

An Introduction to Human Computer Interaction

Mike Sharples

School of Cognitive and Computing Sciences

University of Sussex

mike@cogs.susx.ac.uk

1. Interacting with Computers

1.1. Background to HCI

Human-computer interaction (HCI) is the study of the ways people interact with and through computers. It grew out of work on human factors (the U.S. term) or ergonomics (the European term) concerned with the design and usability of equipment. The aim of the early work in HCI was to provide interfaces to computer programs which hid the low-level operations of the computer and allowed the user to work through direct interaction with familiar entities such as words and diagrams, rather than by writing detailed instructions in a computer language. More recently the scope of HCI has broadened to include interface design, formal models of users and their interactions with computers, evaluation of computer systems, computer-mediated human to human communication and computer-supported cooperative work.

HCI has been greatly influenced by a few visionaries who imagined new ways of working with computers. They should not be seen as isolated prophets. The early work was well known to later researchers and the passionate enthusiasm of people such as Doug Englebart and Alan Kay has both inspired and determined the present-day methods of interaction with computers. Baecker and Buxton (Baecker & Buxton, 1987) provide a valuable historical survey of HCI.

Vannevar Bush and the MEMEX

In 1945 Vannevar Bush, a scientific advisor to the United States Government, published an article in the *Atlantic Monthly* which argued that a successful peacetime research effort would depend on people having ready access to appropriate information. He proposed a future device called the MEMEX which would extend the human memory by presenting and associating diverse pieces of information. An operator could call up text or images stored on microfilm and choose to follow trails of association to other related information, or could add new trails to indicate a train of thought or a cluster of ideas. The MEMEX was never implemented, but the ideas, of using a machine to augment the intellect, of direct manipulation of information, and of trails of association, laid the foundations of multimedia computing.

Ivan Sutherland and Sketchpad

By the early 1960's experimental time-sharing computers were being built which allowed a computer to work on several jobs simultaneously, sharing its time between users according to their activity. Computer time could be switched automatically between the users, allowing people to work at the computer screen in bursts of activity and to solve problems by

interacting with the computer rather than providing pre-prepared programs. Early work on military computer systems had shown that it was possible to display and manipulate images on a computer console, and researchers began to explore the possibilities of graphical interaction between humans and computers.

Ivan Sutherland, a researcher at MIT Lincoln Laboratory, implemented a drawing system named Sketchpad which both demonstrated the power of pictorial interaction and introduced methods which are still being developed for contemporary graphics packages. These included a *hierarchical representation* of pictures composed of sub-pictures, the notion of a *master picture* and transformed *instances* of the picture, and the use of *constraints* to specify attributes and relations (such as a two images being constrained to stay a fixed distance apart when one or other is moved).

Doug Englebart and Augment

Doug Englebart, as a graduate student at the University of California in the 1950's, advocated the computer as an 'augmentation device', offering people new ways to study problems, experiment with ideas and hunches, and test possible solutions. Computer programs, instead of being designed to solve a single problem could be constructed as toolkits, with parts that could be re-used and extended. The synergy which comes from combining these tools into an integrated 'workshop' makes each tool considerably more valuable than if it were used alone, and the combined effort of people working together on the computer augments the abilities of individuals into computer-assisted communities.

With colleagues at Stanford Research Institute in the 1960's Englebart developed NLS (oNLine System, later called NLS/Augment) which assisted people in working together on tasks such as planning, analysis and problem solving. It provided many novel facilities, such as active links between pieces of information, user-controlled windows, filters which displayed files with a specified content, the ability to send electronic mail, and shared-screen conferencing.

Alan Kay and the Dynabook

In the late 1960s Alan Kay wrote about his imagined self-contained knowledge manipulator in a portable package the size and shape of an ordinary notebook. It would be able to store and display text, drawings, music, and animations and the owner would be able to design documents, compose music, create pictures and communicate directly with other people through sound and image. Alan Kay named this notebook computer The Dynabook and in 1972 he suggested that "The Dynabook is now within reach of current technology". As a member of the Xerox Palo Alto Research Center, Alan Kay was uniquely placed to realise his vision.

For ten years he and colleagues developed a series of personal computers (which they called 'interim Dynabooks'), and a computer language called Smalltalk to support the construction and manipulation of dynamic objects. In April 1981 Xerox announced the 8010 Star Information system, a personal computer designed for office use. Although the Star was a desktop machine rather than a portable package, it realised much of the Dynabook vision. The Star was the first commercial computer to be designed around a *graphical user interface*, which offered a consistent analogy of the top of an office desk with surrounding furniture and equipment. The objects on the simulated desktop were represented by pictorial *icons* indicating 'in' and 'out' baskets for electronic mail, file drawers, folders and other

electronic metaphors for objects in an office. Operations on the Star were carried out by *direct manipulation* of objects. For example, to file a document the user would move an icon representing the document over a picture of a file drawer.

The novel aspects of HCI found in the Star computer, such as ‘the desktop metaphor’, ‘windows’, ‘icons’ and ‘pull down menus’, were brought to the mass market in the Apple Macintosh computer. As well as providing a consistent ‘look and feel’ to its range of computers, Apple set down guidelines for designers of Macintosh software which ensured that the Apple’s investment in interface design and human-computer interaction would be reflected in software produced by other companies. By the mid-1980’s companies such as Apple, Microsoft and Aldus were developing software for small business and home use which allowed people with no knowledge of computing, and little training in use of the software, to perform complex and highly interactive tasks such as producing spreadsheets and designing page layouts. The commitment of these companies to good practice in human-computer interaction, and their subsequent commercial success, has meant that research in HCI has an increasingly strong influence on commercial practice.

1.2. Interfaces

The interface to a computer is the combination of screen and interaction devices (such as a keyboard, mouse and light-pen) which enable a user to interact with the computer. A main aim of interface design is to produce interfaces which hide the complexity of the computer, giving a user the impression of working directly on a productive task such as writing a document or creating an illustration. One way to do this is by presenting the interface as an *analogy* to some more familiar system, such as a typewriter or set of illustrator’s tools.

The advantage of an interface analogy is that it provides a short-cut to learning a complex system, by offering concepts and operations that fit with the user’s existing experience (such as ‘cutting’ and ‘pasting’ objects on the screen). Although analogies and metaphors play a central part in interface design, they can lead to faulty reasoning, when the interface does not match the familiar system. (In most interfaces ‘cutting’ and ‘pasting’ an object does not exactly match the physical process. When a screen object is ‘cut’, it ‘disappears’ until the paste operation is performed.) They can also be restrictive, tying the interface to concepts and operations that occur in the non-computer world, rather than exploiting the possibilities of computer technology.

1.3. Interaction devices

As computers are adapted for diverse tasks, from control of industrial processes to art and design, they are being fitted with *interaction devices* which enable objects on the screen to be created, selected, oriented, tracked, and moved. New input devices include trackballs, DataGloves, eyetrackers, digitising tablets and thumb-wheel. The keyboard is still the main method of providing data to a computer, but the design of keyboards is changing to address concerns about fatigue and ‘repetitive strain injury’. New keyboards include wrist supports and keyboards which split down the centre, so that each half can be swung outwards to keep the hand in a comfortable position.

Direct pointing devices such as a light-pen, stylus, or finger on touchscreen can be moved directly to a point on the screen and can perform all the main interaction tasks, but writing on an upright screen can cause severe fatigue. A screen mounted at 30 to 45 degrees from the horizontal is convenient both for pen input and for reading. Indirect pointing devices such as

a mouse or trackball overcome the problem of writing on a vertical screen, but require hand-eye coordination and take up additional desk space. Coordination is not a major problem, however, and a child can learn to use a mouse in a matter of minutes.

There is much debate about the merits of different pointing devices, and factors affecting their use include the time required to select a near or distant target on the screen, ability to carry out fine movement for handwriting or drawing, muscular strain, and durability. Shneiderman (Shneiderman, 1992) proposes a touchscreen for durability in public-access applications, a mouse or trackball for accurate pointing, and a keyboard cursor when there are a small number of targets. Joysticks offer a firm grip and easy movement but are slow and inaccurate in guiding a cursor to a fixed destination.

A major limitation of all the devices described above is that they are intended for moving around a 2-dimensional space on the computer screen. With the advent of 3-dimensional simulations, shown on a computer screen or through helmet-mounted displays, there is a need for input devices with six degrees of movement (movement through three dimensional space plus forward and sideways tilt and rotation). Such devices include the DataGlove, which fit over the hand and can register its position and gestures), the Polhemus tracker (a wand whose position and orientation in space is transmitted to the computer) and the spaceball (a small ball, mounted on a base, which can be twisted and pushed).

1.4. Communicating with a computer

An influential account of human-computer communication is Norman's *execution-evaluation cycle* (Norman, 1986). The user of a system starts with a goal to perform some activity, forms an intention, and then specifies and performs some action. As the system responds, the user perceives the state of the system, interprets that state and evaluates it with respect to the goals and intentions. This leads the user to set further goals and continue the interaction. The user's goals are expressed in psychological terms (such as 'check the spelling of this word') but the system presents its current state in physical terms (such as a list of possible word corrections). The goals and the system state differ in form and content, creating 'gulfs' that need to be bridged to ensure successful communication.

The *gulf of execution* represents the gap between a user's intentions and the means by which the can be carried out on the computer. The user must specify an appropriate sequence of actions and then perform them by physical acts (such as selecting an item from a menu) that the computer can interpret. To bridge the *gulf of evaluation* a user must compare the state of the system, as presented on the computer screen or other output devices, with the original goals and intentions. If the computer does not appear to have satisfied the goal then the user must reformulate it (and possibly first attempt to undo the action which led to the wrong response).

A system designer can narrow the gulfs of execution and evaluation by such means as providing an interaction language which matches the user's intentions (for instance by providing a spell checker, which can be applied to selected pieces of text, and which can be operated by a memorable and easily accessible physical act), by supplying interaction devices that allow a user to translate intentions into actions, and by presenting the system in a form which can be easily interpreted in terms of psychological goals. If the user's psychological attributes map easily onto to the input language then the specification of tasks is made easier. Norman gives an example of poor mappings in everyday objects, such as a room with a panel of light switches where there is no simple relationship between the

position of the switches and the position of the lights, so that it is necessary to discover by trial and error which switch controls which light. For an interface controlling a complex industrial process the difficulties, and dangers, can be considerably greater.

1.5. Styles of communication

The cycle of activity and response creates a dialogue between the human user and the computer which has some similarity with human to human conversation. Successful communication depends on establishing a shared representation and carrying out a dialogue to progress a task.

Command line interaction

The earliest interactive computers communicated via commands typed at a console, and responded with printed output. Command line interaction is still found on powerful environments such as the UNIX operating system because it gives the user direct access to computer operations. These can be strung together into ‘scripts’ to carry out multiple tasks, with the output from one operation being used as input to the next. The penalty for this flexibility and power is that the computer offers little assistance as a conversational partner. The user must remember the commands and the syntax required to combine them into scripts. Command line interaction may be suitable for experienced ‘power users’ of a general-purpose computing environment, but there are many tasks, such as drawing, where issuing commands to the computer is tedious and unintuitive.

Menus and WIMP interfaces

Menus package together commands into a list from which one or more can be selected. Styles of menu include ‘pull down’ (where selecting a header word causes a menu to appear below it), ‘pop up’ where the menu appears beside an item selected on the screen, ‘walking’ (where selecting a menu item can cause a sub-menu to appear beside it), and ‘pie’ (where the menu radiates out from a central point). Menus overcome some of learning and memory problems of command line interfaces, but do not offer the power of combining commands into program-like scripts.

Menus form a part of WIMP (*windows, icons, menus, pointers*) interfaces. The screen displays a number of overlapping, bounded windows and a pointing device such as a mouse directs a cursor to one of the windows. Menus allow commands to be directed to the selected window and icons can represent commands or objects (such as ‘closed’ windows) in pictorial form. Other elements, such as ‘buttons’, ‘palettes’ and ‘dialogue boxes’ allow the user to communicate in a variety of modes.

Natural language

Some tasks, such as dictating a memo, would clearly be made easier by speaking to the computer and having the spoken word translated directly into text. Speech input may also be valuable for public-access systems, where there is limited, information-seeking communication, such as timetable enquiries or tourist information. Early natural language interfaces were limited to single word input, and had to be trained to recognise the speaker’s voice and intonation. More recent systems can recognise continuous, slowly-spoken speech and require less or no training. Natural language offers new forms of interacting with machines, such as by telephone conversation, but the apparent ease of spoken conversation hides the difficulty of conducting a useful dialogue. A computer application is designed to

perform a restricted task, such as giving tourist information, and cannot behave like a human conversational partner. Either the user must adapt to the limited linguistic capabilities of the computer (by guessing what language forms the computer might recognise and rephrasing commands until they are accepted), or the computer must direct the dialogue leaving the user to give limited responses such as ‘yes’, ‘no’ or numerals.

Pen input

Software for the automatic recognition of handwriting has encouraged the development of pen interfaces, which allow informal communication with the computer through writing, sketching and gesturing. A form displayed on the screen can be filled out by writing responses using a light-pen or stylus which are then converted into digits or words. A hand-drawn diagram can be tidied by converting roughly drawn shapes into exact straight lines, circles and boxes. A gesture, such as moving the stylus back and forwards over an object can be interpreted as a command, such as ‘delete this object’. Pen input is slower than typing for simple text input, but it opens possibilities for ‘informal interaction’ with the computer through sketches and gestures.

Direct manipulation

Direct manipulation describes the manipulation by pressing, dragging, and re-shaping, of objects displayed on a screen or in a simulated ‘visual world’. Instead of commanding the computer to perform an action, the user performs the action directly on the simulated object. *Virtual reality*, in which the user is placed in a simulated world presented on a screen or through helmets with miniature displays for each eye, is an extreme form of direct manipulation, where the objects in the simulation can be manipulated as if they were in the real world. But direct manipulation is not restricted to everyday objects. The same interaction techniques can be used to operate on diagrams, charts and documents. Direct manipulation can provide the user with a rapid response to actions and can lessen the gulf of execution and evaluation by allowing the system state to be changed directly, rather than obliquely through commands.

Direct manipulation allows many tasks to be performed quickly and naturally but it does not provide the expressiveness and combinatorial power of language. Even the restricted syntax of command line interfaces allows new expressions to be constructed by combining primitive commands, and natural language offers the subtlety of referring to classes of objects, of indicating conditionality (“if the diagram does not fit the screen then show me the centre portion”), and of indirect reference (“are there any objects which have the following properties ...”).

Methods of communication are being developed which combine the immediacy of direct manipulation and the expressiveness of natural language. These include *programming by example* where the user performs a sequence of operations and the machine records them as a generalised procedure, and direct communication with *agents* which perform complex, general tasks (such as arranging a meeting, or searching multiple databases) on behalf of the user.

1.6. Communicating through the computer

As computers were linked through *local area networks* providing communication between machines in a building or locality, and *wide area networks* such as Internet which connects

academic computers throughout the world, it became possible to use the computer as a communication device between people. Electronic mail is the main method of communicating between computer users. Each user of a computer network has a unique user name (such as *mike@cogs.sussex.ac.uk*) and a message can be sent from any computer on the network. Images and sounds, as well as text, can be sent as mail. *Bulletin boards* provide a means of storing messages on a range of topics. A user can connect to the bulletin board, read messages on a topic and add new ones. A *computer conference* is similar to a bulletin board, but for a restricted group of people exchanging messages on a fixed topic. Normally a computer conference will be moderated by a conference leader who invites participants, chooses a topic and keeps the discussion going by posing questions or countering responses. Each contribution is in the form of a written message posted to the conference, so a new participant can read previous messages to catch up with the discussion.

Techniques for compressing data, by automatically removing redundant information, and high-bandwidth data lines now allow direct communication from computer to computer by voice and moving image. These offer a styles of communication ranging from *videophones* to *awareness* systems (where windows on the screen give an impression of activities in other offices or buildings).

The computer acts as a *mediator* of person to person communication. It may be designed to be unobtrusive, as in a *video tunnel* with a camera directly behind the computer screen, so that users can converse and hold eye contact as if face to face, or it may offer a *shared medium* such as a space on the screen where two or more users can draw and write. The study of *computer supported cooperative work* is concerned with the design and use of computers to facilitate shared working.

2. The psychology of computer use

Most interactive computer systems are designed to assist tasks which engage the mind, so cognitive psychology has the potential to assist in their design and evaluation. There have, however, been surprisingly few psychological theories which make predictions about human-computer interaction, or which can be used directly to guide design. The findings of cognitive psychology have generally been used either as broad (and often inaccurate) guidelines, such as “show seven plus or minus two items on the screen”, or as general frameworks, such as the work on mental models leading to the design of systems based on deliberate metaphor (such as the ‘desktop metaphor’ of the Apple Macintosh interface).

Landauer, in an influential paper on ‘Relations between Cognitive Psychology and Computer System Design’ (Landauer, 1987), suggests four principal ways in which cognitive psychology can interact with computer system invention and design:

1. We may apply existing knowledge and theory directly to design problems.
2. We may apply psychological ideas and theories to the creation of new models, analyses and engineering tools.
3. We may apply methods of empirical research and data analysis to the evaluation of designs, design alternatives, and design principles.
4. We can use problems encountered in design to inform and guide our research into the fundamentals of mental life.

2.1. Applying psychology to design

A computer system which interacts with a human user should take account of the properties and limitations of the human mind and body. The nature of the human mind has been studied by successive generations of psychologists, and some results of this work are directly applicable to system design.

Memory

Findings from research on human memory that could influence computer system design include *recency* and *primacy effects*, *chunking*, and the *associativity* of semantic memory.

Information in short term memory decays rapidly and if items are presented in serial order, those items towards the end of the series will be remembered well for a short period of time (the recency effect) and items at the start of the series will be remembered well for a longer period of time (the primacy effect), and items in the middle of the list will be less well recalled after a short or longer delay. This suggests that list presentations, such as ‘pull-down’ screen menus, should be organised so that less important items are placed in the middle of the list or, if all items are important, then the user should be given assistance in recall and selection.

Miller’s work (Miller, 1956) suggests that human retention of short term information is limited to around seven *meaningful chunks*. Thus, if information on a computer screen can be grouped into meaningful chunks then it is likely to be more memorable than as disparate items. As a simple example, presenting the time as 19:55:32 is likely to be more memorable than as 195532. There has been considerable work on how to present computer commands in ways that are meaningful and well-structured. Suggested guidelines (from (Shneiderman, 1992)) include the following: choose names for commands which are meaningful and distinctive; group the commands into categories easily-understood categories; present the commands in a consistent format such as action–object (e.g. ‘delete–word’); support consistent rules of abbreviation (preferably truncation to one letter); offer frequent users the ability to create ‘macros’ which group a series of commands; limit the number of commands and the ways of accomplishing a task.

The final guideline above seems counter-intuitive, and certainly limiting the commands and functions too far could impede use, but *interference effects* can occur when there are alternative ways to perform tasks. For example, in Microsoft Word a selected piece of text can be turned to italic by (among other ways) selecting ‘italic’ from the ‘format’ menu, pressing the –i keys, or pressing –shift–i. Some of these methods work with other word processors, some don’t or they invoke other commands. Ensuring consistency of appearance and operation within a program and across related programs is a major aim of interface design. Manufacturers such as Apple, IBM and Sun Microsystems provide detailed *interface guidelines* to developers of software for their products, to ensure a consistent *look and feel* and to set standards for command names and methods of interaction.

Long-term memory appears to be organised around concepts sharing similar properties, and related by association. This associativity can be mirrored in the computer through *hypermedia* systems which present information as texts or images, with explicit links to associated items. Thus, a hypermedia guide to an art gallery might display a painting on the screen and provide buttons allowing the user to display other paintings by the same artist, or in the same style, or on display in the same room. The computer can also offer aids to

learning or invention by providing *external representations* of associative memory, allowing learner, writer, or designer to set down ideas as visual notes on the screen and to associate them by drawing labelled links.

Perception

A computer assaults the senses and demands a high level of attention. Unlike a book, it has an active light-emitting display; unlike a television, it requires the user to sit close to the screen and interact with it. Helmet mounted virtual-reality displays have been claimed to cause fatigue and disorientation after ten minutes of use. The ability of the computer to abuse the senses means that especial care should be given to designing interfaces which complement human perception. *Visual acuity* is the ability to distinguish fine detail. Acuity increases with brightness so, in general, a bright screen will be more readable than a dark one, and dark characters on a white background will be more readable than light characters on a dark background. But high luminance displays are perceived to flicker and, since flicker is more noticeable in peripheral vision, then the larger the screen the more it will appear to flicker.

The choice of colour for displays is fraught with difficulty. When colour displays first appeared, programmers splattered primary colours on the screen, hoping to make them more attractive, but generally causing confusion and headache. Different colours are perceived as having different depths (for example red appears to stand out, while blue recedes) and strong differences in visual depth can be fatiguing and distracting. Blue lines on a black background, or yellow lines on a white background, are hard to distinguish. Eight percent of the European male population is colour-blind and cannot distinguish red from green. In general, colour should be used sparingly, and for a purpose, such as to distinguish items in a list, or to indicate important information.

Maguire (Maguire, 1990) collated the following guidelines for the use of colour in displays:

1. A neutral colour helps set off a full colour.
2. Colours near in hue (on the hue circle) form pleasing combinations and will lend a predominant tone.
3. Complimentary colours (opposite on the hue circle) contrast and give a warm-cool effect.
4. Colours which are 120° apart on the hue circle are easy to distinguish but often appear garish.
5. Colour edging enhances polygons.
6. common denominations of colours should be used where possible. For example, red: danger or stop, yellow: caution or slow, green: OK or go.

Sound has been little used in interfaces, but experiments on *auditory icons* (which simulate everyday sounds in computer simulations of, for example, industrial processes), *earcons* (which use structured combinations of notes to represent actions and objects), suggest that sounds can be effective in complimenting visual presentations and in providing warnings and monitoring of off-screen processes. Dix et al. (1993) provide a good overview of multi-sensory systems.

Motor skills and selection

Studies of reaction time have some application to computer design, particularly in computer interfaces to machinery, but most interaction with computers is pre-meditated, the user

forms an intention and then specifies and performs an action, so considerations of speed and accuracy are generally more important.

Fitts' Law (Fitts, 1954) states that the time to move to a target on a screen is a function of the size of the target and the distance to be moved, i.e.:

$$\text{Movement time} = k + c \log_2 (d/s)$$

where d is the distance to target, s is the size of the target and k and c are empirically determined constants. This suggests that, in general, targets on the screen such as buttons should be large and close to the user's start point.

Landauer (1987) combines Fitts' Law with Hick's Law, which states that mean response time in simple decision tasks is a linear function of the transmitted information, to predict the time a user takes to find an item in a hierarchical menu. For example, an interface to a dictionary might ask the user to choose from (1) apple – mother, (2) motor – trouble, or (3) under – zebra. If the user selects one then it would offer a further selection such as (1) apple – cap, (2) car – keep, (3) key – mother, and so on. If b items are indicated at each level of the menu hierarchy, then what is the optimal structure for the menu? Hick's law suggests that broader menus (with more items per level) require a longer time to select an item, and Fitts' law suggests that more alternatives will require targets that are smaller and harder to select. If the decision and movement components are performed in sequence then their mean times will add to give a mean choice time per step of $k + c \log_2 b$. But the total search time is equal to the mean time per step, multiplied by the number of steps. A broad menu will have a longer choice time, but fewer steps. Combining the various factors produces the equation

$$T = k \log_b N + c \log_2 N$$

where N is the total number of items at the end of the menu. It indicates that b , the number of alternatives, only affects $\log_b N$, the number of steps, and so search time will be minimised by using as broad menus as feasible within the constraints of space and display. Studies of menu-selection tasks broadly support the finding and this is one of the few direct applications of experimental psychology to interface design.

Skill acquisition

An experienced user of a computer system will not perform the same actions and require the same assistance as a novice. Computer use is an acquired skill and studies of skill acquisition suggest a progression from learning facts and applying general procedures, through the acquisition of directly applicable procedures, to the automatic execution of tasks. An expert in a skill such as chess does not consider many alternatives and make explicit plans to the extent that a novice might, but instead relies on a meaningful pattern of pieces to prompt a good move.

The implications for human-computer interaction are that novice computer users should be offered help in learning the basic objects and actions of the system, through tutorials and context-specific help. The Macintosh interface, for example, offers *balloon help* whereby 'balloons' offering simple advice pop up over objects on the screen. As they gain expertise, users need practice and assistance in forming a *command set* of the most useful commands (which will depend on the user's needs and abilities). Experienced *power users* can be

offered ways to reduce time and effort by, for example, invoking commands directly from the keyboard rather than using a mouse and menu.

2.2. Using psychology to create new models

The earliest use of computers was to solve specific problems, in mathematics, engineering, ballistics, logic, and so the research psychology of problem-solving was applied and adapted to the study of computer use. Newell and Simon's 'information processing psychology' (Newell & Simon, 1972) both described human problem solving in information-processing terms and was applied to the study of problem solving with computers. It led directly to the development of cognitive models of problem solving, from high-level descriptions of the writing process (Flower & Hayes, 1980) to a keystroke-level model of interaction (Card, Moran, & Newell, 1983). These models have been used to predict behaviour in problem-solving with computers and to inform the design of more effective and comprehensive computer systems.

More recently it has been recognised that people use computers for tasks such as design, exploration, and communication that do not easily fit the problem-solving mould. New models are being developed which draw on studies of human creativity, exploration, and social interaction. The modelling of human-computer interaction is discussed in section 3.

2.3. Applying psychology to evaluation

The computer is a piece of equipment and many of the experimental psychology techniques for evaluating human interaction with equipment are applicable to human-computer interaction. For example, Reece (ref) has carried out an elegant investigation to compare the relative merits of pen and paper, a computer word processor, and a simulated speech-driven word processor using standard experimental methods. Shneiderman (1992, p. 18) indicates five measurable human factors which are central to evaluation: time to learn; speed of performance; rate of errors by users; retention over time; subjective satisfaction.

Where the computer differs from everyday equipment is in its flexibility of function and appearance, resulting in a large space of possible system designs and presentations. Much of the work in human-computer evaluation is not concerned with testing clear hypotheses and comparing fixed products, but in assessing the quality of early designs. New techniques such as *Wizard of Oz* evaluation (where a human pretends to be the computer, for example by mimicking a speech-understanding system) and *design space analysis* for organising and comparing alternative designs, have been devised to evaluate prototypes and design possibilities rather than finished systems. Section 6 covers the topic of system evaluation in more detail.

2.4. Applying studies of human-computer interaction to psychology

The versatility of the computer makes it a valuable device for studying the psychology of learning and problem solving. As well providing a tool for problem solving and a medium for communication, it can automatically record data and analyse patterns of interactions.

The study of children as programmers has provided insight into children's development of problem-solving skills (such as the use of physical devices like the Logo 'turtle' to catalyse understanding (Papert, 1980)).

Investigations of programming as a skill have revealed differences between the problem-solving behaviour of novices and experts (Jeffries, Turner, Polson, & Atwood, 1981). Both

novices and experts tend to start from a general problem goal and refine it into sub-goals and down to lines of program code, but novices tend to expand one part of the problem down to its lowest level before starting on the next part, whereas experts generally consider the entire problem before developing it to a deeper level.

Computer games have formed the basis of studies of motivation (Malone, 1981) and collaboration and conflict (Hewitt, Gilbert, Jirotko, & Wilbur, 1990).

As people begin to use computers as an everyday tool to augment their intellect, to extend their memory and to distribute cognition amongst a group of co-workers, this has led to the study of computer-augmented and computer-mediated cognition.

3. Modelling Human-Computer Interaction

Models of human-computer interaction serve much the same purposes as architectural, scientific or engineering models. They can be used to predict behaviour, to assist in design and to evaluate competing theories and designs. Where they differ is in the importance given to modelling human cognition. Understanding the psychology of the computer user is important to creating systems which solve interesting problems, respond in appropriate ways, and are engaging and easy to use.

There has been much confusion in the literature about the types, names and purposes of models of HCI. Young (Young, 1983) attempted to sort out the mess by suggesting two types of models – models of the user, and models of the computer – which may be held by different entities: designers, researchers, computer systems, and users. Thus, a researcher may develop a model of a user in order to understand the psychology of computer use, or a computer may hold a rudimentary model of its user so as to offer individualised help or guidance. The models with most significance to HCI are the designer's model of the computer, the designer's model of the user, and the user's model of the computer.

3.1. The designer's model of the computer

A software designer has direct control over the form and function of the software, but of not the people who use it, so designers' models of the computer tend to be more detailed and formal than those of its users. Software descriptions can broadly be divided into *syntactic models* concerned with the structure of the dialogue between user and machine, and *semantic models* which describe what the user's actions do.

The simplest syntactic models describe the desired states of the computer system and transitions between them. They can be represented as *state transition diagrams*, with each state shown as a circle and each transition as a labelled, arrowed line. The transitions correspond to actions performed by the user, such as selecting from a menu or pressing a key. Figure 1 shows a state transition diagram for a simplified bank Automated Teller Machine.

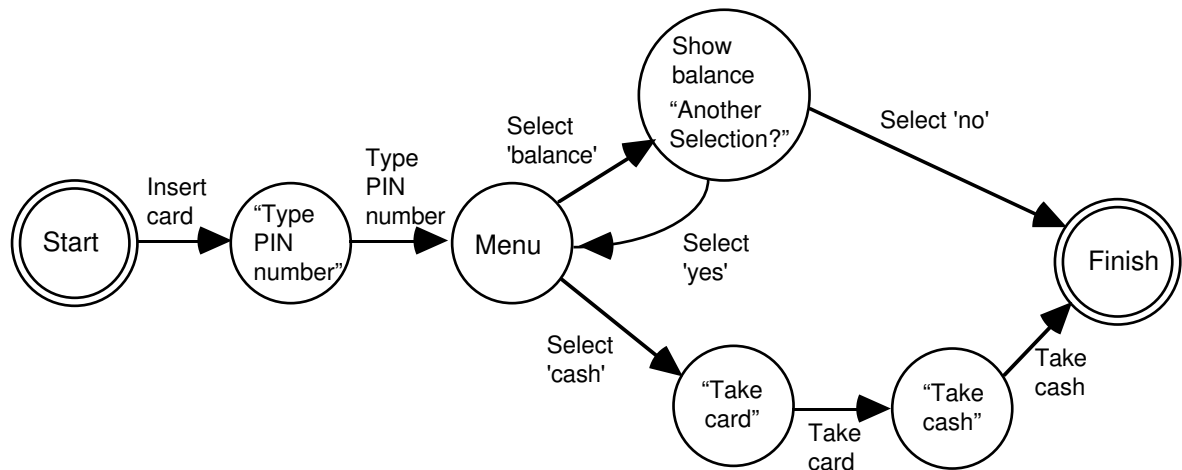


Figure 1. A state transition diagram for a simplified Automated Teller Machine

A state transition network is easy to create and understand, but for a large system it can become complex and tangled. A *hierarchical state transition network* deals with complexity by showing sub-systems as ‘boxes’ in the diagram where each box can be described as a separate diagram. Another problem is that the diagrams do not show events and dialogues which can occur concurrently or which interrupt each other, such as might appear in a computer-aided design package with concurrently active windows, menus, palettes, and forms. Harel’s *state charts* provide a way of showing complex, concurrent systems as diagrams.

An alternative way of representing the syntax of computer systems is as formal grammars or rules. The *BNF* notation is typically used to describe the structure of programming languages, but has been applied to HCI. Condition-action rules, or *Production rules*, of the form

IF (condition) THEN (action)

can specify the events which change the system state or cause program actions such as

Insert card → Wait-for-PIN < display “Type PIN number” >

where the event of inserting a card causes the ATM to display the message “Type PIN number” and puts it into a state of waiting for a PIN number to be typed. Production rules are better than state transition networks at describing concurrent events, but less appropriate for specifying sequences of events.

Syntactic models of interaction with the computer specify not just the transitions between states of the computer system, but also what functions are invoked at each transition. The most straightforward way is to add functional descriptions to syntactic representations such as state transition diagrams or production rules. These could be specified in a notation of formal semantics, or as simply as calls to programming language procedures. More sophisticated formalisms, such as *Z* and *VDM*, provide a complete formal notation to

describe the semantics of programs. They were developed to reason about the behaviour of programs, but have been adapted for producing precise specifications of interactive software. Producing a formal description of an interactive system can be difficult and time-consuming, and it may tell what a designer most wants to know, such as “will the program be easy and intuitive to use?”.

3.2. The designer’s model of the user

The ideal user as far as the software designer is concerned is someone who behaves in a predictable manner, forming clear goals and then carrying them out in a series of steps that can be programmed as basic operations for the computer. In practice, no user behaves in this way. Real users are unpredictable, have conflicting goals (or may not even be able to express their needs in terms of clear goals), change their minds, make mistakes and may end up muddled and maddened. Not surprisingly, it is easier to model an ideal user than a real one, and an idealised model can help the designer to produce software which performs well with competent users, but too little attention has been given in HCI to individual differences in software use.

The most influential approach to modelling the user is Card, Moran and Newell’s GOMS (goals, operators, methods, and selection rules) (Card, et al., 1983) method. GOMS characterises a user of a computer system as an information processor, with interacting perceptual, motor and cognitive subsystems. It assumes that the user acts rationally by forming *goals* which define some state to be achieved. A goal can be achieved by one or more alternative *methods*, where a method is a sequence of elementary perceptual, motor or cognitive *operations* (such as ‘Move cursor over word’, ‘Double click mouse button’, ‘Press delete key’). *Selection rules* choose between alternative methods to accomplish a goal. Using the GOMS formalism a high-level goal, such as ‘edit manuscript’ can be broken down into a sequences of methods and operators, where each operation can be provided by an elemental computer procedure.

GOMS can make some broad predictions of human behaviour. The depth of the sub-goals gives an indication of the load on a user’s short term memory, and the length of the sequence of operations provides a rough measure of the time taken to carry out a method.

Card *et al.* (Card, Moran, & Newell, 1980) demonstrated the power of GOMS by using it to predict the time expert users will take to execute tasks on a computer. They proposed that, at the level of individual motor actions, the time taken to execute a computer task can be described as the sum of the times taken in keystroking, pointing (e.g. with a mouse), homing (on a target), drawing, thinking (by the user), and responding (by the system). They encoded methods, such as editing text with keyboard and mouse-operated editors, in the GOMS notation and produced estimates of individual operations by means such as Fitts’ Law. This enabled them to predict the times taken to carry out tasks such as replacing a 5-letter word with another 5-letter word. Comparing their predictions with the performance of human users, they found that the time required for experts to perform a unit task can be predicted to within 20 per cent.

Kieras and Poulson (Kieras & Polson, 1985) clarified and enriched the GOMS model by providing two parallel descriptions, of the user and the computer. The user’s knowledge is expressed as condition-action rules where the condition can match goals, information in the user’s memory, or information provided by the computer. If a pattern matches successfully, the action part of the rule is performed, to alter the representation of the user’s memory or

carry out some operation on the computer. For example, a simple rule for moving the cursor to a given line might be (from (Dix, et al., 1993, p. 199)).

```
(MOVE-UP
IF      (AND      (TEST-GOAL move-up
                  (TEST-CURSOR-BELOW %LINE) )
THEN (      (DO-KEYSTROKE 'K') ) )
```

meaning that if the user has the goal of 'move-up' and the current position of the cursor is below the desired line on the screen then perform the operation of pressing key 'k' (which in the particular text editor moves the cursor up a line). This rule would 'fire' repeatedly so long as the condition is satisfied, i.e. until the cursor is at the desired line. To describe the states and operations of the computer Kieras and Poulson use a form of state transition network. Combining a model of the user with one of the computer system enables the method to describe mappings and mismatches between a user's goals and the provisions of the computer.

Other formalisms for modelling the user include Task Action Grammars, and Programmable User Models.

GOMS and its associates offer fine-grain accounts of a user's interaction with the computer. But particularly in the early stages of design, a system designer may be attempting to understand the task and choose between alternative designs. In that case the designer needs a more general model of problem-solving and user behaviour which can help to ensure that the software is comprehensive and suited to a broad range of users. General cognitive models of problem-solving tasks, such as Flower and Hayes' account of the writing process (Flower & Hayes, 1980), can provide frameworks and justifications for design.

3.3. The user's model of the computer

People spontaneously employ *mental models* to assist their understanding of complex systems. Mental models are a simplifications, often in terms of an analogy such as a telephone as a 'speaking tube'. They frame our understanding of the world and enable people to operate and make predictions about new technology, by relating it to prior experience.

The term *system image* (Norman, 1986) describes the user's conception of a computer system, which may be in terms of an analogy or a set of functions. If the user is given a system image before meeting the system, then the image will determine how the user conceives and begins to operate the equipment. Designers can draw on mental models by designing software to fit an analogy. Thus, electronic mail has deliberate similarities to the postal system, with 'mail boxes', 'addresses', and 'postmasters'. A system image need not represent the way the system is actually constructed, but a well-chosen system image should enable the user to interact with the computer in a way that meets with expectations. It allows the user to make predictions about the system's behaviour and to be satisfied that the system is not going to carry out unexpected actions.

So long as the computer works according to the system image then the user is led treat it a familiar object, but if the system image ever fails then the user is left stranded, without the depth of knowledge needed to repair the problem. For example, if a computer network fails

to deliver an electronic mail message then the message is usually returned to the sender with detailed diagnosis of the error. The diagnostic information is given in terms of gateway computers and message handling programs, far from the image of electronic mail as a postal system. Studies of breakdowns in human-computer interaction, and the user's subsequent attempts at repair can be valuable in revealing people's (sometimes surprising) mental models of technology.

4. System development

The conventional approach to computer system development is to split it into a number of discrete stages, leading up to the finished working product. Definitions of the stages vary, but they usually include:

requirements analysis, which describes the purpose of the system along with the main tasks it has to perform,

system design or *functional specification* which specifies how the tasks are to be performed, and breaks the system down into components which either need to be programmed or adapted from existing software,

detailed design where each component of the system is described in sufficient detail that a programmer can code it,

implementation in a suitable programming language,

integration and testing of the different pieces of program,

maintenance of the system, by correcting errors and updating the software to cope with changes in hardware and requirements.

This staged or *waterfall* approach was introduced to assist the development of large corporate systems such as payroll packages, which typically have tight specifications, large teams of programmers, and very little interactivity. The basic principles, of ensuring that the software meets requirements and that it should be well integrated and tested, apply equally to payroll packages and painting programs, but the main difference is that the operation of highly interactive systems cannot be fully specified in advance of implementation. Any reasonably complex interactive program will be used in ways unforeseen by its designers. For example, spreadsheet programs were originally introduced to ease the job of repetitious calculation, but since their introduction users have adopted spreadsheets for forecasting, visual presentation of information, database management, timetabling, and many other purposes.

User-centred design is a general term for a design process which considers the needs of users in the design of interactive systems. A major industry has grown around the waterfall method of software development, with structured design methods such as JSD (Jackson Structured Design) and SSADM (Structured Systems Analysis and Design Methodology) promoting a strict discipline of design and testing. By contrast user-centred design is inherently messy, based on imprecise psychological models, and unpredictable, relying on

designers to respond to users' demands. Not surprisingly, there have been attempts to reconcile the two, by adapting the waterfall model to the design of interactive systems.

4.1. Adaptations of the conventional software development method

Usability engineering provides information to guide software development by specifying *usability metrics* which the finished system should satisfy. Usability metrics are normally quantitative measures in the areas of *learnability* (the time and effort needed for learning to use the system), *throughput* (the number of tasks accomplished in a given time, and the errors made), *flexibility* (the extent to which users can adapt the system to new ways of working as they become more experienced in its use), and *attitude* (the attitude of the users towards the system). The collection of measures, such as 'time taken to install the system', can be specified in advance as part of the system design and they provide guidelines inside which the programmers should work.

Usability metrics are only one type of guideline. Other more qualitative guidelines can specify general attributes of the system, such as the colours to be used or the appearance of objects on the screen, or they can be very specific, for example indicating the phrases to be used in error messages. At their best, guidelines reflect good design practice, match what is known about the needs of the users, and give the software a consistent 'look and feel'. The major software manufacturers issue books of guidelines to their software developers, such as Apple's *Human Interface Guidelines* (Apple Computer Inc., 1987).

User Interface Design Environments automate or assist the work of designing user interfaces by providing a kit of 'interface parts' such as buttons, menus and forms, each with appropriate functionality, which can be fitted together into a complete user interface. Since each of the components adheres to interface guidelines, a design environment can provide a short cut to design and an assurance of basic quality.

4.2. Departures from the conventional software development method

The departures from the conventional method of software development can be classed under the general heading of *user centred design*, since they give high priority to the needs of the eventual users, either by analysing the tasks that typical users will perform, or by involving actual users in the design process. The second approach is the more radical because it requires an empirical 'trial and error' process, with designers testing out early prototypes of the software on users and making repeated revisions to the plans. In terms of the 'waterfall model' of software design, the later stages of the design process provide information to the designer which may result in revisions of the earlier stages and a series of evolving prototype systems. In some cases, such as for the development of 'walk up and use' programs for the general public, users may be recruited as part of the design team, offering advice and criticising mockups and prototypes.

Task Analysis is the study of what people do when they carry out tasks. It differs from user models such as GOMS in that the emphasis is on observable behavior, rather what might be happening in the user's mind. Task analysis has a long history in the study of work behaviour. Its more modern variants, such as *TAKD* (task analysis for knowledge description) (Diaper, 1989) combine a behavioural description with an analysis of the knowledge that people require to perform the task. The start of a task analysis is usually to break down an observed task into a hierarchy of sub-tasks accompanied by plans which

describe in what order and under what conditions the sub-tasks are performed. Figure 2 shows a task hierarchy for making cups of tea (from (Dix, et al., 1993)).

Figure 2. Task hierarchy for making cups of tea (from (Dix, et al., 1993))

The analysis may then go on to produce taxonomies of the objects and actions involved in the task, accounts of the goals of a person engaged in the task, and plans for accomplishing a goal.

To carry out a task analysis requires detailed observation and recording of people engaged in appropriate activities. A video recording can reveal the structure and sequence of activities, but other methods are needed to infer a user's goals and plans. One general and useful technique is *protocol analysis*. A person is asked to talk about their intentions and plans at the same time as performing the task. These *concurrent verbal protocols* are recorded and matched against the activities, to give an understanding of not only *what* activity is being performed, but also *why* it is being carried out, for what purpose and to achieve what goal.

Task analysis can inform system design by showing the structure and sequence of a task, such as composing a document or designing a building, which could be supported by a computer. It can also be carried out for people working with computers, to measure usability or to study problems and breakdowns in performance.

Rapid prototyping is one of a number of methods to involve eventual users in the early stages of design. One or more mockups or prototypes are constructed which exemplify some aspects of the system, such as alternative interfaces or different presentations of the data. The prototypes may be quite simple. They could be in the form of sketches of an

interface, ‘story-boards’ or ‘slide-shows’ to show a sequence of actions. These prototypes are shown to users, for comment, more systematic surveys and comparisons, or to raise users’ understanding and enable them to contribute to further design. Rapid prototyping environments, such as HyperCard, allow screen displays to be constructed rapidly using screen drawing tools, and then given limited functions, such as allowing objects on the screen to become buttons which, when pressed, cause a move to another screen display. The original rapid prototyping tools required the mockups to be thrown away once they had been tested, but more recent development environments offer more powerful programming languages, allowing the mockups and early prototypes to be developed into full systems. Rapid prototyping is one method employed in the *iterative* approach to software development, which maintains that, unlike conventional software development, the design of interactive systems is not primarily concerned with specifying the functions and input-output behaviour of the computer, but with fitting the computer into the plans and interleaved activities of a human user. Not only are the plans and actions of users of paramount importance to the software designer, but they are inherently ill-defined. The art of iterative system design is to *anticipate* the many ways in which the system might be used, and to *adapt* the system to suit the users’ needs as they become identified.

5. Evaluation

Evaluation of HCI is concerned with making judgements and measurements of the usability and potential usability of interactive computer systems. an evaluation can be carried out at any point in the software lifecycle: to choose between competing requirements, to decide what features and functions should be included in a proposed system, to validate the system with respect to the requirements, to verify that the system meets the specification, to compare different systems of implementations, or to provide metrics such as learnability, ease of use, and speed of operation. The type of evaluation can range from formal measures of usability, to ensure that the system meets industry standards, to offering the system for informal comment.

The Open University Guide to Usability (The Open University, 1990) classifies evaluation methods into five categories:

Analytic evaluation uses semi-formal methods, such as task analysis, to predict the performance of expert users of a system. It allows designers to predict the performance of users in terms of the physical activities and cognitive operations they must carry out to perform a given task. It can be applied early in the development process, reducing the need for iterative prototyping, but it should be used with care since it is intended to model the performance of an ideal user carrying out tasks without error.

Expert evaluation calls on people who are experienced in interface design or human factors to make judgements about the usability of the system and to suggest improvements. This can range from demonstrating the system to colleagues for informal comment, to presenting a set of design issues which need to be resolved.

Observational evaluation involves collecting data on how users interact with a prototype of finished system. It may be an controlled study, carried out in a usability laboratory, or a naturalistic study, with the system set up in the workplace of a potential user. The data

can come from direct observation and field notes, video recordings, automatic logging of the user's interactions, or analysis of verbal protocols.

Survey evaluation asks users to give their opinions of a system that they have used, through a structured interview or questionnaire.

Experimental evaluation requires the investigator to form hypotheses about the interaction which are then tested using the methods of experimental psychology. A typical experimental evaluation might be to compare the time taken to perform a given task, such as editing a document, using two different interfaces.

The different approaches to evaluation reflect the paradigms of psychology, and come laden with prejudice and anecdotal evidence about the 'best' way to develop and assess interactive software.

A team from Hewlett-Packard laboratories performed a careful comparison of four methods of user interface evaluation (Jeffries, Miller, Wharton, & Uyeda, 1991): heuristic evaluation, software guidelines, cognitive walkthroughs and usability testing. Heuristic evaluation is similar to 'expert evaluation' above, where specialists in user interfaces study the interface in depth and look for properties that they know, from experience, will lead to usability problems. The use of published guidelines is not strictly a method of evaluation, but it does provide specific recommendations about interface design which can lead to the identification of problems. A cognitive walkthrough is a type of analytic evaluation where the developers of an interface carry out a series of tasks that a typical user might perform. They compare the actions of the interface with the user's goals and knowledge, and record discrepancies between the user's expectations and the steps required by the interface. Usability testing begins by analyzing the users' needs and setting usability criteria for the system. Experimental evaluation is then used to test whether the system meets the criteria.

The comparison of the different techniques found that heuristic evaluation produced the best results. It found the most problems with the interface (50% of the known problems compared with around 14% for each of the other methods) at the least cost in terms of time spent on the evaluation. It does, however, depend on finding several people with the knowledge and experience necessary to apply the technique.

The main divergence in approach to evaluation is between laboratory studies, carried out in a *usability laboratory*, under carefully controlled conditions, and *situated studies* of people interacting with computer systems in their normal working environment.

The main advantage of usability engineering is that it sets explicit measures of usability (such as the time taken to perform a task, the number of errors, the number of commands used, and the ratio of successful completions to failures) which can be used to judge the value of the system once it is finished. The concomitant problem is that the measures may not capture the real qualities and failings of the system, and may divert the designer into fixing minor defects rather than considering alternative solutions. For example, usability measures applied to the design of a word processor may help in re-arranging items on a menu (for some word processors this is equivalent to re-arranging deckchairs on the Titanic), but they do not address the deeper issues of "how do people write?" and "what computer tools can best support the writing process?"

Situated studies can call on the techniques of ethnography and social psychology to give a rich picture of the ways in which technology mediates the activity of individuals and groups.

They can reveal breakdowns in understanding or communication caused by system failure or human misconception, and they can suggest new areas of concern to a system developer (such as the importance of tactile feedback). But the amount of data collected from such studies can be overwhelming, and there may be no guarantee that a detailed investigation of one working environment can be applied to the technology set in a different situation.

The best conclusion is that usability measures and situated studies can inform system design, but there is no substitute for the heuristic knowledge and experience of experts in HCI.

6. Beyond Human-Computer Interaction

The phrase Human Computer Interaction invokes an image of a person sitting at a console, communicating with a computer program. That image fits only one aspect of computer use and is becoming increasingly inappropriate. Lawrence Tesler, in one of a perceptive series of articles on computer networks in *Scientific American*, (Tesler, 1991) describes four paradigms of computing, associated with the four decades from the 1960s to 1990s (see figure 3).

	BATCH	TIME-SHARING	DESKTOP	NETWORK
DECADE	1960s	1970s	1980s	1990s
TECHNOLOGY	MEDIUM-SCALE INTEGRATION	LARGE-SCALE INTEGRATION	VERY LARGE SCALE	ULTRA LARGE SCALE
LOCATION	COMPUTER ROOM	TERMINAL ROOM	DESKTOP	MOBILE
USERS	EXPERTS	SPECIALISTS	INDIVIDUALS	GROUPS
USER STATUS	SUBSERVIENCE	DEPENDENCE	INDEPENDENCE	FREEDOM
DATA	ALPHA-NUMERIC	TEXT, VECTOR	FONTS, GRAPHICS	SCRIPT, VOICE
OBJECTIVE	CALCULATE	ACCESS	PRESENT	COMMUNICATE
USER ACTIVITY	PUNCH & TRY (SUBMIT)	REMEMBER & TYPE (INTERACT)	SEE & POINT (DRIVE)	ASK & TELL (DELEGATE)
OPERATION	PROCESS	EDIT	LAYOUT	ORCHESTRATE
INTER-CONNECT	PERIPHERALS	TERMINALS	DESKTOPS	PALMTOPS
APPLICATIONS	CUSTOM	STANDARD	GENERIC	COMPONENTS
LANGUAGES	COBOL, FORTRAN	PL/1, BASIC	PASCAL, C	OBJECT ORIENTED

Figure 3. The Four Paradigms of Computing (from (Tesler, 1991))

The 1990s is seeing the emergence of fully mobile computers, in communication by wireless networks. They can integrate communication services (such as fax and electronic mail), personal organisers (such as diaries and meeting schedulers), and work tools (such as document processors and expert systems). Instead of interacting with a computer, a person will work with many tools and applications, situated on a number of computers throughout the network, and will communicate through the computers to other people.

6.1. Computer Supported Cooperative Work

Computer Supported Cooperative Work (CSCW) is the term used to describe a group of people working together, with the aid of computers, on a common task. It can range from a group of people in a meeting room using laptop computers to contribute to a shared 'electronic whiteboard', to a conference conducted over electronic mail. CSCW allows people in physically dispersed organisations to work together, and a study by Sproull and Kiesler (Sproull & Kiesler, 1991) showed that using a network induced the participants to talk more frankly and more equally. People who were reluctant to contribute to a face-to-face

meeting reported feeling more lively and confident when expressing themselves in a computer conference.

CSCW brings a social and organizational dimension to human-computer interaction. Computer-mediated communication is less rich than face-to-face discussion, lacking the visual cues that signal, for example, impatience, annoyance, and dissent. Partly as a consequence, people in a computer conference took longer to reach a consensus than an equivalent face-to-face group, and they tended to express more extreme opinions. On a wider scale, communication by computer can alter patterns of work and the exercise of power within an organisation. The free sharing of information and frank exchange of views which typifies network discussion may threaten established management and information channels within an organization. On a larger scale, worldwide networks provide an 'information anarchy' across political boundaries.

The Internet already links together several million users on more than 300,000 computers in around 30 countries. An Internet user can send mail to any other user, connect to remote machines, and share work across the network. The Mosaic software gives a global hypertext network to each Internet user, so that a single menu of options might offer satellite weather maps, catalogues from major academic libraries in the United States, and a program to search for the email address of any user of the Internet. Clicking one of these options will automatically connect to the relevant computer and run software to display a map, show a library catalogue index or return an email address.

The information carried on Internet and the services it provides are not managed by a central agency, but are contributed by individual sites and users. Access to such a wealth and diversity of information can be seen as overwhelming (there are thousands of 'newsgroups' on the Internet, each contributing a hundred or more messages a day), threatening (people throughout the world can exchange politically and socially inflammatory messages), or liberating (computer networks cut across established boundaries of power and influence). New social and technical challenges include finding ways to allow colleagues who have never met face to face to work together productively, developing 'information rejection' methods to filter out unwanted information, and providing mechanisms to manage the electronic pathways, to stop them being choked by an exponential growth in communication or by the deliberate sabotage of computer viruses.

6.3. Autonomous agents

As networks grow beyond the comprehension of any individual, computers will play an increasing part in retrieving information and managing interaction. *Autonomous agents* are programs which assist with some computer-based task, such as arranging a meeting, finding a document across computer networks, or filtering email according to priority. A user might ask an agent to arrange a meeting with selected colleagues around a given date. Once the agent is activated, it carries out the job of arranging a meeting autonomously. It consults the on-line calendars of the colleagues, finds a range of dates and times, suggests a suitable time and venue to the colleagues by mail and, when agreed, confirms the arrangement.

The style of interaction with an autonomous agent is very different to a command language such as DOS, or the point and click of direct manipulation. Rather than issuing commands, the user specifies needs and constraints, for example by partially completing a form or by indicating ranges on a scale. The interaction is one of request and negotiation. An

autonomous agent requesting a meeting with the company director may need to ask and respond in a very different way to one which negotiates with a colleague. Computer etiquette and the micropolitics of organizations are new challenges for software designers.

6.2. Ubiquitous computing

In time, the computer may fade away altogether. An average house already has ten or more microcomputers. They are embedded in washing machines, video recorders, telephones, remote controllers, thermostats, clocks, microwave ovens, door bells, cameras, toasters, and electric shavers. We do not notice them as computers because they are programmed to perform a restricted task and because their interface or system image is that of a familiar consumer device. However, they use general-purpose computer chips and in the future they will perform a wider range of tasks, such as providing diagnostic information, or communicating with other domestic appliances.

Ubiquitous computing can mean giving everyday objects computing power while retaining their familiar forms. A pad of paper might transmit what is written on it, a telephone might provide a printed transcript of a conversation, or a window might show an image of the scene outside some hours before. Or it can refer to the seamless integration of computer and non-computer objects. Xerox EuroPARC are developing a DigitalDesk to demonstrate the movement of information between paper and computer. They have suspended a video camera and a projector above an ordinary desk, so that electronic documents can be projected onto the desk, and paper ones can automatically be digitised into computer text. Software connected to the camera will be able to recognise hand gestures, so that a person working at the desk can move the projected documents around just as they would push sheets of paper.

Perhaps the consumer device which will change the most as it is invaded by the computer is television. Digital television will combine the interactivity of computer games with the image quality of television opening a vista of 'edutainment' from interactive soap operas to self-guided tours of the great museums. It will demand a new range of skills, combining television production, electronic art, and software design and it may provide a new meaning to 'human-computer interaction', as millions of people simultaneously participate in an interactive television production.

6.4. The challenge of human-computer interaction

Human-computer interaction is a new discipline, and it has had little opportunity to mature, because the nature of human interaction with computers is changing so rapidly. At worst the literature on HCI is a set of anecdotes and good intentions, but at best it amalgamates the psychology and technology of computing, begins to turn software design from an art into a science, and offers guidelines for good practice in developing and deploying computer systems. It has responded to new topics such as user-centred design, and computer supported cooperative working and it has informed the design more usable computers such as the Apple Macintosh. It must now address both technical issues such as network integration and compatibility of software, and the broad areas of social interaction, organizational computing, and the ubiquitous computer.

7. References

- Apple Computer Inc. (1987) *Apple Human Interface Guidelines: The Apple Desktop Interface*. Reading, MA: Addison-Wesley.
- Baecker, R. M., and Buxton, W. A. S. (1987) 'An Historical and Intellectual Perspective.' In R. M. Baecker & W. A. S. Buxton (Eds.), *Readings in Human-Computer Interaction* San Mateo, California: Morgan Kaufman Publishers.
- Card, S., Moran, T. P., and Newell, A. (1983) *The Psychology of Human-Computer Interaction*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Card, S. K., Moran, T. P., and Newell, A. (1980) 'The keystroke-level model for user performance time with interactive systems.' *Communication of the ACM*, 23, 396–410.
- Diaper, D. (1989) 'Task Analysis for Knowledge Descriptions (TAKD); the method and an example.' In D. Diaper (Eds.), *Task analysis for Human-Computer Interaction* (pp. 108–159). Chichester: Ellis-Horwood.
- Dix, A., Finlay, J., Abowd, G., and Beale, R. (1993) *Human-Computer Interaction*. New York: Prentice Hall.
- Fitts, P. M. (1954) 'The information capacity of the human motor system in controlling amplitude of movement.' *Journal of Experimental Psychology*, 47, 381–391.
- Flower, L. S., and Hayes, J. R. (1980) 'The dynamics of composing: making plans and juggling constraints.' In L. Gregg & E. Steinberg (Eds.), *Cognitive Processes in Writing: an Interdisciplinary Approach* (pp. 31-49). Hillsdale NJ: Lawrence Erlbaum.
- Hewitt, B., Gilbert, N., Jirotko, M., and Wilbur, S. (1990) *Theories of multi-party interaction*. Technical Report, Social and Computer Sciences Research Group, University of Surrey and Queen Mary and Westfield Colleges, University of London.
- Jeffries, R., Miller, J. R., Wharton, C., and Uyeda, K. M. (1991) 'User interface evaluation in the real world: a comparison of four techniques.' In *ACM CHI '91* (pp. 119–124). New Orleans, LA: ACM, New York.
- Jeffries, R., Turner, A. A., Polson, P. G., and Atwood, M. E. (1981) 'The processes involved in designing software.' In J. R. Anderson (Eds.), *Cognitive Skills and their Acquisition* Hillsdale, N.J.: Lawrence Erlbaum.
- Kieras, D. E., and Polson, P. G. (1985) 'An approach to the formal analysis of user complexity.' *International Journal of Man-Machine Studies*, 22, 365–394.
- Landauer, T. K. (1987) 'Relations between cognitive psychology and computer system design.' In J. M. Carroll (Eds.), *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction* (pp. 1–25). Cambridge, Ma.: MIT Press.
- Maguire, M. C. (1990) 'A review of human factors guidelines and techniques for the design of graphical human-computer interfaces.' In J. Preece & L. Keller (Eds.), *Human-Computer Interaction* (pp. 161–184). Hemel Hempstead: Prentice Hall.
- Malone, T. W. (1981, December 1981) 'What makes computer games fun?' *BYTE*, p. 258–277.
- Miller, G. A. (1956) 'The magical number seven, plus or minus two: some limits on our capacity for processing information.' *Psychological Review*, 63, 81–97.
- Newell, A., and Simon, H. (1972) *Human Problem Solving*. Englewood Cliffs: N.J.: Prentice Hall.
- Norman, D. A. (1986) 'Cognitive Engineering.' In D. A. Norman & S. W. Draper (Eds.), *User Centred System Design* Hillsdale, New Jersey: Lawrence Erlbaum.

- Papert, S. (1980) *Mindstorms: Children, Computers and Powerful Ideas*. New York: Basic Books.
- Shneiderman, B. (1992) *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Reading, Massachusetts: Addison-Wesley.
- Sproull, L., and Kiesler, S. (1991, September 1991) 'Computers, networks and work.' *Scientific American*, p. 84–91.
- Tesler, L. G. (1991, September 1991) 'Networked computing in the 1990s.' *Scientific American*, p. 54-61.
- The Open University (1990) *A Guide to Usability*. Milton Keynes: The Open University.
- Young, R. M. (1983) 'Surrogates and mappings: two kinds of conceptual models for interactive devices.' In D. Gentner & A. L. Stevens (Eds.), *Mental Models* (pp. 35–52). Hillsdale, N.J.: Lawrence Erlbaum.