategies for effective human-computer interaction.* Addison-Wesley Publishing Co. (2nd Edition).

Shneiderman, B. (1994). Dynamic queries for visual information seeking. *IEEE Software*, (Nov.), **6**, 70-77.

Stearns, G., (1989). Agents and the HP NewWave application program interface. *Hewlett-Packard Journal,* **40**, 32-37.

Stenton, P. (1990). Designing cooperative interfaces: Tailoring the channel. In J. R. Galliers (ed.) *Proceedings of the First Belief Representation and Agent Architectures Workshop,* Computer Laboratory Technical Report No. 194, 193-197. University of Cambridge.

Strausfield, L. (1995). Financial viewpoints: Using point of view to enable understanding of information. *Companion Proceedings of CHI '95: Human Factors in Computing Systems*, 208-209. New York: ACM.

Streeter, L. (1995). Interface styles: Direct manipulation versus social interactions. *Companion Proceedings of CHI '95: Human Factors in Computing Systems*, 178. New York: ACM.

Tabachneck, H. J. M., & Simon, H. A. (1994). What you see is what you get - But do you get what you see? *Companion Proceedings of CHI '94: Human Factors in Computing Systems*, 293-294. New York, ACM.

Te'eni, D. (1990). Direct manipulation as a source of cognitive feedback: human-computer experiment with a judgment task. *International Journal of Man-Machine Studies,* **33**, 453-466.

Tullis, T. S., & Kodimer, M. L. (1992). A comparison of direct manipulation, selection and data entry techniques for reordering fields in a table. *Proceedings of the Human Factors Society 36th Annual Meeting,* **1**, 298-302. Santa Monica, CA: Human Factors Society.

Ulich, E., Rauterberg, M., Mott,T., Greutmann, T., & Strohmo, O. (1991). Task orientation and user-orientated dialogue design. *International Journal of Human Computer Interaction,* **3**, 117-144.

Venolia, D. (1993). Facile 3D direct manipulation. *Proceedings of INTERCHI '93,* 31-36. New York: ACM.

Walker, M. A., (1989). Natural language in a desktop environment. *Proceedings of 3rd International Conference on Human-Computer Interaction,* 502-509.

Whittaker, S. J. (1990). Next generation interfaces. *AAAI Spring Symposium on KBHCC,* 127-131. Stanford, CA.

Whittaker, S. J., Geelhoed, E., & Robinson, E. (1993). Shared workspaces: How do they work and when are they useful? *International Journal of Man-Machine Studies,* **39**, 813-842.

Whiteside, J., Jones, S., Levy, P., & Wixon, D. (1985). User performance with command, menu and iconic interfaces. *Proceedings of CHI '85: Human Factors in Computing Systems*, 185-191. New York: ACM.

Williamson, C., & Shneiderman, B. (1992) The Dynamic HomeFinder: Evaluating dynamic queries in a real-estate information exploration system. *Proceedings of ACM SIGIR '92*, 338-346.

Wixon D., & Good, M. (1987). Interface style and eclecticism: Moving beyond categorical approaches. *Proceedings of the Human Factors Society 31st Annual Meeting*, 571-575.

Wooldridge, M., Jennings, N. R., (1995). Intelligent agents: Theory and practice. *Knowledge Engineering Review*, in press.

Yankelovich, N., Levow, G. A., & Marx, M. (1995). Designing Speechacts: Issues in speech user interfaces. *Proceedings of CHI '95: Human Factors in Computing Systems*, 369-376. New York: ACM.

Ziegler, J. E., & Fahnrich, K. P. (1988). Direct manipulation. In M. Helander (Ed.) *Handbook of Human Computer Interaction,*123-135. Amsterdam: North Holland.

Zimmerman, T. G., Smith, J. R., Paradiso, J. A., Allport, D., & Gershenfeld, N. (1995). Applying electric field sensing to human-computer interfaces. *Proceedings of CHI '95: Human Factors in Computing Systems*, 280-287. New York: ACM.