HEWLETT
PACKARD

The History and Future of Direct Manipulation

David Frohlich
Information Systems Laboratory
HP Laboratories Bristol
HPL-93-47
June, 1993

direct manipulation, direct conversation, user interface, metaphor, modality

The earliest interactive computer systems were based on a conversational mode of interaction, in which user and computer communicated through the exchange of linguistic utterances. Since the advent of 'direct manipulation' technology, there has been a tendency to develop and promote an alternative mode of interaction, based on the user's manipulation of computer displayed objects. This paper reviews recent developments in the implementation and understanding of direct manipulation interfaces which point to various limitations in manipulative interaction which might be overcome through the selective re-introduction of conversational interaction. A re-conceptualization of the direct manipulation philosophy is suggested to accommodate these developments in which directness is said to be a property of both action and language based systems. A social definition of directness is proposed as the basis of a new philosophy of *graceful interaction* and a number of practical guidelines are offered to reduce the incidence of clumsy manipulation.

1. Introduction

One of the most significant developments of the 1980's in human computer interaction (HCI) was the emergence of 'direct manipulation' as a theoretical concept and design practice. At the heart of this development is the promotion of graphic and manual forms of interaction over and above more abstract and linguistic ones, on the grounds that the former place less load on the human cognitive system and are preferred by users. This philosophy has had a massive and largely beneficial impact on the face of personal computing and continues to exert a strong influence on the design of interactive software today.

However the technological and intellectual landscape of HCI in the '90s is changing and there is a need to review the direct manipulation 'story' in the light of these changes. Technologically, the supremacy of the personal computer as we know it may be coming to an end, with the advent of specialised and portable appliances for both computation and personal communication. Intellectually, there is a concern to understand the social as well as the cognitive attributes of computers as they are used within some real work setting. Given these changes, we might ask if is it still appropriate to hold on to the direct manipulation philosophy in its present form, and if not, what new philosophy should replace it.

The purpose of this paper is to answer this question. It is argued that direct manipulation was a concept for its time, and that a number of recent developments in theory and practice have begun to highlight new values for conversational forms of interaction which need to be integrated with the original insights. The paper begins by considering the origin and development of direct manipulation, and goes on to present a reconceptualisation of the concept prior to a discussion of implications for the field.

2. Original observations on direct manipulation

Direct manipulation was a term coined by Shneiderman (1982, 1983) to refer to a style of interaction characterised by the following three properties:

1. Continuous representation of the object of interest,
2. Physical actions or labelled button presses instead of complex syntax, and
3. Rapid incremental reversible operations whose impact on the object of interest is immediately visible.

These properties were said to be present in a variety of systems ranging from video games and spreadsheet packages to computer aided design systems and certain office systems.

Display editors were said to be a good example of these new-style interfaces, since they presented a full page of text in the format in which it was to be printed out (what you see is what you get), and allowed the user to manipulate that format directly by manually inserting space or text and moving blocks of material physically on the screen. In this way they presented the object of interest to users (i.e. the page as it would be printed) and supported rapid, reversible actions on that object through simple physical activity.

In effect Shneiderman was noting the intuitive appeal of physically manipulating graphical representations of objects instead of formulating complex linguistic commands, and he claimed a variety of benefits for such manipulation; including ease of learning, ease of use, retention of learning, reduction and ease of error correction, reduction of anxiety and greater system comprehension. A further observation was that the trick to creating such systems was to 'come up with an appropriate physical model of reality' from which to design a visual interface language (1982, p253).

At the time, the new systems referred to by Shneiderman contrasted stongly with more traditional database, statistics, transaction processing, and text editing systems based on command languages and character line displays. In these systems, users would interact with the computer by writing short programs, entering local instructions or responding to detailed questions, always in some 'programming language' of sorts. Hence the title of Shneiderman's 1983 article: *Direct manipulation: A step beyond programming languages*. Viewed in this context, manipulation-based interaction was indeed an exciting development, and it is easy to understand why Shneiderman concluded his 1983 article with such optimism:

> "The future of direct manipulation is promising. Tasks that could have been performed only with tedious command or programming languages may soon be accessible through lively, enjoyable interactive systems that reduce learning time, speed performance and increase satisfaction", (p68).

3. Theoretical developments since 1983

3.1 Introduction

Thinking about direct manipulation has itself developed and changed since 1983 through a series of papers discussing the relative merits of action and language based interactions. These are worth reviewing briefly here since they are unanimous in equating direct manipulation systems with the representation of model worlds, and in noting the value of alternative or additional forms of linguisitic interaction. They come primarily from Hutchins, Hollan & Norman (1986), Hutchins (1989) and Laurel (1990), and from two schools of researchers at Nijmegen University and Hewlett Packard Laboratories.

3.2 Hutchins, Hollan & Norman (1986)

In 1986 Hutchins, Hollan and Norman published an analysis of Shneiderman's original claims. They concurred with Shneiderman about the attractiveness of direct manipulation systems, but sought to explain the psychological basis of this attractiveness. Taking Shneiderman's definition of direct manipulation as a starting point, Hutchins et al went on to examine the notion of 'directness' in some detail, proposing that it was related both to the psychological *distance* between user goals and user actions at the interface, and to the psychological *engagement* of feeling oneself to be controlling the computer directly rather than through some hidden intermediary. The notion of distance relied heavily on an application of Norman's (1986) seven stage theory of action to the domain of human computer interaction. The notion of engagement was based on Laurel's (1986) ideas about the mimetic and metaphorical nature of interfaces.

Essentially distance refers to the mis-match between the way a user normally thinks about a problem domain and the way it is represented by a computer. Systems which reduce distance reduce this mis-match and the associated mental effort of working out what can be done in the system (semantic distance) and how to do it (articulatory distance). Engagement on the other hand refers to a particular style of representation based on a model world metaphor rather than on a conversational metaphor for interaction. Systems can encourage a feeling of engagement by depicting objects of interest graphically and allowing users to manipulate them physically rather than by instruction.

This leads to a space of interactive systems varying in distance and engagement (see Figure 1). Direct manipulation systems minimise distance and maximise engagement and are said to present the most 'direct' kind of interfaces to users. Interestingly, Hutchins et al also acknowledge the possibility of poorly designed model world systems (low level world) and well designed conversational systems (high level language) in this diagram, to indicate that direct manipulation is not a panacea for all interface design problems. More specifically they note the difficulty of carrying out repetitive operations, identifying sets of objects and specifying very precise values through manipulation (p118).

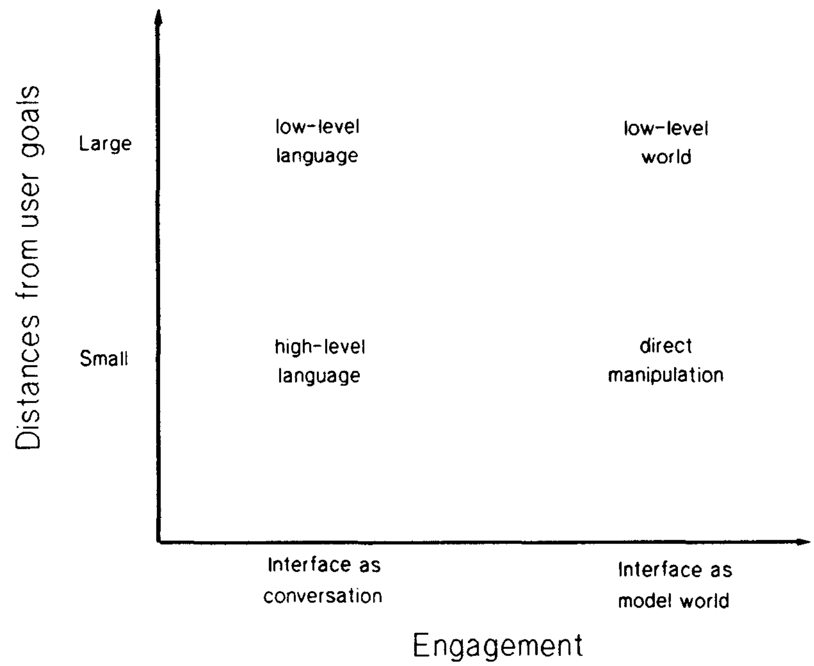

Figure 1. A space of interfaces. Figure 5.8 in Hutchins, Hollan & Norman (1986). Redrawn with permission.

Thus Hutchin's et al conclude that direct manipulation systems are likely to have benefits insofar as they simplify the mapping between goals and actions at the interface, but at the cost of losing the power of abstraction in interaction: "Basically, the systems will be good and powerful for some purposes, poor and weak for others", Hutchins et al (1986, p119).

3.3 Hutchins (1989)

Hutchins (1989) promotes mixed mode interaction in a discussion of *Metaphors for interface design*. He describes three kinds of metaphors commonly used in interface design; activity metaphors representing the broad goals of some computer based activity, mode of interaction metaphors representing the general relationship between the user and the computer, and task domain metaphors representing the structure of particular computer based tasks. Concentrating on mode of interaction metaphors, Hutchins goes on to describe four types which he calls the conversation, declaration, model-world and collaborative manipulation metaphors. Direct manipulation systems are said to be those which support the model-world metaphor where "expressions have the character of actions taken in the world of interest", (p20).

In a departure from the original theory, Hutchins argues that direct manipulation systems have the special psychological property of eliminating the reference gap or 'referential distance' between the expressions in an interface language and the objects to which they refer. It was first argued that direct manipulation systems helped to reduce the referential ('articulatory') distance between the form and meaning of an expression inherent in the conversation metaphor. Hutchins now claims that they serve to remove it altogether, by collapsing the linguistic description of an action on some object into an action which is itself applied to a representation of the object.

Paradoxically, the lack of referential distance within direct manipulation systems is shown by Hutchins to be a weakness as well as a strength. By definition, some referential distance is required to support abstract reference to unseen objects, which might itself be useful for particular tasks. The example Hutchins gives is wanting to perform some action on every word in his paper begining with the letter s. It might be easier to specify this action using a single conversational expression refering to the character of the target words, rather than by searching for each target word itself and performing the action manually.

This leads Hutchins to propose a collaborative manipulation metaphor for human computer interaction. This is characterised by the combination of conversation with an intelligent agent with the manipulation of objects in a model world. In Hutchin's own words, the collaborative manipulation metaphor suggests that:
"The user should be able to have a conversation about the world with the agent, and both the user and the agent should be able to manipulate the shared world", (p24).

3.4 Laurel (1990)

Quite independently, Laurel (1990) makes the same point in an article promoting *Interface agents: metaphors with character*. Laurel was commissioned by Apple Computers to edit a collection of readings on interface design intended to train Apple employees. Her own chapter highlights a tension in the rest of the book between those who argue for the further elaboration of the model world metaphor and those who argue for a re-introduction of the conversational metaphor. Interface agents are an implied focus for conversational interaction or action through an intermediary, being defined as "a character, enacted by the computer, who acts on behalf of the user in a virtual (computer-based) environment", (p356).

This represents a development of Laurel's 1986 paper which promoted 'first-person' interfaces, in which the user feels him or herself to be acting directly on some model world (rather than indirectly through some hidden intermediary). Like Hutchins, Laurel begins to note a number of computer-based tasks which are difficult to perform directly in first-person mode and might be more easily achieved in second-person mode 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.

These ideas on the relative advantages and disadvantages of manipulative and conversational forms of interaction can also be found elsewhere in the interface design literature. In particular, two distinct schools of interface researchers from Nijmegen University and Hewlett-Packard Laboratories provide various comparisons and critiques in support of mixed mode interfaces.

3.5 The Nijmegen school

Two papers from Nijmegen build on Norman's (1986) Seven Stage Theory of Action to evaluate conversational and model-world systems in terms of their 'directness' (Claassen, Bos & Huls 1990, Desain 1988).

Desain (1988) examines ways of minimising the semantic and syntactic distance between user goals and actions at an interface, in systems based on both the conversational and model-world metaphors. He shows how each mode of interaction can facilitate directness in different but complimentary ways. For example, conversational systems are said to be good for supporting abstract refering expressions and varying amounts of precision and detail. In contrast, model-world systems are said to be good for supporting anaphoric reference and a consistent user model.

Claassen, Bos & Huls (1990) collapse Hutchins' 4 mode-of-interaction metaphors into 2, the 'conversation mode' and the 'manipulation mode', before going on to summarise what they see as their main disadvantages. The conversation mode is said to be bad for helping the user maintain a mental model of the system, describing spatial structures, learning to use, and refering to entities. The manipulation mode is said to be bad for describing functional or causal properties, providing help, supporting abstractions, interrupting the user, and refering to actions and invisible objects.

Both these works conclude strongly that *engagement*, in the first-person sense, is not always necessary to ensure directness, since some tasks can be performed more 'directly' by instructing an intermediary in language than by manipulating objects in some model world.

3.5 The Hewlett-Packard school

This last point is also stressed in a number of papers from Hewlett-Packard research staff writing on mixed mode interaction (eg. Brennan 1990, Frohlich 1991, Jones 1991, Stenton 1990, Walker 1989, Whittaker 1990). All these authors point to the shortcomings of manipulation based interaction and promote the use of conversational interfaces to compensate for these in a mixed language/action environment.

For example, Whittaker (1990) lists three major problems with direct manipulation interfaces as defined by Shneiderman. Ensuring that objects of interest are continuously visible is said to be difficult or impossible in complex or distributed applications. Giving immediate feedback about the results of action is said to be impossible for actions with delayed results and undesirable for background tasks. Finally, making actions incremental is said to be tedious and inefficient for repetitive tasks.

Walker (1989) begins to list some of the properties of language which could overcome these problems in a desktop office environment. These include definite description, discourse reference, temporal specification, quantification, coordination, negation, comparison and sorting expressions.

These kinds of comparisons led Brenan (1990) to promote conversation as a form of direct manipulation, and Stenton (1990) to propose mixing the conversational and manipulative modes of interaction by design. Frohlich (1991) has begun to map out the design space of possibilities for such mixing, and to note that it is already taking place even in classic 'direct manipulation' systems such as the Apple Macintosh.

4. Major technological developments since 1983

4.1 Introduction

Since the advent of the earliest direct manipulation systems noted by Shneiderman there have been roughly three significant developments in technology (insofar as the direct manipulation philosophy is concerned). The first development is a continuation of the trend to design model world or *virtual world systems*. The second development is a counter-trend to extend the scope and usability of conversational or *virtual partner systems*. And the third development is a growing tendancy to combine both action and language based interaction in so called *mixed mode systems*. A brief flavour of each development is given in the following three sub-sections.

4.2 Virtual world systems

The ultimate extension of the direct manipulation philosophy can be seen in the development of virtual world systems. These are model-world systems so compelling as to cause the illusion of true 'immersion' inside the world. In other words users are psychologically projected into an artificial reality which they can affect through their own actions.

The sense of presense or immersion in the world is a critical factor which goes beyond mere engagement with objects of interest. This can be illustrated with reference to Computer Aided Design systems for architects. Conventional CAD systems might allow architects to draw and then manipulate building plans displayed on a graphics workstation. Virtual reality CAD systems would be able to present a stereoscopic display of the planned building for evaluation by allowing users to move around in it (see below). Only in the second case does the user experience being inside the scene depicted for manipulation.

Historically the development of virtual reality systems can be traced back to the work of Sutherland at MIT in the early 1960's. Sutherland developed the first interactive graphics system called Sketchpad around 1963 and followed this up with the first head mounted stereoscopic display in 1965. However, virtual reality system development accelerated rapidly from the early 1980's onwards with the establishment of military projects like the US Airforce Supercockpit programme in 1982 and the NASA VIEW Project in 1985. These projects recognised the value of simulating dangerous or inaccessible environments for operator training and/or remote control, and contributed to the drive for better input and output technologies for human computer interaction. Foley (1987) lists some of the new interface technologies to emerge during this period, which include 3-D interactive colour graphics, stereoscopic head mounted displays, and motion sensitive input devices such as the DataGlove.

The application of virtual reality systems now extends far beyond the military field. Brooks (1987) lists a variety of early civilian applications, including the WALKTHROUGH system which allows architects and their clients to move through a 3-D representation of a building generated from plan and elevation drawings, by manipulating a pair of joysticks, a 6-D sensor or by walking on a treadmill steered with handlebars. Other examples given are GROPE to support molecule docking for chemists, GRINCH to support electron density interpretation for biochemists, and ANATOMY RECONSTRUCTION to support the interpretation of anatomical representations for radiologists. In addition to Computer Aided Design, mollecular modelling and medical applications, virtual reality technology is also being applied within the leisure and entertainment industry to create realistic special effects for films and interactive games.

However, despite modern developments in virtual reality toolkits and cheaper input and output devices virtual reality technology has not yet achieved widespread commercialisation.

There are some who doubt that it ever will because of the implicit expectation that people will be prepared to transfer a large part of their everyday activities into the artificial environment of the computer. Rather they argue that computers should be used to invisibly enhance the real world environment that people are already familiar with operating (Weiser 1991). This is a nice twist to the virtual reality story which is sometimes referred to as *real virtuality* or *ubiquitous computing* because it is based on the notion that computers and computer-generated objects should become part of the furniture and fabric of everyday life. Thus Weiser (1991) goes on to describe work at Xerox PARC in which researchers are creating enhanced office environments containing wall sized displays in meeting rooms, pen-based scratchpads for viewing and marking electronic documents, miniature tabs which represent project folders, and active badges which keep track of your movements in the building.

Another example of this approach is represented by work on the DIGITAL DESK and MARCEL systems which monitor real paperwork with a camera suspended above a desk and project computer displayed images onto the same desk surface (Wellner 1991, Newman & Wellner 1992). In this way people can momentarily conjure up and manipulate virtual computer generated objects such as a calculator, or perform instant computations like foreign word translations by simply gesturing at textual material in the real world.

In the long term, the real virtuality movement seems set to make the bigger impact, since industry commentators are already noticing the way in which computation is creeping into ordinary households in the guise of enhancements to existing devices like phones, televisions and video cassette recorders (Miller 1990). However, from the point of view of developments in interface design, both virtual reality and real virtuality systems share the characteristic of promoting manipulation-based interaction with virtual material objects. This contrasts with systems described in the next section which promote language-based interaction with virtual immaterial subjects.

4.3 Virtual partner systems

In parallel to the development of virtual world systems, there has been substantial progress in the development of what might be called *virtual partner systems*. These are systems based on a conversational rather than a manipulative mode of interaction, which employ linguistic utterances rather than physical actions as the primary means of communication. Consequently they tend to project system facilities in terms of the knowledge of some virtual partner rather than as the behaviour of objects within some virtual world. Typical virtual partner systems are knowledge based or expert systems designed to answer user questions within some domain (e.g MYCIN), natural language information retrieval systems (e.g. INTELLECT) and speech input-output systems (e.g. VODIS).

Expert and knowledge based system technology has matured to the point where such systems are being routinely embedded in more conventional applications as diagnostic or advisory modules. Current research is therefore focussed on methods of integration and maintenance (e.g. Merry 1992), on improving the quality and helpfulness of advice and explanation (e.g. Gilbert 1989, Stenton 1987) and on making the dialogue more conversational (e.g. Frohlich & Luff 1990). An offshoot of this technology lies in the development of intelligent tutoring systems which are often designed to *ask* questions, critique actions and explain concepts (e.g. Cawsey 1989). To do this these systems usually monitor and converse with users about their actions or understandings of some domain. They may therefore be perceived as virtual partners working alongside the user in some wider activity which is also represented by the system (see Section 4.4).

Natural language technology has been steadily improving for some time, and currently allows the construction of reliable but restricted (text-based) natural language interfaces to knowledge bases, databases and other applications. These are becoming increasingly important in information retrieval applications since users can no longer be expected to navigate through the growing volume of information available to them on personal and corporate databases around the world. One thing which makes natural language interaction feasible today is the fact that users have been found to automatically adapt their style of language use to that of the computer (Brennan 1991).

Speech input and output has suddenly become highly valued with the growth of telecommunications services and the miniaturisation of devices. In the first case, the possibility of providing telephone enquiry services by computer is acting as a massive impetus to the development of better speech recognition and synthesis technologies (c.f. the SUNDIAL system, Peckham 1991). In the second case, the design of smaller and smaller devices is making voice and/or handwriting input into an attractive option. This is related to the shrinking of input devices and screens which makes it more difficult to design convincing graphical representations for manipulation based interaction and to the replacement of keyboards by pens. Whenever voice or handwritten input is interpreted by the device the user is likely to think of it as a virtual partner.

As with developments in virtual world systems there is a twist to developments of virtual partner systems. This is represented by the proliferation of text and video based systems for computer mediated communication. Electronic mail, text and videoconferencing systems are arguably virtual partner systems by virtue of introducing other human agents into the human computer interaction. However, many of these systems support interpersonal communication in the context of some shared workspace and might therefore be better described as mixed mode systems (see Section 4.4)

4.4 Mixed mode systems

Finally, in contrast to both the above trends there is evidence of an upsurge in the design of 'mixed mode' systems which attempt to integrate language and action modes of interaction in a single system (Taylor 1988, Taylor, Neel & Bouwhuis 1989).

Perhaps the simplest examples of mixed mode systems are those in which the language and action components of a system have been deliberately separated into distinct virtual partner or virtual world *metaphors*. By definition, these systems depict an explicit interface agent within some virtual world. This is true of AGENTS in Hewlett-Packard's NewWave office environment (Stearns 1989), GUIDES in an Apple Computers' hypermedia system (Oren, Salomon, Kreitman & Don 1990), and ANGELS in MCC's HITS Knowledge Editor (Terveen 1990). Agents are currently designed as software robots for automating routine activities performed on the NewWave desktop, and can be instructed by demonstration or through a specialised 'Agent Task Language'. Guides recommend routes through a hypermedia database on American History between 1800 and 1850, based on their likely interests as stock characters of the period. Angels provide advice on the use of a knowledge representation and maintenance tool, making suggestions and engaging the user in question answer sequences about editing operations.

The other class of systems which mix interface agents with virtual world representations are certain computer mediated communication systems (see also Section 4.3). In these cases the 'interface agents' are not agent-like bundles of system features but other people to whom the system connects the user in some way (e.g. the COCO system, Tang 1993). Much of the value of these systems lies in the fact that people can communicate remotely about some shared visual environment which is also represented by the system (Whittaker, Geelhoed & Robinson 1993), and there is a wealth of reported experimentation and experience in this field which begins to suggest ways of mixing language and action based forms of interaction to good effect (e.g. McCarthy & Miles 1990, Whittaker, Brennan & Clark 1991).

In other mixed mode systems, the separation between the conversational and manipulative components is less clear; being more closely integrated at a micro rather than a macro level of analysis. Thus for example, in most 'desktop' office systems Windows, Icons, Menus and a Pointer are used together to manipulate a virtual office world. This leads to a now common WIMP-style interface in which physical manipulations are performed through *sequences* of mixed mode expressions. For example in Hewlett-Packard's NewWave office environment one way of duplicating a document involves clicking over it's icon with a mouse-driven cursor, before applying 'Copy' and 'Paste' operations to it through a menu.

Although action might be said to be the *dominant* mode in such systems, they cannot technically be called direct manipulation systems because of the way in which linguistic commands on menus are invoked in combination with manipulative actions like pointing, dragging and dropping (Frohlich 1991). Desktop office systems might be contrasted with certain hypertext systems in which *language* is dominant but presented for manipulation within some visuo-spatial framework. Electronic form filling systems such as the FORMS HELPER are mixed mode in this second sense since they involve carrying out a conversation with a virtual partner within the framework of a stack of virtual pages (e.g. Frohlich 1987).

In yet other systems, language is offered as an *alternative* to action as a method of interaction, as in the ACORD system (Lee & Zeevat 1990). This 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.

This kind of approach is being extended by supporting various *combinations* of language and action in a single input expression. Thus in the ACORD system itself, pronominal reference in a sentence like 'It goes to the depot' can be performed by the user clicking on the 'it' in question. Other examples include the manipulation of graphical tree formalisms to change the focus of a natural language question about circuit boards in SHOPTALK (Cohen et al 1989), and the use of previous result windows to constrain the interpretation of subsequent natural language database queries in SAMi (Nelson & Stenton 1991).

Another way in which the language and action modalities can be combined is to split them across the input and output interfaces. This results in what might be called *cross modal* interfaces of either a Language-In/Action-Out or Action-In/Language-Out form (Frohlich 1991). An example of the former is the use of voice commands in the VIrtual Environment Workstation or VIEW system (Fisher, Wenzel, Coler & McGreevy 1988), while an example of the latter would be represented by a system summary or critique of user actions performed within an intelligent tutoring system.

All these forms of mixed mode interaction show how the character of pure virtual partner or virtual world systems can be altered by introducing elements of the complementary mode into the input or output interface. Not all of these forms lead to the kind of 'clean' collaborative manipulation described by Hutchins (1989) in which the user can both perform actions on some virtual world and delegate actions to an explicit interface agent having equal access to the same world. Many mixtures of language and action are possible, some of which appear to result in quite novel forms of physical and social activity which have no correspondence in the natural world. Perhaps the most challenging mixtures lie at the boundary between the physical and the social, where language is somehow manipulated as a material in its own right, as in Aron's (1991) hyperspeech system for presenting interview data, or where sequences of actions become welded together into higher level language-like units, as in visual programming by example (e.g. Kurlander & Feiner 1992).

5. Towards a new philosophy

5.1 Separating directness and manipulation

What then, should be made of these developments in the application and modification of ideas on direct manipulation interfaces? What implications do they have for the direct manipulation philosophy?

In the first instance they appear to point to a clear separation between issues of directness and issues of modality in interface design; in the simplest case to a separation between directness and manipulation.

Manipulation refers to an action-based mode of interaction between people and computers, in which, as Hutchins says, expressions at the interface take on the character of actions within a world of interest to the user. This action-based form of interaction encourages a certain quality of interaction which has been described as 'engagement', where the user feels him or herself to be the primary agent effecting changes in this world, rather than being 'once removed' as it were, in the position of an onlooker instructing another agent to carry out such changes on his or her behalf.

Being 'once removed' typifies a second person form of interaction which corresponds to an alternative conversational modality. Here expressions at the interface take on the character of utterances in a conversation (Frohlich 1991). This leads to a language-based form of interaction which exploits the instrumental and regulatory functions of language (Haliday 1970) in changing the world by controlling the behaviour of some virtual partner to which the input expressions are addressed. It is this partner who is seen, metaphorically speaking, to carry out the requested actions and return their results. The growing debate about interface agents seems to turn on the extent to which this virtual partner should be characterised explicitly at the interface.

Directness seems to have been a term originally intended by Shneiderman to refer to the first personness of interaction through manipulation. However, Hutchins, Hollan and Norman have now substantively changed the meaning of 'directness' by dividing it into distance and engagement; using engagement to refer to first personness. Distance is elaborated at considerable length to refer to the cognitive complexity of mapping psychological variables of interest to the user, such as goals, to physical variables of interest to the computer, such as actions at the interface. Systems that are designed in such a way as to minimise this distance are now described as direct. Given this new interpretation of directness as minimised distance, it is quite legitimate and indeed instructive, to apply it to both manipulative *and* conversational systems as Hutchins et al suggested (see again Figure 1) and as the Nijmegen and Hewlett-Packard schools have done.

What these schools have begun to discover is that all combinations of directness and modality are indeed possible. In contrast to Shneiderman's examples of indirect conversation ('low-level language' in Figure 1) and direct manipulation, researchers from each school have begun to provide examples of indirect manipulation ('low-level world' in Figure 1) and direct conversation ('high-level language' in Figure 1) - see also Cohen 1992. In fact, the entire debate about the relative advantages and disadvantages of language versus action based interfaces turns on an attempt to explicate the conditions under which each is most direct.

Fortunately there is a symmetry to these conditions which means that language is often well designed to overcome the problems with manipulation as a method of interaction, and vice versa. This complementarity of physical and social activity should not be seen as an accidental feature of system design but rather as a deep property of human collaboration and perhaps even of life itself (c.f. Pattee 1977). Its effect on technological developments in system design appears to have been to polarise the application of manipulative interaction around tasks which are naturally *visual* in character, and the application of conversational interaction around tasks which are not. Mixed mode interaction tends to be applied to support complex tasks in which only a subset of operations can be represented visually (see Figure 2).

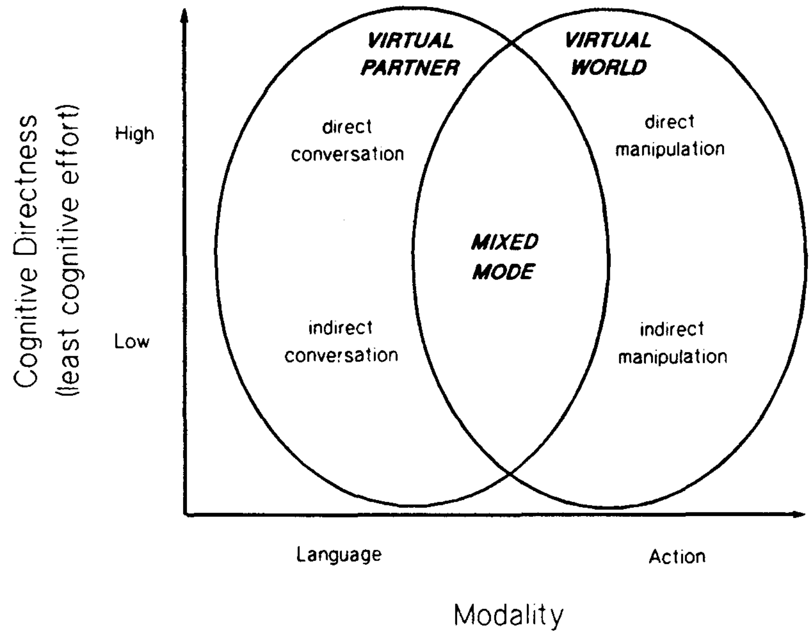

Figure 2. Direct and indirect interaction as a function of cognitive directness and interface modality.

Separating directness from manipulation in this way now allows us to talk more radically of what might be called *direct interaction*; where the feeling of doing things directly with a computer system can result both from the 'rapport' which can develop between a user and some virtual partner, as well as from the 'engagement' between a user and some virtual world. The concept of rapport between user and system is a productive one since it calls attention to some of the social dimensions of directness which have so far been overlooked in discussions of direct manipulation. These are explored in the next section to reveal a puzzle for any theory of direct interaction, which is that high levels of rapport in conversational interaction seem to be associated with high levels of *indirectness* between participants.

5.2 Rapport and indirectness in conversation

There appear to be three kinds of indirectness associated with conversational interaction.

The first type of indirectness simply refers to the high cognitive effort involved in describing certain concepts in language. This indirectness is present at least in the expression of visuo-spatial relations such as when you are explaining directions to someone over the phone. This particular task is subjectively hard work, and is performed more easily and efficiently in mixed mode fashion with reference to a shared visual workspace or map (Whittker et al 1992). Its relevance to HCI should be clear from the discussion above in which it was suggested that this kind of indirectness might be minimised by freely mixing and matching linguistic and manipulative modes of interaction to suit the task at hand.

A second type of indirectness in conversation seems to be associated with conventions of politeness where people may not always speak frankly and openly about certain issues because it would be considered rude to do so. These kinds of inhibitions and constraints appear to be relaxed with increases in rapport between parties who subsequently feel they can talk directly on a wide range of subjects. This kind of directness is probably irrelevant to the HCI situation in which there are few, if any, inhibitions about the kind of subjects which might be broached in interaction. Indeed such disinhibition is a valuable attribute of computer interviewing systems to which people are often more willing to divulge sensitive personal information (e.g. Waterton & Duffy 1984).

A third kind of indirectness in conversation stems from the fact that participants commonly manage to mean much more than they actually say. In doing this they rarely say directly what they mean. This phenomenon is worth explaining in a little more detail since I want to argue that it is seldom recognised but highly relevant to the HCI situation.

Garfinkel (1967) provides a dramatic demonstration of this kind of indirectness in two informal experiments he asked his students to perform. In the first experiment he asked students to report conversations they had had by writing on the left hand side of a sheet of paper what was actually said and on the right hand side what they and their partners understood they were talking about. Many more matters were understood than were spoken about as can be seen from the following two opening turns from one such report (pp38-39):

HUSBAND: Dana succeeded in putting This afternoon as I was bringing Dana, our

	a penny in a parking meter today without being picked up.	four-year-old son, home from the nursery school, he succeeded in reaching high enough to put a penny in a meter parking zone, whereas before he has always had to be picked up to reach that high.
WIFE:	Did you take him to the record store?	Since he put a penny in a meter that means that you stopped while he was with you. I know that you stopped at the record store either on the way to get him or on the way back. Was it on the way back, so that he was with you or did you stop there on the way to get him and somewhere else on the way back?

In the second experiment Garfinkel asked students to try get co-participants to spell out directly the sense of their remarks. This turned out to be both impossible and distressing since people tended to become aggressive in sanctioning the request, as shown below (p42):

(S) Hi, Ray. How is your girl friend feeling?

(E) What do you mean, "How is she feeling?" Do you mean physical or mental?

(S) I mean how is she feeling? What's the matter with you?

Garfinkel used these findings to show that the use and interpretation of language relies on a good deal of taken-for-granted knowledge possessed by speakers and hearers relating to the biography, relationship, and purposes of participants and to the form and placement of utterances in the conversation. Other versions of the indirect character of language can be found in philosophy and linguistics. For example Searle (1975) points out the form and importance of indirect speech acts in which the literal and intended meaning of an utterance appear to come apart, and Grice (1975) describes a cooperative principle for conversation whereby participants expect each other to be as honest, relevant, informative and brief as possible. This leads people to design economical contributions to conversation which, for practical purposes, both refer and add to the background knowledge or common ground held between them (Clark & Schaefer 1989, Clark & Brennan 1991). This can only work because of the highly interactive system of turn taking in conversation and the many sequential opportunities it provides for the repair of misunderstanding (Sacks, Schegloff & Jefferson 1974, Schegloff, Jefferson & Sacks 1977).

As far as rapport between parties is concerned, this appears to be associated with highly indirect conversations in which there are large gains in common ground for fewer turns, and where misunderstandings, trouble and repair is minimised between parties. Conversely, indications of a breakdown in rapport can be found in the extensive interactional work involved in coping with dissaffiliative situations like disagreements or corrections where one party must spell out a trouble they are having with another's talk. Utterances in these areas are typically, prefaced, delayed, hesitant, and mitigated (e.g. Pomerantz 1984).

All this is a puzzle for the original notion of direct interaction in HCI, since we have now described a pervasive form of *indirectness* in conversation which seems to underlie highly efficient interaction. Furthermore, its very efficiency appears to rely on *more* rather than less cognitive work being performed by the interactants.

5.3 Twists, turns and graceful interaction

The puzzle of desirable indirectness in interaction is solved if we shift to a more social definition of directness as interaction in which there is least *collaborative* effort expended to achieve a users' goals. Activities which are cognitively indirect can then be seen as socially direct in that they have the effect of minimising the *joint* work carried out by system and user entailed in achieving task success.

This is a much more promising basis on which to build a philosophy of interaction since it clarifies the relation of 'directness' to usability which was always an unknown and unreliable quantity in the direct manipulation story (Ankrah, Frohlich & Gilbert 1990). Collaborative effort, unlike cognitive effort, can be measured in terms of the actual interactional work involved in carrying out some activity. In practice this work may amount to the number and severity of *twists* and *turns* in the interaction. Usability may simply be proportional to the total number of twists and turns in all interactions experienced with a system.

I'd like to refer to this view and its associated prediction as a philosophy of graceful interaction (after Hayes & Reddy 1983) since it is centrally concerned with the conditions underlying elegant interaction per se, rather than with those related to the reduction of cognitive effort in human interactants. Gracefulness rather than directness captures the flavour of this shift in emphasis from the interactant to the interaction, and further serves to distinguish the specific social definition of directness promoted above from all the other definitions of directness mentioned in the review. Unlike Hayes & Reddy (1983) who used the term to refer to certain internal properties of a gracefully interacting system, I use gracefulness to refer to the dynamic property of a human computer interaction that would normally be recognised as 'pleasing or attractive, especially in form, movement or action' (Concise Oxford Dictionary). The opposite of graceful interaction is graceless or clumsy interaction in which there is an awkwardness of movement or shape (Concise Oxford Dictionary).

Both gracefulness and clumsiness are equally likely to be found in interactions with language and/or action based systems as shown in Figure 3. However I suggest that the selection and combination of modalities to achieve graceful interaction and avoid clumsy interaction will serve as a better heuristic for design than any number of attempts to maximise the cognitive directness of user activity. At the end of the day it may matter less that the user has a consistent mental model of the user interface and more that he or she can accomplish useful interactive work, together with the system, in a way which is somehow smooth, elegant and trouble-free. This may mean sacrificing some attractiveness in the visual appearance of a user interface for a more important attractiveness in the turn-by-turn operation of the computer system by its users.

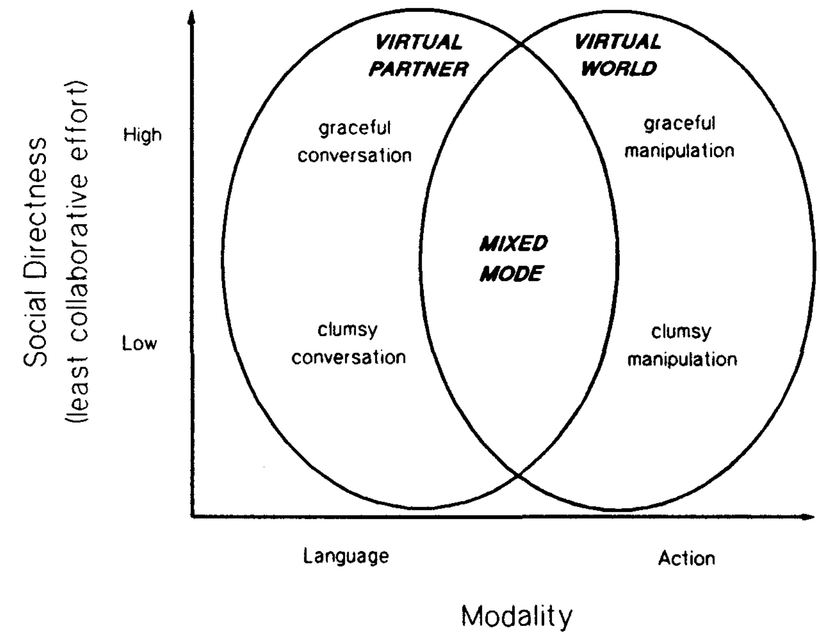

Figure 3. Graceful and clumsy interaction as a function of social directness and interface modality.

At the very least the philosophy of graceful interaction should re-focus design attention on a much neglected aspect of human computer interaction; namely the dynamics of the interaction itself (see also Frohlich, Drew & Monk 1993, Payne 1990). In order to understand what goes to make up a graceful interaction in contrast to a clumsy one it will be necessary to characterise interactional dynamics in a more technical and social sense than is generally attempted. Measuring twists and turns are two good starting points in this enterprise since they are non-trivial properties of an interaction which can only be identified by careful analysis of sequential interaction data. Thus a turn might be defined as a bounded 'contribution' to an interaction by one party preceeding a turn transition to the other party. A twist might be defined as a repair side sequence to the main interaction attending to some trouble detected earlier. Neither turns nor twists are straightforward to measure since their identification depends on making certain decisions about what counts as a true interactional 'boundary' and about when and where departures from the 'main' sequence begin and end (Frohlich, Drew & Monk 1993, Luff & Frohlich 1991). These are ultimately decisions about a user or system's *interpretation* of prior interactional events, situated in the ongoing flow of actions exchanged between them. To recognise that 'usability' is somehow related to the experience of interaction at this micro level of analysis is the first step in seeing and then fixing a clumsiness of interaction experienced at a more macro level.

6. The future of manipulation-based interaction

6.1 Selective application

So far I have reviewed the history of direct manipulation and made some recommendations about how its associated design philosophy might be updated. An important part of these recommendations is that the use of manipulation-based interaction should be further constrained rather than further expanded in the future, to take account of its limitations in supporting certain applications, technologies and tasks. Failure to constrain the use of manipulation in these contexts will result in forms of clumsy manipulation where the joint user and system work required to carry out individual operations is greater than it needs to be (see again Figure 3).

This doesn't mean that manipulation-based interaction has no future and should not continue to grow and develop within virtual world and mixed mode systems. Rather it implies that the real future for such technology lies in a more selective application within these systems, which builds on its strengths in helping users to visualise and explore novel virtual environments whilst avoiding its weaknesses in filtering and presenting non-visual information.

In this final section of the paper I want to draw together the main lessons of the review for constraining manipulation based interaction. There are really two classes of constraints. The first class are mode of interaction constraints related to the way a user is expected to think about the machine. The second class are task constraints related to the way a user is expected to carry out basic interaction tasks.

6.2 Mode of interaction constraints

Although it may be difficult to project a consistent mental model of any system to the user, the overall impression of how the system offers itself to the user is certainly a matter for consideration and design. In particular, the appropriateness of language or action metaphors for the application should be explicitly thought out, including options for mixed, multi and cross modal frameworks (see Section 4.4).

A pure implementation of the action modality should be avoided in situations where core information is typically non-visual, or where the computer is expected to take on an active rather than passive role in the interaction. In both instances the following forms of *agency* may be involved and better represented by a virtual partner than a virtual world:

1. Informing and reminding

In systems whose state may change as a result of external influences it is useful to be able to inform the user of changes. Alternatively users may want to ask the system to remind them of a future appointment or event.

2. Responding to interrogation

Users may want to learn about the current state or contents of a system without changing its state. In a material world there is no-one to ask and every attempt at interrogation causes a change in the world. An extension of this is the ability to make 'what if' interrogations to explore possible user actions and system states safely.

3. Helping and advising

In some circumstances users may benefit from spontaneous help in operating the system or from advice on how to improve their use of the system. In other cases the end value of the system may be in the provision of advice on some specialist topic and no further 'actions' may be required.

4. *Delegation and problem solving*

Cooperative problem solving might be achieved in some interactions by handing over responsibility for large parts of a task to the computer. Giving the system a high level task description for execution is an example of this, as is drawing on statistical reporting or rule-based diagnoses in data analysis.

6.3 Interaction task constraints

In one sense, the recent history of thinking on direct manipulation is a history of discovering tasks which are difficult to do by manipulation. Here is a brief list of them, for which linguistic or mixed mode expressions may be more effective:

1. *Referring to parts of the previous interaction*

Because everything in a model world is represented in the 'here and now', actions referring to the past or future are difficult to control. Referring to a previous segment of interaction is something you might want to do if you encounter trouble at the current 'turn' to see what went wrong and maybe selectively undo or redo prior activity. Another example of back-reference is in identifying a previous action or event for re-use.

2. *Scheduling actions to take place in the future*

Similarly referring to actions and events in the future is difficult through manipulation. This is necessary when controlling actions with delayed consequences or specifying actions to take place later or repeatedly. A way of describing time *intervals*, external event *triggers* and action *policies* seems to be required for this.

3. *Identifying unseen objects*

One the greatest drawbacks of manipulating what you can see as an interaction method is that objects of interest may lie out of view. This is really a version of the storage and retrieval problem which is so troublesome in the physical world. Getting unseen objects into view in order to interact with them may be tedious particularly in large applications with several layers of visual space, with many objects or with small viewing windows. Furthermore in applications where the object contents are unknown it may not even be worth looking. What is required here is a way of describing unseen objects by abstract properties such as their name, contents, or origin.

4. Identifying groups of objects

A related difficulty is in grouping objects which may be scattered spatially in the material world but nevertheless related in some way. Again the ability to describe categories of objects by property combinations would obviate the need to do extensive manual searching and visual comparisons.

5. Performing repetitive actions

If the representation of objects of interest in a model world system is too low level, users may find themselves performing similar manual actions over and over again. For example making the same spelling correction throughout a text document is difficult if done purely by manually searching and changing instances of the miss-spelt word. The ability to define and invoke 'macro' actions would be useful here to let the system handle some of this repetition.

6. Doing more than one thing at a time

The immediacy of feedback from manipulative user actions means that actions are generally performed one at a time. This reduces the possibility of concatonating instructions to execute several actions which are then carried out at the same time. Users might find this more efficient in cases where a string of operations have to be performed on the same object (e.g. forwarding a message to several people and printing it out.

7. Specifying actions very precisely

Physically manipulating objects to very high degrees of precision is difficult, for example in drawing an image to exact specifications. The ability to input numerical values is useful here.

6.4 Conclusions

Direct manipulation was clearly a large and important step beyond programming languages (Shneiderman 1983). However, it is not a panacea for human computer interaction as even the earliest of commentators point out (Hutchins, Hollan & Norman 1986) since there are situations in which manipulation is clumsy to perform (see Sections 6.2 and 6.3 above). The key to taking the next step beyond manipulation lies in recognising that these situations are the very same ones in which language is more graceful as an interaction method, and that a one-sided 'programming' of computers by people can be replaced by a two-sided conversation conducted in the context of some visual framework. Given the very many ways in which language and action might be combined at the interface, the options for new forms of manipulation based interaction are expanding rather than contracting at present; although the overall effect of exercising these options may be a general reduction in the dominance of manipulation as an interaction method.

Acknowledgements

I thank numerous colleagues and friends for their inputs to this paper, through critical reading and extensive discussion of earlier drafts. These include Phil Stenton, Steve Whittaker, Nick Haddock, Andy Hunter, Nigel Derrett and Bill Sharpe from Hewlett-Packard Research Laboratories Bristol, Andrew Monk, Greg Abowd and Alan Dix from the University of York, and Edwin Bos from the University of Nijmegen. I would also like to thank Steve Payne from the University of Cardiff and two anonymous reviewers for comments which prompted me to think more deeply about the (in)directness of conversational interaction in revising the paper.

Support

This work was carried out at Hewlett-Packard Research Laboratories Bristol, UK, and in the Departments of Sociology and Psychology at the University of York, UK. The latter placement was funded by a Cognitive Science/HCI Initiative of the joint British Research Councils through Special Training Fellowship number G106/511.

References

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 (Amsterdam: North Holland).

Arons, B. 1991. Hyperspeech: Navigating in speech-only hypermedia. *Proceedings of Hypertext '91*, 133-146.

Brooks, F.P.Jr. 1987. Grasping reality through illusion: Interactive graphics serving science. *Proceedings of the CHI '88 Conference on Human Factors in Computing Systems*, 1-11 (New York: ACM).

Brennan, S.E. 1990. Conversation as Direct manipulation: An iconoclastic view. In B. Laurel (Ed.) *The art of human computer interface design* (pp. 393-404) (Menlo Park, California : Addison-Wesley).

Cawsey, A. 1989. Explanatory dialogues. *Interacting with Computers 1*, 69-92.

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

Clark, H.H. & Brennan, S.E. 1991. Grounding in communication. In L.B. Resnick, J. Levine & S.D. Behrend (Eds.) *Perspectives on socially shared cognition* (Washington D.C.: American Psychological Association).

Clark, H.H. & Schaefer, E.F. 1989. Contributing to discourse. *Cognitive Science 13*, 259-294.

Cohen, P.R. 1992. The role of natural language in a multimodal interface. *Proceedings of the 5th Annual Symposium on User Interface Software and Technology (UIST '92)*, 143-149 New York: ACM).

Cohen, P.R., Dalrymple, M., Moran, D.B., Pereira, F.N.C., Sullivan, J.W., Gargan, Jr.R.A., Schlossberg, J.L., & Tyler, S.W. 1989. Synergistic use of direct manipulation and natural language. *Proceedings of the CHI '89 Conference on 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.

Fisher, S., Wenzel, E.M., Coler, C. & McGreevy, M. 1988. Virtual Interface Environment Workstations. *Proceedings of the Human Factors Society 32nd Annual Meeting.*

Foley, J.D. 1987. Interfaces for advanced computing. *Scientific American* (Oct.). 83-90.

Frohlich, D.M. 1987. On the re-organisation of form-filling behaviour in an electronic medium. *Information Design Journal 5*, 111-128.

Frohlich, D.M. 1991. The design space of interfaces. In L. Kjelldahl (Ed.) *Multimedia - Principles, systems and applications* (Berlin: Springer-Verlag).

Frohlich, D.M., Drew, P., & Monk A. (1993) The management of repair in human computer interaction. Under review.

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* (London: Academic Press).

Garfinkel, H. 1967. *Studies in ethnomethodology* (New Jersey: Prentice Hall).

Gilbert, G.N. (1989) Explanation and dialogue. *Knowledge Engineering Review, 4*, 235-247.

Grice, H.P. 1975. Logic and conversation. In P. Cole & J. Morgan (Eds.) *Syntax and semantics 3: Speech Acts*, 41-58.

Halliday M.A.K. 1970. Language structure and language function. In J. Lyons (Ed.) *New horizons in linguistics* (Harmondsworth: Penguin Books).

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.

Hutchins, E.L. 1989. Metaphors for interface design. In M.M. Taylor, F. Neel & D.G Bouwhuis (Eds.) *The structure of multimodal dialogue* (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*, (pp. 87-124) (Hillsdale, New Jersey: Lawrence Erlbaum Associates).

Kurlander, D. & Feiner, S. 1992. A history-based macro by example system. *Proceedings of the 5th Annual Symposium on User Interface Software and Technology (UIST '92)*, 99-106 (New York: ACM).

Laurel, B. 1986. Interface as memesis. In D.A. Norman & S.W. Draper (Eds.) *User centered system design*, (pp. 67-86) (Hillsdale, New Jersey: Lawrence Erlbaum Associates).

Laurel, B. 1990. Interface Agents: metaphors with character. In B. Laurel (Ed.) *The art of human computer interface design*, 355-366 (Menlo Park, California: Addison-Wesley).

Lee, J., & Zeevat, H. 1990. Integrating natural language and graphics in dialogue. *Proceedings of INTERACT '90*, 479-484 (Amsterdam: Elsevier Science).

Luff, P., & Frohlich, D.M. 1991. Mixed initiative interaction. In T. Bench-Capon (Ed.) *Knowledge based systems and legal applications*, (pp265-294) (London: Academic Press).

McCarthy, J.C. & Miles, V.C. 1990. Elaborating communication channels in conferencer. *Proceedings of the IFIP Working Group 8.4 Conference on Multi-user Interfaces and Applications*.

Merry, M. 1992. Rubicon - A computer configuration system. *Proceedings of Expert Systems '92* (London: IEE & BCS-ESSIG).

Miller, M.W. 1990. PC technology is sneaking into the home disguised as household gadgets. *Wall Street Journal*, 10th July.

Nelson, A.L. & Stenton, S.P. 1992. Dialogue modelling for information access. *ASLIB Proceedings 44*, 275-281.

Newman, W. & Wellner, P. 1992. A desk supporting computer-based interaction with paper documents. *Proceedings of the CHI '92 Conference on Human Factors in Computing Systems*, 587-592 (New York: ACM).

Norman, D.A. 1986. Cognitive engineering. In D.A Norman & S.W. Draper (Eds.) *User centered system design* (Hillsdale, New Jersey: Lawrence Erlbaum Associates).

Oren, T., Salomon, G., Kreitman, K. & Don, A. 1990. Guides: Characterising the interface. In B. Laurel (Ed.) *The art of human computer interface design*, (pp. 367-382) (Menlo Park, California : Addison-Wesley).

Pattee, H.H. 1977. Dynamic and linguistic modes of complex systems. *International Journal of General Systems 3*, 259-266.

Payne, S.J. 1990. Looking HCI in the I. *Proceedings of INTERACT '90*, 185-191. (Amsterdam: Elsevier Science).

Peckham, J. 1991. Speech understanding and dialogue over the telephone: An overview of the ESPRIT SUNDIAL project. *Proceedings of the DARPA Workshop on Speech and Language*, 14-27 (Los Altos, CA: Morgan Kaufmann).

Pomerantz, A. 1984. Agreeing and disagreeing with assessments: Some features of preferred/dispreferred turn shapes. In J.M. Atkinson & J. Heritage (Eds.) *Structures of social action: Studies in conversation analysis* (pp57-101) (Cambridge: Cambrige University Press).

Sacks, H., Schegloff, E.A. & Jefferson, G. 1974. A simplest systematics for the organisation of turn-taking for conversation. *Language 50*, 696-735.

Schegloff, E.A., Jefferson, G. & Sacks, H. 1977. The preferences for self-correction in the organisation of repair in conversation. *Language 53*, 361-382.

Searle, J.R. 1975. Indirect speech acts. In P. Cole & J. Morgan (Eds.) *Syntax and Semantics*

3: Speech Acts, 59-82.

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

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

Stearns, G. (1989). Agents and the HP NewWave Application Program Interface. *Hewlett-Packard Journal, 40*, 32-37.

Stenton, S. P. (1987). Dialogue management for co-operative knowledge based systems. *Knowledge Engineering Review, 2*, 99-122.

Stenton, S.P. (1990). Designing cooperative interfaces: Tailoring the channel. In J.R. Galliers (Ed.) *Proceedings of the first belief representation and agent architectures workshop, March 1990.* (Computer Laboratory Technical Report No. 194, 193-197). University of Cambridge.

Tang, J. 1993. A study of the use of a desktop conferencing technology to support a distributed team. Under review.

Taylor, M.M. (1988). Editorial: Special issue on multimodal computer-human interaction. *International Journal of Man-Machine Studies, 28*, 99-100.

Taylor, M.M., Neel, F., & Bouwhuis, D.G. (Eds.) (1989). *The structure of multimodal dialogue* (Amsterdam: North-Holland).

Terveen, L.G. (1990) Resources and conventions for person-computer interaction. *Proceedings of the AAAI Workshop on Complex Systems, Ethnomethodology and Interaction Analysis*, (pp81-100), Boston, July 1990

Walker, M.A. (1989). Natural language in a desk-top environment. *Proceedings of HCI89 3rd International Conference on Human-Computer Interaction*, 502-509. Boston, Massachusetts: ?

Waterton, J.J. & Duffy, J.C. 1984. A comparison of computer interviewing techniques and traditional methods in the collection of self-report alcohol consumption data in a field study. *International Statistical Review 52*, 173-182.

Weiser, M. 1991. The computer for the 21st century. *Scientific American, Sept.*, 66-75.

Wellner, P. 1991. The digital calculator: Tangible manipulation on a desk top display. *Proceedings of the 5th Annual Symposium on User Interface Software and Technology (UIST '91).*

Whittaker, S. (1990). Next generation interfaces. *Proceedings of the AAAI Spring Symposium on Knowledge Based Human Computer Communication,* 127-131.

Whittaker, S.J., Brennan, S.E., & Clark, H.H., (1991). Coordinating activity: An analysis of interaction in computer-supported cooperative work. *Proceedings of CHI '91* (New York: ACM).

Whittaker, S.J., Geelhoed, E. & Robinson, E. 1993. Shared workspaces: How do they work and when are they useful? Under review.