Applied Ergonomics 1988, 19.1, 55–66

The psychology of personal information management

M. Lansdale

Department of Human Sciences, University of Technology, Loughborough, UK.

A requirement of 'The Office of the Future' is that it provides us with an effective way of storing and retrieving information. But existing IT products go nowhere near supporting the variety of activities which can be observed in paper-based offices, and it is not surprising that concepts of the 'paperless office' are as far off as they were when the idea was first mooted. This paper illustrates how many of the issues involved in the automation of information management are essentially psychological in nature. These principally devolve upon the processes of recall, recognition and categorisation. Examples of existing information management techniques show how there is a trend to automate with a view to simulating office practices, or to develop according to the availability of technological solutions. Both of these are inefficient with respect to the user's psychological needs. A framework for developing user-oriented information management systems is discussed and relevant research issues presented.

Keywords: Psychology, information, management

Introduction

We all keep information in our work and domestic lives. It may be books, notes, folders, diaries, personal records, files or whatever. This is personal information not necessarily in the sense that it is private, but that we have it for our own use. We own it, and would feel deprived if it were taken away. The primary reason (there may be others) for keeping this information is to be able to retrieve and use it in the future. This article considers the process of managing this information: the methods and procedures by which we handle, categorise and retrieve information on a day-to-day basis.

As a psychologist, the motivation for my interest is that this represents an area in which psychological theory can contribute to the satisfaction of a clear technological need. We are witnessing in the emergence of Information Technology, a facility in which all the pressures are to create, store and process more and more information. But the purpose of IT should be to increase the quality, not merely the quantity, of available information. If all it achieves is to increase the volume of our filed information without an associated improvement in retrievability, then the reverse must be expected. The psychological aspect of this is that the management of personal information, by definition, involves psychological processes: The information to be retrieved has already been handled, categorised and filed away by the individual. It is therefore reasonable to ask how much the problems of information management represent a shortfall in these processes and the ability to remember what was done. Consider some examples:

"My boss wants to see all the project reviews I have carried out over the last six months. The trouble is, they are filed under each of the individual projects. It will take me ages to work through and dig them all out".

"I know what the thing looked like: it has a blue and white stripe at the top, but I can't remember how I filed it, or even what it was about."

"Yes I remember that paper. It came at the same time as the product audit. I can't remember what happened to it, though."

"The document I want is the French Finance Committee's minutes, but I've tried looking under 'French', and 'Committees', and it's not there. Perhaps it's under 'France'. I don't know. I may as well search through the lot."

I would not like to suggest that all the problems of information management are psychological in nature. But these examples illustrate two general points which will re-appear throughout this paper and which represent important issues in the development of information management tools on computers. First is that there is a general problem in categorising items, both in terms of deciding which categorisations to use, and in remembering later exactly what label was assigned to that categorisation. The second is that we remember far more about documents than can be used in retrieval procedures. Clearly, if the first of these could be ameliorated, and the second exploited, powerful tools will emerge.

Studies of the 'Natural history' of offices

One approach to the psychology of information management is to study the natural history of offices (eg: Cole, 1982; Malone, 1983). The notion is that, by studying the behaviour of people in natural situations, we may come to understand the underlying psychological principles. The small study by Malone is a good example of this. Ten office workers were interviewed about their jobs, their information management habits and their ability to access target documents. The disposition of their offices in terms of layout, distribution of papers, filing methods, etc, were also examined. Although the sample is small, the results are sufficiently suggestive to be worthy of discussion.

A principal distinction that Malone draws is between 'neat' and 'messy' offices. Neat offices are characterised by a structured filing system and reasonably well defined categorisation structures. Messy offices, on the other hand, appear to be less organised, with piles of documents, overlapping papers on the desk, and a generally unstructured layout. We should not be surprised to find out that occupants of neat offices reported fewer difficulties in information retrieval, overlooked fewer things they had to do, and were better able to find specified target documents on request.

Clearly some of this is personal style. Some people are tidier than others. However, this may not be the whole story. There is some evidence to suggest that the nature of the person's job may influence this: people with more proceduralised jobs, such as purchasing agents, tended to have tidier desks. Given the small numbers in the sample, the statistical significance of this remains a matter of future research, but it is not hard to see why it could be true. A highly proceduralised job is likely to result in well defined flows of information and rigid procedures. More flexible jobs (some of Malone's subjects were research scientists) cannot: people in these types of jobs are expected to carry out several tasks at a time and need to be in a position to respond flexibly to new information demands.

A second point of note in Malone's analysis is that the 'mess' in untidy offices is not entirely arbitrary, but reflects a response on the part of the individual to particular needs or difficulties. Three in particular are identified. First, documents are left lying around in conspicuous places as reminders that something has to be done with them. Second, some tasks require that several documents are kept close and 'active' while the task is being carried out. Third, Malone suggests that on many occasions the person is reluctant to file information away either because they cannot decide how to categorise it, or because they are not confident in their ability to retrieve it later. As a result, the compensating strategy is to pile documents around the offices in relatively unstructured files.

If one accepts this view, people do not slip into messy offices purely as a matter of personal slovenliness. There is a mismatch between what the person needs to do, such as act upon several documents at a time, or to make sure not to forget to respond to a particular memo, or to ensure that something does not get filed in such a way as to make it irretrievable in the future, and the facility to do so. All these things create the pressure to spread documents around the office in more or less unstructured piles. The 'technology' (in this case paper-based) does not adequately support these needs, and ultimately the strategy of using piles of information is counter-effective.

Problems in the interpretation of office behaviour

To summarise, there are interesting things we can learn from observing natural information management behaviour. These identify needs such as having reminders and the ability to work off several documents at once. No doubt more detailed systematic studies will reveal more. However, there is also a note of caution required. It might seem straightforward to suppose that we can translate observed strategies of information handling from existing paper-based methods to computers. This is, as we shall see later, the basic philosophy of many leading producers of office automation equipment such as Xerox (Smith *et al.* 1982).

In principle, however, this must be a mistake. Let us consider the use of piles identified in Malone's study. No one would suggest the introduction of unstructured 'piles' of documents in a computer environment. (I say this with the thought that somewhere someone probably has, much in the way that someone thought of building planes that flapped their wings.) Apart from anything else, they are evidently counter-productive. How, then, do we decide which aspects of office behaviour to emulate in office automation and which to avoid?

The problem is that the strategies used by people in one technology need not apply to another. A good analogy is to consider a ball game in which two teams were trying to get a ball into each other's goal. If the ball is small and hard, then the players need sticks to knock the ball around, and the game is hockey. If the ball is larger and made of inflated leather, then the players can use their heads and feet to move the ball around, giving us soccer. Although generically similar, the two games are different because the 'technology' of the game is different. Because of this, the strategies the teams use may be distinctly different and the way players do things in hockey is quite different to the way they do them in soccer, even though, ultimately, their objectives are identical.

The point is subtle but critical to a psychological analysis of information management. The piles that Malone reports are not, in a simple sense, representative of a *need* in the user. Quite the reverse, in fact. They are a compensating strategy for the problems of classification. In using piles, the worker is making a trade-off along several dimensions of difficulty. To avoid the process of classification, which we will discuss in detail below, he puts objects in a particular place. With this he forgoes the opportunity to retrieve the document by any simple classification-based search. But he has other things up his sleeve. First he can remember what it looked like. Consequently, a strategy of scanning the piles (they are, of course, open to inspection) may identify the target by visual recognition. As we will see in a later section, there is every reason to suppose this is a viable strategy. Second, he may remember where it was left, thereby reducing the area for search. Third, the use of piles has an implicit element of ordering by time: the most recent documents are near the top of the piles. This may also be useful information.

Eventually, these strategies are overwhelmed by the sheer volume of papers around the office. However, we can assume that for a while they are actually of some benefit to the user. No formal evidence exists to support this assertion, but any number of anecdotal reports can be found to the effect that people actually prefer spreading their information around in piles. But merely to invoke the concept of *piles* in such systems without an understanding of the role played by recognition memory, or memory for the spatial and temporal attributes of documents, may be to miss the point entirely: the piles may be the visible manifestation of an information handling strategy, but the psychology that underlies them may be much less direct. Only an understanding of this psychology can lead to genuine advances in information retrieval systems.

Concentrating upon piles is to caricature what happens when procedures from office practices are reapplied to computers, but it illustrates the point well. In a later section, discussing particular computer systems, we will return to this point. For the moment, I will leave the reader with the thought that concepts such as in-trays and filing cabinets, artefacts of a paper-based technology, may transfer sensibly to a computer-based system, because they have a genuine function and fulfil a cognitive need. On the other hand they may be the trappings of a constrained and outmoded technology which have no relevance in computer-based systems.

Psychological issues in information management

The purpose of the previous section was to establish the point that the underlying psychology of office behaviour, and in particular in information management, is subtle and complex. A central aspect is that people are making trade-offs between the strategies they can use which shift the emphasis upon the psychological processes that have to be employed. For example, the use of reminders avoids the need to have to remember what to do, or to maintain a checklist of things to do, by relying upon our ability to scan and recognise information relatively easily. The idea of looking at behaviour as a balance of trade-offs is not new, but appears to have particular value in approaches to human-computer interaction. Norman (1983), for example, describes similar trade-off strategies in the use of calculators, and the trade-offs between different methods can be used as an evaluative method (e g, Grudin and Maclean, 1984).

One way of seeing this process is as a shift of resources by the individual to minimise the difficult psychological tasks and to exploit the things we are good at. If we are going to make any sense of this, some exposition of what we are good at and what we find difficult is called for. For the purposes of this paper, I intend to consider the issues identified above; namely, the problems of classification and the role of memory.

Classification

The problem of categorisation has been thrown into the limelight by the introduction of databases such as Prestel, generically known as Videotex. These provide access to information via a sequence of stages in which the user makes a selection of a particular category from a range of options, known as a menu. Initially, the menu has very general topics. At the top level of Prestel, for example, the options are to choose between things such as Education, Travel, News, Agriculture and Insurance. Having made a selection, the menus become increasingly specific. For example, after a few choices under Banking, one can be confronted by choices between things such as the different services offered under a particular bank's unit trust scheme. The idea of this system was that information would be available via the simple process of selecting the desired area of interest from the range of options shown.

However, it very rapidly became evident that all was not well with this method. Users were making too many mistakes by way of selecting the wrong category on the menus, and it was turning out that it was by no means easy to produce sets of options which were not ambiguous (eg, see Dumais and Landaeur, 1983). The first thought was that the designers of these options were not being careful enough. Perhaps, with thought, potentially ambiguous terms could be retitled, or perhaps those areas of the database which were confusable could be more adequately 'signposted'. It was not a success.

What went wrong? Quite simply, information does not fall happily into neat categorisation structures which can then be implemented on a system by using simple labels. Dumais and Landauer (1983) identify two classes of problems which cause this. First, it is impossible to generate category names which will be used unambiguously. This means that there is an inherent likelihood that the categories used to retrieve information will not correspond to those actually used in filing. Secondly, information in the real world falls into several overlapping and fuzzy categories, which means that any categorisation of an item of information can only be relevant to certain aspects of it, even if it can be used accurately.

The problem of accurate use of names arises from the fact that any one name can be used to express several ideas; equally, any one idea might be expressed by several names. As will be seen in the following sections, the human mind is sensitive to the meaning of information, but not necessarily to the details of how it is communicated. The structure of language tolerates this because the context of language disambiguates most confusions. A word like 'table', for example, is perfectly clear in a sentence like: 'He put the plate on the table'. However, the use of words for formal procedures such as categorisation of information means that many of the contextual cues are absent. Put another way, language, which evolved as a form of communication between cavemen is being put to a use for which it was never designed. There seems to me to be no good reason why we should ever expect a formulation of words which will circumvent this problem. At best we can minimise it, but we can have no confidence that merely thinking about the choice of words will produce sufficient accuracy for the purposes of information management.

The problem of overlapping classification of documents is prominent in Malone's study as a prime motivation for placing information into piles. The reasoning goes thus: placing a document into a filing system under one category places the information out of reach if retrieval is required for some other reason. For example, if you file a document about 'medical insurance' under 'medical', then a survey of information relating to 'insurance' is unlikely to produce it. With this clear problem in mind, Malone argues, people do not categorise information because it avoids the difficulty of making a decision between a number of evils, and avoids the consequences of having made it.

Memory

We have already established that memory is used in retrieving personal information, in that we remember filenames (sometimes), when we received information, what it looked like, and perhaps many other things. It seems reasonable to conclude that if we could determine what people could remember and what they could not, then we are in a position to ask (a) how existing methods exploit or fail to exploit memory; and (b) what we can do within the constraints of new technology which would provide the most usable system. Memory is a complex phenomenon and it is not appropriate here to enter into an academic review of the subject. As it is, it will be necessary to get involved in some rather obscure issues. For clarity, I will structure these around three observations:

- the workings of memory are such as to retain the meaning and gist of events, but not necessarily detailed information about them;
- we retain more information about events than we may be able to recall at any one time, and that the ability to recall this information depends critically upon what we are thinking about at the time;
- certain procedures, such as the use of mnemonics, produce apparently supernormal memory performance which may be suggestive of techniques to be used in the future.

Memory for details

It is clear we find it easier to remember the meaning of particular events rather than the details. For example, Sachs (1967) showed sentences that we can accurately recall what they mean, but tend not to recall the particular syntax and construction used to express it. Studies of memory for stories by Bartlett (1932) also show how the recollection of the story is essentially a reconstruction of what the story was about, rather than a memory for what was said.

This is not to say that we cannot remember details, but the ability to do so is not purely a matter of attention and rote learning: it is also a matter of understanding those details in a wider context. The classic example of this is the Chase and Simon (1973) study of the recall of chess positions. Taking positions from real play, chess experts were able to remember and reconstruct the positions of the pieces rather well, and certainly much better than novice players. If, however, the pieces were placed around the board in a random way, their memory was no better than the novices. Their ability to remember the details was in some sense governed by their understanding of the 'meaning' of positions in the real game which the novices did not have. When the pieces were randomised, this prior understanding of the configuration and strategic significance of the piece's positions is lost, and with it the expert's advantage in remembering the details.

Thus it appears that while memory for arbitrarily detailed information can be poor, if we understand the information and can organise it within a wider scheme, then recall can be improved markedly. For the purposes of information management, this implies that the process of filing a document should be a well-defined event in which the choices made fall into a clear pattern of organisation which the person understands well.

Here we can see a strong reason why the organisation in a well-defined, proceduralised job confers an advantage over the flexible, non-proceduralised job. In the former, information management takes place within a strict organisation. In the latter, it may not be possible to provide such a framework. One would also expect to find that the methods of information handling in a tightly procedural job do not change over time, and that the vocabulary used to describe information categories would be better defined. The categorisations used will therefore be less prone to inaccuracy due to synonomy or the fuzziness of filing concepts. This may not be true for the more flexible information manager for whom the organisation of his information needs will change as a function of what he is doing. The worker with the flexible information demands is therefore at a mnemonic disadvantage when it comes to details of how exactly he categorised a document. On the other hand, we will find in the next section that the diversity of information and the changes in context between tasks gives the flexible worker an advantage in remembering more about other aspects of information.

The process of recall

Everyone has had the experience of not being able to recall information they know they have in memory, such as somebody's name, for example. Laboratory studies of this phenomenon have been directed towards showing that this process is not arbitrary: the ability to recall information depends upon a critical relationship between how the information is held in memory and what we are thinking about when we are trying to retrieve it. An experiment by Tulving and Thomson (1973) demonstrates this well. Subjects were shown words that had to be remembered, such as *jam*, in the context of another word, such as 'traffic'. Later, subjects were given words, such as *marmalade* and asked to produce words closely related to them, of which *jam* might be one. These words were chosen so as to be likely to produce as responses words which were in the original to-be-remembered list. Later, the subjects were asked to recognise which, if any, of the words they had generated were words they had been originally asked to remember. Finally, they were given a test in which they were given the words which appeared with the target words as prompts for recall – eg, traffic –? What happens in this experiment is that words such as *jam* may not be recognised in the context of the closely related words such as *marmalade*. Yet some can still be recalled, nominally a much harder mnemonic process, when given the original context word 'traffic'.

The interpretation of this experiment lies in the postulation of two related processes. First, when information is committed to memory it undergoes an encoding process in which the information is interpreted. This process is what takes place when expert chess players are assimilating details about the position of the pieces, for example. This encoding process seems to be determined by the context in which the information is committed to memory. Thus the word 'traffic' places a particular interpretation upon what is an ambiguous word. Secondly, at the time of recall, the context of what we are thinking about provides a prompt for memory. If this context matches the way in which the information was initially interpreted, then recall is successful, but if it does not, then recall fails. Hence the word *jam* is not recognised in the context of *marmalade* because it does not match the meaning of what is held in memory. 'Traffic' as a cue, however, does, and recall succeeds. In technical terms, therefore, human memory is said to be content-addressable. For memory to succeed, the meaning of the retrieval cue must match what is stored in memory.

Theorists of human memory may argue with the sophistication of this model, but it captures the essence of many memory phenomena. For example, it explains why we fail to recognise the butcher in the bus queue: we are only used to recognising the butcher in his whites behind the counter. It also indicates why our memory for detail is so much better if placed in the context of a wider scheme of things, as in the example of the chess players. Presumably the encoding of the information in a coherent framework makes it more likely that future prompts for recall will be in a form which is already represented in memory.

When discussing information management we will see that this property of memory is particularly relevant to how we can use existing systems and what we might exploit in future ones. The process of information retrieval in the human mind is fundamentally different from filing or library systems in which items are accessed by location rather than by their meaning. As it stands, therefore, there is good reason to believe the informal observations that what is remembered about documents is the meaning of their content and contextual information such as what they looked like, what one was doing at the time, and so on. As we shall see, computer systems do not use this information but rely upon the user remembering filenames and/or the categorisation of required information. What we are good at is being ignored, and what we are required to do is a difficult and flawed psychological process.

Exceptional memory

There are circumstances in which memory appears to be especially effective. One of these is the use of mnemonic techniques of imagery (imagining pictures in the mind's eye). One mnemonic technique, for example, known as the method of locii, involves learning a list of objects by mentally moving along a well known route 'leaving' objects at strategic places. Recall is then a process of going over this path again 'reading-off' the objects as they are encountered. This was the method taught by the ancient Greeks as a form of mental discipline. Another mnemonic method involves learning a rhyme of the form 'one for a bun, two for a shoe . . . ' etc and remembering objects by forming mental images in which the object is seen interacting with the appropriate object. Thus, if the first object was a telephone box, one might imagine a large bun with a telephone bow on top. This is known as the 'peg' technique (eg, see Baddeley, 1976).

High levels of performance are also seen in some experiments looking at memory for pictures. One type of experiment in particular has served to establish a reputation for picture memory. Subjects are shown a large number of photographs of various scenes and are subsequently asked to identify them among other, new pictures. For example, Standing *et al* (1970) showed subjects over 2500 pictures over the course of two days. Subsequently, they were shown two pictures, one of which had been seen, and one which had not. Accuracy in recognising the original pictures remained at very high levels even over considerable periods of time.

The common theme in these lines of research is the use of visual information. It ties in well with anecdotal remarks about people's ability to remember incidental features about documents such as their typeface, where they were located, dog-eared pages, or the coffee stains on the front page. This has led to many researchers recommending the use of computer interfaces which rely upon visual and spatial features to categorise information. For example:

"The human brain processes visual information and language quite differently... While there are many unanswered questions about how visual information is recalled, studies have shown that most people can recognise previously viewed images with almost perfect accuracy. The speed of image processing and the accuracy of image recognition are two factors on which an iconic-based man-machine interface can capitalise. (Lodding, 1983).

Taken in the extreme, this view probably represents a significant lack of understanding of what these experiments mean. Earlier, I discussed the dangers in assuming that trends in office behaviour, such as the use of piles, were, *de facto*, useful strategies to apply to computer-based information management. The crux of the argument was that the overt behaviour and performance (ie, the success of the strategy) were the result of a trade-off between the constraints of the technology and the underlying psychological processes. Consequently, if the technology were to change, and with it the constraints, then the match between the psychological processes and the technology would result in different trade-offs. Consequently, strategies would change also, and what constituted a reasonable strategy in one technology need not in another.

A similar argument applies to the interpretation of psychological experiments. Just because particular experiments on the memory for pictures appears to produce high levels of recall, it does not follow that tasks involving pictorial memory will always do so. It could be that the particular circumstances of the experiment allowed for successful strategies which may not be applicable elsewhere.

There are, in fact, a number of reasons why experiments on picture memory might produce high levels of performance without resorting to the idea that picture memory is inherently exceptional or involves distinct psychological mechanisms (eg, Paivio, 1971). For one thing, the recognition test requires of the experimental subjects only that they remember enough to *differentiate* between the target picture and the distractor. Given that these experiments use pictures of a variety of complex scenes, this need not be very much information at all. Also, while pictures of harbour scenes or landscapes may appear to contain considerable detail, compared to most memory experiments this detail is encapsulated within a meaningful context. In a recognition test, the representation of the picture provides an ideal prompt for memory in that this context is re-established. We therefore have the ideal combination of encoding and retrieval circumstances being employed upon a task which requires minimal recall.

The performance of memory when using imagery can be similarly explained without resort to the notion that visio-spatial memory is exceptional. In the method of locii, for example, it could be argued that the process of imagining an object in a particular place and then mentally going back to that place to 'find' it is one in which there is a close relationship between the to-be-remembered object and the prompt being used to recall it: the place where it was put.

Experiments on visio-spatial filing methods

Consideration of the theoretical processes in imagery and memory for pictures, therefore, leads to a clear debate between two points of view. One says that memory for pictures, or in general visio-spatial information, is exceptional, presumably because it relies upon distinct psychological processes. This is sometimes known as the 'dual code' hypothesis, since it relies on the postulation of different methods of coding visual and verbal imformation in the mind. The other view is that, for the experiments described, the general mnemonic processes of encoding and retrieval are precisely the same ones used for other kinds of memory. In this case, however, they are being used to particular effect. The practical difference is this. If the former point of view is correct, then information management tools using visio-spatial information will automatically be accompanied with better recall. If the latter view is correct, then recall will depend upon how well the prompts for recall match the way the information was encoded, and will not be related to whether the information to be remembered was visual, verbal or whatever.

This provides the starting point for a series of experiments carried out by myself and some students recently (Lansdale, 1985a; Stroud, 1986; Simpson, 1986) Subjects were asked to imagine that they were working in an employment agency. First of all they were required to file a number of job adverts which may subsequently be needed when a 'client' came in looking for a job. They could do this either by using a visio-spatial method of filing, or a more traditional method. In the traditional method, subjects assigned three labels to the document chosen from three sets of 12 labels, the first being the descriptors of the possible job categories, the second being a set of colour names, and the third being a range of adjectives. Thus an advert might be filed under 'fashion red-high', for example. The visio-spatial method was closely comparable: subjects first specified which of a number of shapes (chosen to be relevant to the advert types) was to be chosen, it was then coloured in one of 12 possible colours, and finally it was placed in one of 12 distinct locations on the advert itself; rather in the manner of a coloured watermark. Subsequently, after a delay, the subjects were asked to remember how they filed particular adverts, which meant recreating either the three labels in the verbal method or the shape, colour and location chosen in the visio-spatial method.

While this situation may seem somewhat artificial, it is designed to address the question of whether or not there are inherent benefits from using visio-spatial information. The two methods being used were arranged so as to be entirely comparable in the *amount* of information to be remembered: all that differed is whether the to-be-remembered information is visually or verbally mediated. If the dual-code hypothesis is to be a useful theory for the design of information management tools, then we would expect that the subjects' performance with the visio-spatial method would be better, on balance, than with the verbal system. If the hypothesis does not hold up, then we would expect that the subjects' performance would be determined by their ability to establish a meaningful link between the advert they are being asked to remember (which is used as a retrieval prompt), and the labels or visio-spatial attributes they associated to it.

It turns out that the subjects do not perform better with visio-spatial information. Indeed, being asked to remember an arbitrary location seems particularly difficult, confirming a similar finding by Dumais and Jones (1985). Further, there is evidence that the strength of the subjects' memory is indeed dependent upon their ability to establish a strong associative link between the chosen attributes and the content of the advert, as might be expected in a unitary view of memory. Thus the subjects showed particularly good recall for the names of the job categories or their equivalent shapes. This can hardly be said to resolve the dual-code *vs* unitary coding controversy. That battle will continue in the journals of theoretical psychology (eg, see Kosslyn, 1981; Pylyshyn, 1981). What it does confirm, in my mind at least, is that assertions made about the inherent value of visio-spatial information represent a simplistic view of human cognition and no guarantee of good design.

As a final point in this issue, it is important to emphasise that this does not mean that visio-spatial information should be dismissed in considering future information systems. The reason why Standing *et al's* (1970) subjects recognise pictures well may be controversial, but the fact of it is not. It can well be imagined that future systems, among other things, will provide a means of rapidly scanning several documents with a view to recognising the wanted item. This is probably a key element in existing office behaviour, and technologies such as videodisc, which can store and rapidly manipulate pictures of documents, offer precisely this facility. What *is* being dismissed is that visio-spatial methods are automatically a panacea for human-computer interaction.

Let me summarise this review of the cognitive psychology of information management. We can see that humans are not good at categorising information. We can also see that when committing information to memory, they do so not by categorising that information in the way a librarian would place a book on the shelves, but by interpreting that information in the context in which it appears. We then find that the ability to retrieve this information depends upon being given information, or thinking about things, which are directly related to that specific interpretation. Information which is logically related to the required memory will not succeed in eliciting recall unless it is also related to the way in which that information was interpreted. This model gives some insight to the conditions under which we can remember details and when we cannot. Finally, we can look at circumstances in which people seem to remember rather well. Here we find that the success of recall is more probably related to the particular circumstances under which they are being asked to perform than resorting to concepts of supernormal psychological powers. Some experimental set-ups seem to provide a set of circumstances under which the general principles described above seem to work efficiently.

We now want to transform this understanding, such as it is, to more practical matters. What do we know which relates to the design of personal information management systems? As an approach to this question, it is useful to evaluate the current state of affairs. This provides some context for assessing the existing problems and also serves to highlight where these problems actually occur. The next section, therefore, reviews existing and proposed futuristic information management systems in the light of what we have discussed.

Computer-based information management methods

Direct access

The oldest and by far the most common method of accessing files is by direct access. That is, retrieval is achieved by giving the system an exact identifier of the information required. This may be a file name, a number or a location. Its difficulties of use are also the most obvious. First, the user has to know exactly what information is required, and secondly has to know exactly how to identify it to the system. If either or both of these are misssing, then the only alternative is to undertake a search, which may be very time consuming. We have already seen that people find the memorising of filenames difficult and are reluctant to categorise items under a single label. Carroll (1982) has shown that a common response to this is to use compound filenames which actually incorporate abbreviations of several concepts. The name for this particular document on my word processing system, for example, is INFOMAN2, standing for 'Information Management Paper, second section'. Carroll is also able to show, for reasons that are becoming repetitive, that in relatively short periods of time, people's filenaming patterns produce ambiguous and inconsistent patterns which lead to retrieval difficulties.

Another obvious thing to note about such systems is that they are very restrictive. With the system with which I am writing I can only have one document upon the screen at once, and to get it I have to remember its exact filename. (In fact, what I do is to look at the directory listing of filenames and guess which is the most likely looking name, sometimes having several tries before the file that opens is the one I want.) The situation is rather like having a filing cabinet in which all your information is stored, and your desk and the office you are in is empty – not a single document is visible. When you want something, you give a secretary a name. She does not ask you any questions, she merely goes and tries to find the thing you want. If she finds it, it drops on your desk; if she cannot find it, she tells you. She does not tell you if there is another file with a similar name. When you have finished with the document, it is taken away and replaced by another one.

It all paints rather a bleak picture. If a secretary put such a barrier between you and your filing system, it would be regarded as hostile behaviour. But until recently, this is exactly what potential purchasers of office automation were being asked to buy, without exception. Indeed, the concept of the 'paperless office' was built around such facilities. It is not surprising that the introduction of OA, and the move towards the elimination of paper as an information medium, has been far slower than might have been predicted. Reminding ourselves of the kinds of things office workers actually do with information, such as keeping several documents 'active', reminders, and exploiting the flexibility of piles around the office, we can see that these systems are not sufficiently flexible. It simply is not realistic to expect that people will rely upon such a restricted domain of procedures to carry out the full range of their information management needs. One might also note, as indirect confirmation of this point, that OA has had most success in those proceduralised tasks such as accounting or invoice management.

Some attempt to alleviate these problems and add functionality to information management has been provided by relational databases. These allow the use of several keywords per document and the ability to search by using combinations of them. However, automated search processes around keywords suffer from the same problem that causes difficulties of categorisation: words have several meanings and also several synonyms. To find particular documents by relational keyword search means either that enough keywords are used to ensure that the necessary items are retrieved (in which case many of the items identified are irrelevant and the user has to sift through the list of retrieved items), or it has to be expected that, with fewer keywords, some relevant documents may not be accessed. This is the 'recall versus precision' dilemma (see Christie, 1985).

However, we have still not exhausted the difficulties with such systems. As a working environment, these systems present information in a uniformly undifferentiated way: all that changes between one document and another is the filename and its content. In terms of visual appearance, every document looks like every other one: simply text on a screen. Information about documents such as their length, structure or relation to other documents, features which are often implicit in paper-based documents' appearance, are either missing or at best hard to get at. A common complaint of traditional OA equipment is, therefore, that it is impoverished of cues: the extrinsic attributes of information which add to their diversity and our ability to differentiate them in memory. This has led to the notion of 'cue enrichment' (eg, see Cole, 1982). By this is meant the reintroduction of attributes which put back information into the interface which can be used to help manage the information. The next section discusses one approach to this.

'Desk-top' computers

Since the introduction of the Xerox 'Star' system, and Apple's Lisa and Macintosh which followed soon after, a strong fashion has developed for designing computer interfaces around the 'desk-top' metaphor. In such interfaces, the user is confronted with a display of windows and graphic images, known as icons. These icons represent familiar office objects such as documents, waste bins, folders, in-trays, filing cabinets and so on, whose function is largely the same as their real counterparts. Actions with these objects are, therefore, very literal Printing may be executed by moving a document icon to a printer icon. Information is destroyed by moving it to the trash can. As far as is possible, the functionality of the interface is derived from the desk-top metaphor (eg, see Smith *et al*, 1982). It is important to recognise that the design philosophy of these machines is not principally to support information management, but to make the machines easy to use. In this respect, there is no doubt that this style of computer design is distinctly more successful in this respect than its predecessor.

Their facility in certain aspects of information management is clearly helpful. By providing the ability to show more than one document at once (they are represented as overlapping 'windows rather analagous to sheets of paper), and by providing the means to keep other information closely available in 'notepads' or 'scrapbooks' or in unopened icons, these clearly go some way towards supporting tasks requiring more than one document. It is relatively simple to scan around the 'desk-top' and pick up and put down documents as required.

However, in areas of information management which involve longer term storage and retrieval, their facilities do not provide added value. When it comes to filing, these systems resort to rather traditional methods: the files need a name, they can be placed in a folder, and in a particular 'filing cabinet' or on a particular disc. Ultimately, to retrieve this information, we are back to remembering filenames and the categorisations we used to file the information. The added visual aspects of the interface do not provide any help here: icons of a particular type (such as documents) all have exactly the same appearance. The only way they are differentiated is by a filename underneath them. Thus these interfaces are functionally identical in this respect to the traditional methods described in the last section. The added value of these 'revolutionary' devices is entirely invested in the ability to interact with the documents while they are on the desktop.

Spatial data management

A rather futuristic system which attempts to get away from traditional filenaming and categorical systems is the Spatial Data Management system developed at MIT (Bolt, 1979). In this system, information is not stored in the normal way, but as an interactive, three-dimensional image, known as the 'Dataworld' which is shown on a wall-sized display. At the top level, when the user is 'looking down' on the dataworld, the items of information are represented as crude graphical miniaturisations of themselves, or icons. To find and interact with the data, the user 'zooms in' upon the area of the dataworld he is interested in. This area then becomes magnified to show increasing detailed information as the user gets 'closer' to it. The philosophy of this system is very clear. SDMS is built around the notion that people navigate around their office information systems by using crude locational and visual cues which focus their search upon the relevant areas of their filing space. In this view, documents are recognised by their visual appearance and the only aspect of categorisation is invested in where documents are placed in the dataworld, and next to what.

We have seen that there is plenty of reason to think people can and do use visual appearance as an aid to information scanning and retrieval in offices. We also know that recognition memory for complex visual stimuli is very good. Clearly, the ability to search visually a range of documents could be an important facility provided by say, videodisc. But is a system based *purely* on the SDMS principle viable? No evaluation has been carried out, but I am confident that it is not. And circumstantial evidence as to why not is already available. First, looking at Malone's messy offices, we can speculate that the build-up of uncategorised information in piles eventually overloads the ability to locate information by memory for its location and visual scan. Every now and again a 'tidy up' is essential. Looking at the dataworld, it is hard to believe that the same will not be true for it also. Further, it is questionable whether there is enough space (or could be) to physically hold the volume of information the average office worker maintains. Jarrett (1982), for example, estimates this at 20 000 pages and rising at 2000 per annum.

The seond type of evidence, which is far more damaging to the philosophy, questions the ability of users to remember arbitrary locations (Dumais and Jones, 1985; Lansdale, 1985b; Simpson, 1986; Stroud, 1986). But the dataworld is essentially a two-dimensional space for which there is no implicit organisation or meaning. I see no reason why, on that basis alone, users should expect to find it easy to remember where they left information. Suppose we argue that users can impose meaning and organisation on to this arbitrary space? There is some evidence from experiments I have carried out that this is so. One of the strategies adopted by subjects in trying to remember locations of a coloured object on a document, when that location was originally under their control, was to use the locations to categorise aspects of the documents. For example, in filing adverts on education, some subjects would consistently use a particular location, such as the top left-hand corner, or would restrict their choice of locations to, say, the right-hand side. In effect, this strategy is one of mapping categorisations of documents on to locations on the screen. We could therefore imagine that users of SDMS could segment the dataworld in such a way as to give meaning to where items were placed. However, in doing so, the problems of categorisation are reintroduced: what happens when a document could be in more than one place is precisely the same problem as when a document could be filed under several keywords. We should expect, therefore, that one of the problems of SDMS will be that users will find themselves 'hunting' around the dataworld looking for clues as to where they left particular documents.

I have no doubts that using dataworld provides the user with a basis for developing some sophisticated strategies which exploit aspects of cognition such as speed of scan and accuracy of visual recognition. However, if we look at behaviour in general, and office behaviour in particular, it is characteristic of humans that they will exploit *any* strategy which is of some use. Dataworld is restrictive in this sense. It was intended only to illustrate the power of visio-spatial navigation, but in doing so (if indeed it does) it has not left it at all clear how all the other strategies can be integrated with it. As we have seen, it is not even clear whether the strategies it will support will be all that effective. In the final sections I will address the problem of how an information management system will be built which exploits the things we are good at while at the same time does not restrict the user to a set of strategies.

Looking to the future

I have spent some time talking about the need to design information management systems to match the needs of the people who use them. I have also discussed, at length, aspects of the psychology of those users. Finally, I have tried to show how approaches to design of information systems have fallen short in some way or another. But how do we proceed? How do we design for the future? Ultimately, what we need is to turn this descriptive information into a specification for future systems.

The problem with this is that it is rather like being asked to imagine a colour one has never seen. We have the mental apparatus to see things in terms of how they are, or how they have been, but not how they will be. If the reader feels the need for confirmation of this remark, let him look at science fiction from earlier decades, such as from Chaplin's *Modern Times* or Jules Verne. We can see quite clearly that their vision of the future is quite firmly rooted in their grasp of the existing technology. The knobs, dials and wheels are bigger, the rockets are more powerful and the cars are smoother, but they are not *different*. How could they possibly have imagined the future we live today?

Another general remark worth making is that the history of technological development is littered with products which failed not for economic reasons, but because nobody really wanted the functionality they provided. I cite as examples of this quadrophonic hi-fi, 3-d cinema, and radio controlled cat-flaps. The patents offices are full of thousands more potty ideas. They occur because the technological capability is outstripping a good understanding of what people really want. Information technology is rapidly becoming a ripe area for hunting these flops for precisely this reason.

How then do we proceed to something different and more functional? It must by by an iterative process rather than purely by the adoption of some new scheme or technology. It is true that new possibilities of the technology will be recognised: some of these, by natural selection, will succeed and some will not. These will lead to methods which the user will have available for all aspects of his job, including information management. From the other side, clear user needs (such as reminders, for example) will emerge from behavioural studies of office activities for which suitable interface devices will be envisaged. Again, some of these will succeed and some will not. Thus far, this is how the development has been driven: by technological development and by the automation of existing strategies.

One of the main points of this paper is that developing systems in this way is inefficient, because there is no strong reason to believe that they will produce a good solution to the problem of adequately automating our information management needs. They certainly have not done so yet. Also, this paper has been directed to making the point that the underlying psychology of information management cannot be directly inferred from users' behaviour in offices, because that behaviour is largely adapted to overcoming the problems being created by the mismatch between the facilities provided, the users' need, and their cognitive capacities.

Psychological study of office behaviour can therefore contribute to the development of information management tools much more directly. By specifying the strategies we are good at, and by identifying areas where we are weak, it should be possible to provide software devices to assist the user. In the former case, such devices are 'cognitive enhancers' and in the latter 'cognitive prostheses'. In general, we can think of them as information management tools, implying a range of different methods for different needs, which the interface will provide for the user. I will call these Cognitive Interface Tools (CITs) because their principal purpose is to support the cognitive processes of information management. Without testing and development, one cannot tell what tools people will want or what they will look like. The final section is devoted to research issues designed to tackle this. However, to give some flesh to these somewhat abstract remarks, it might be useful to give from my own experiments upon information retrieval systems (Lansdale, 1985a) two results which demonstrate the value of such a tool.

The first of these was based upon the notion that people remember chronological information about information: what else was happening at roughly the same time. As in the more recent experiments I have already described, users were asked to remember visual attributes associated with a document. Retrieval was achieved by telling the system which shape and colour, etc, appeared with the required document. However, in one particular experiment (Lansdale, 1985a, Experiment 2), as an aid to retrieval users were presented a graphical display of *when* documents associated with the visual attributes they had specified were received. The success of retrieval with this chronological information was considerably better than without, confirming the hypothesis that the test subjects had information about the sequence of events which could be used to effect. It seems likely that the ability to remember the chronological significance of documents is of general value. It also seems likely that the subjects were able to use the configural information this provided. For example, if the user remembered that the document he wanted had *red* and *building* associated with it, then he could concentrate on those areas of the database where these two were present together. It is clear from the subject's performance with this graphical display that this was a significant way for him to limit his search. Therefore, an interface tool which makes this information available in an easy-to-use form could well have considerable utility.

The second significant result arises from looking at how people remember the attributes of colour, shape and location, or in the verbal systems, the three labels associated with documents. It turns out that these associations between the document and the three attributes assigned are *independent*. That is to say, the likelihood of remembering one of them is not affected by whether anything else has been remembered. For reasons too detailed to go into here, this means that most of what people remember about a particular document is *partial*: they remember some of what was associated, but not all. Systems which require complete accuracy (memory for all three elements) therefore fail. By providing a system in which users can express their partial knowledge, computers can specify only those documents which correspond to it, thereby providing a much smaller subset of information to be searched.

The previous example of the graphical/chronological display succeeds exactly in this respect because the user can focus upon those areas in the sequence which correspond to where the partial recall he has is observed. In the example above, for instance, the user can concentrate upon documents showing that they were associated with *red* and *building*.

The purpose of this section has been to illustrate that the tools a future information management system might use will probably be quite unlike anything we might have imagined. We have also seen that in these circumstances it is very hard to iterate towards a design without some kind of guideline or framework to indicate the way to go. The final section discusses a framework for research into information retrieval which is aimed at providing this.

Conclusions: A framework for future research

Throughout the paper I have resisted the temptation to produce a model of information management. It seems to me that one of the characteristics of behaviour in this area is its flexibility and the relevance of several theoretical issues to it. Also, information management is not a single process, but a collation of several. To model such an amorphous area is therefore in danger of crystallising and simplifying the subject to the point where anything useful or insightful is squeezed out by the need to produce the model at all. I would venture that this is rather a common failing in psychological research. However, having presented the material in a relatively unstructured form, and, I hope, having established what a complex and subtle thing is behaviour in this area, I feel that presenting a framework may be possible without misleading the reader (or myself) into thinking that it or its assumptions has any value beyond helping us to grasp the range of issues and focus upon illuminating and cost-effective lines of future research.

The major simplification I intend to make in this respect is to differentiate between information *handling* and information *retrieval*, and to concentrate upon the latter. By information handling I refer to all those processes which relate to what people do when they have got information, such as handling several documents at a time on the 'desktop', or leaving some around as reminders, and so on. This behaviour is largely idiosyncratic, and is in any case becoming better supported by modern interfacing techniques. Future research should concentrate upon how well these interfaces support this behaviour and exactly what the microstructure of that behaviour is, so that it may be even better supported. This will be crucial to the acceptability of future electronic offices. However, the problem of information storage and retrieval represents in my mind a far more general, and therefore potentially profitable, problem. We have seen that the modern iconic interfaces have nothing to offer in this direction, and novel approaches such as SDMS cannot sidestep the problems of the categorisation of information merely by slipping into a different and revolutionary style of man-machine interaction. Spreading out to other systems, such as that proposed by Spence and Apperley (1982), we see a similar picture: some ingenious devices to approach the problems of information handling, but basically straightforward classification methods for information storage and retrieval.

The framework I propose for the process of information retrieval is very simple: it is that *every* attempt at retrieving information involves two distinct psychological processes: recall-directed search, followed by recognition-based scanning. Recall-directed search refers to the use of memory about the required item to get as close to the document as possible. At its most exact, this will give direct access, such as remembering a filename or in which sleeve of a filing cabinet drawer it is located. As it becomes less accurate, memory might identify an area of the database: a drawer of a filing cabinet, a shelf, or perhaps a particular computer directory. In this case, more or less recognition-based scanning within that area is required depending upon the exactitude of recall. Recognition-based scanning is therefore the process we undertake when recall has failed to produce a unique item, and the amount we have to undertake depends upon how specific the recall was.

I propose that the observed behaviour with any information storage and retrieval system can be seen as a trade-off between these two processes of recall and recognition. The users can be said to be balancing the two processes to maintain, in their eyes, reasonable utility. In other words, the strategies people adopt can be seen as a way of shifting this trade-off to their advantage. For example, by leaving documents around an office in piles, people are using their recall for where it was placed, when it was placed there, and what it looked like. The scanning process can therefore be reasonably efficient, at least until the office becomes clogged with papers, because recall can not only specify where to look, but also visual aspects of the target which help in recognising the target.

This is also a way of looking at the usability of the traditional 'direct access' system (see above), and why it fails to provide an adequate solution to the storage and retrieval needs of the electronic office. The recall process is deficient because, as we have seen, the users' ability to categorise documents with the appropriate filenames, and their ability to remember those filenames, is limited. On the other hand, the user cannot fall back upon scanning methods: it is enormously clumsy to search directories and open files one at a time to inspect their contents. Having done so, they can only be differentiated by reading them. The user simply has nowhere to go, no way of using the machine which will provide an acceptable trade-off between the problems of recall and recognition.

Thinking in terms of this trade-off also suggests a general method of approaching the design of an information management system. A future system should be as near to an optimisation of the two processes as possible: recall processes should allow the users to use whatever memory they have to limit the area of the database to be searched; and then the information within this area of the database should be represented in such a way as to maximally assist the search process. Optimising both processes (as far as is possible) should give any users enough room to manoeuvre such that they can accommodate their special needs and idiosyncracies. This leads me to a number of specific research issues.

Multiple keywords

We have already noted that there is a fundamental problem in categorisation. However, we might imagine that the use of multiple categorisations and systems that are sensitive to synonyms will considerably improve matters, at least in that it does not provide a difficult decision to the user as to which categorisation to use. However, it might be expected that increasing the number of keywords or categorisations of data, and permitting access by synonyms, will produce a great deal of unwanted recall, putting greater stress upon the recognition process. Research should therefore be directed towards increasing precision in this respect.

Recent research claims to have achieved some success in this direction already. Gomez and Lochbaum (1984) have demonstrated a system of multiple keywords which actually *reduced* the number of unwanted retrievals in a simulated information system. However, we need to be a little careful in the interpretation of their results: the subjects were provided with an intermediate feedback in the system in which they were shown, for each keyword entered, an array showing which documents were associated to that keyword. With several keywords giving such feedback, the subject is at liberty to use the information about the *conjunction* of the keywords which is far more informative as to the likely target than anything else. The situation is, in fact, very similar to the use of a chronological display in my own experiments described in the previous section. Gomez and Lochbaum may therefore be ascribing functionality to the keyword system which is actually attributable to the subject's ability to use their method of feedback to good effect.

A person's ability to use multiple sources of information to focus his or her searching strategy to a smaller number of potential targets seems particularly sophisticated (eg, Lansdale, 1985b). This suggests that the inaccuracies of keyword systems, which might be expected to be magnified with multiple categorisations, can be ameliorated by the use of suitable interface devices (in my terminology, CITs) which allow the user to use the configural information inherent in them. Exactly how such devices will appear is a matter for future research.

How easy should filing be?

One aspect of information retrieval which is important to emphasise is that although difficulties may be apparent at the retrieval stage, in that one is not able to find something, it is as sensible to see this as a problem of *storage* as of retrieval. We cannot find information because we did not file it in a form which was going to be appropriate for future recall, or perhaps because we did not pay enough attention to how we classified it at the time. It is a natural reaction for people not to spend a great deal of time filing information because it has no immediate pay-back and because they want to get on with the next piece of work (eg. Cole, 1982). This is at least one negative reason why users resort to unstructured piles: it gets the user quickly out of the problem of how to file documents.

However, the success of information retrieval depends upon the user doing the right things at the time of the storage, so that the information is filed sensibly and because his memory for the document, and how he filed it, will be stronger. For example, Lansdale (1985a) and Simpson (1986) show that users assigning retrieval tags to documents (pictures and words) have much stronger memory for those tags if they assign them themselves as opposed to having them assigned by the system. This ties in well with well-established educational principles that memory is much more robust for self-generated material than when it has been provided for the subject (eg, Bower, 1970).

This leaves us with an important dilemma. The more we ask the user to do at the process of information storage, the less likely he is to do it, creating retrieval problems. On the other hand, the more we automate the process of storage and take responsibility away from the user, the less he is going to remember, and therefore the less he is going to be able to retrieve. This leads to two classes of research issues which follow. First, we can ask how we might facilitate necessary tasks so that they are not so onerous that users will not carry them out. This is a matter of finding out what it is we want the user to do at the stage of storage and considering how the software could be devised to do this most easily. Alternatively, we can resign ourselves to automated systems in which the user's memory is intrinsically poorer. In this case we are interested to ask questions as to what, if anything, is best remembered under these conditions and how we can best exploit what is remembered. The key point here is that any future research must be sensitive to the trade-off between the positive aspects of easy filing facilities and the negative effects they have upon the user's memory.

What attributes should electronic documents have?

The final set of research issues falls into the category of asking what it is that people can remember about electronic documents, above and beyond the categorisations we may ascribe. This is, of course, rather like asking the length of a piece of string, because what they remember is to some extent dependent upon what is provided: if documents are all monochromatic, there is no sense in asking whether we remember their colour.

The first question for research in this area is therefore this: what attributes do we ascribe to information for the purposes of managing information retrieval? The list of candidates is fairly varied: we can use names or compound names; any number of physical attributes such as length, format, colour or typeface; or we can try experimental methods such as icons or the visual formats of SDMS. I am sure there are more. The whole philosophy of 'cue enrichment' (eg, see Cole, 1982) is one of using these dimensions to create a sufficiently rich and varied electronic environment to provide more recallable and recognisable aspects of documents.

Different attributes will have different functions. For example, if one were to use the size of a document as a coding attribute, it is not likely to be as useful for recall as it is for recognition. Our relatively imprecise ability to remember a quantity such as size, which can vary between a wide range of possibilities, means that it would not make a particularly precise recall cue. On the other hand, as a way of rapidly differentiating between a number of potential targets during a recognition-scanning process, a representation of size might prove very useful. It may therefore be that consideration of what attributes to confer upon documents may depend upon whether one is contemplating supporting the recall or recognition-scanning process. Consequently, we might also expect that some aspects of electronic documents will be there as attributes designed to support one process, and some the other.

The second question to ask about these attributes, having decided upon their individual utility, is to ask how they perform in a real system which employs several. Which, if any, are more likely to be remembered, and under what conditions? Suppose it is true, for example (and I have no particular reason for thinking so), that particular attributes of documents, such as their typeface, are remembered to a certain level almost automatically. Then clearly this information could be used in preference to other attributes which may be remembered rather better, but only under conditions where the person involved has to think quite hard as to how to remember it. Which is adopted depends upon how the user of an eventual system sees the trade-off between ease of encoding and success of retrieval

Preliminary research on the memorability of attributes such as shape, colour and location has already begun with interesting results, as already described (Lansdale, 1985a; Dumais and Jones, 1983). But considerably more work is required before we are in a position to specify to designers which attributes should be used and how, to support all-round information management performance.

References

Baddeley, A.D. 1976, The psychology of memory. Harper and Row, London, UK.

Bartlett, F.C. 1932, Remembering. Cambridge University Press, Cambridge, UK.

Bolt, R.A. 1979, Spatial data management. MIT Press, Boston Library of Congress No 78–78256.

Bower, G.H. 1970, Analysis of a mnemonic device. *Amer Scientist*, 58, 496–510.

Carroll, J.M. 1982, Learning, using and designing filenames and command paradigms. *Behaviour & Info Tech*, 1 (4), 327–346.

Chase, W.G., and Simon, H.A. 1973, Perception in chess. *Cognit Psychol*, 4, 55–81.

Christie, B. 1985, Shared information systems. In: Christie (Ed) Human factors of information technology in the office. Wiley, Chichester, UK.

Cole, I. 1982, Human aspects of office filing: implications for the electric office. Proc Human Factors Soc, 26th Ann Meeting, Seattle, USA.

Dumais, S.T., and Landauer, T.K. 1983, Using examples to describe categories. Proc CHI '83, Boston, USA.

Dumais, S.T., and Jones, W.P. 1985, A comparison of symbolic and spatial filing. Proc CHI '85, San Francisco, USA.

Gomez, L.M., and Lochbaum, C.C. 1984, People can retrieve more objects with enriched keyword vocabularies. But is there a human performance cost? Proc Interact '84, 1.429–1.433. London, UK.

Grudin, J., and Maclean, A. 1984, Adapting a psychophysical method to measure performance and preference trade-offs in human-computer interaction. Proc Interact '84, London, UK.

Jarrett, D. 1982, The electronic office: A management guide to the office of the future. Gower Publishing Co, Aldershot, UK.

Kosslyn, S.M. 1981, The medium and the message in mental imagery: a theory. *Psychol Rev*, **88**, 46–65.

Lansdale, M.W. 1985a, Psychological issues in the use of icons in information retrieval. Manuscript prepared for publication.

Lansdale, M.W. 1985b, Strategy and structure in cognition. Paper presented to ESRC International Workshop on Modelling Cognition, Lancaster, UK.

Lodding, K.N. 1983, Iconic interfacing, Computer graphics and applications.

Malone, T.W. 1983, How do people organise their desks? Implications for the design of office information systems. *ACM Trans Office Info Systems*, 1(1), 99–112.

Norman, D.A. 1983, Some observations on mental models. In: Gentner and Stevens(Eds) Mental models. Erlbaum, Hillsdale, New Jersey, USA.

Paivio, A. 1971, Imagery and verbal processes. Holt, Rinehart and Winston, New York, USA.

Pylyshyn, Z.W. 1981, The imagery debate: Analogue media versus tacit knowledge. *Psychol Rev*, **87**, 16–45.

Sachs, J.S. 1967, Recognition memory for syntactic and semantic aspects of connected discourse. *Perception & Psychophys*, **2**. 437–442.

Simpson, M. 1986, Unpublished final year project, Department of Human Sciences, Loughborough University, UK.

Smith, D.C., Irby, C., Kimball, R. and Verplank, W. 1982, Designing the Star User interface. *BYTE*, April, 242–282.

Spence, R. and Apperley, M. 1982, Data base navigation: an office environment for the professional. *Behaviour & Info Tech*, 1(1), 43–54.

Standing, L., Conezio, J. and Haber, R.N. 1970, Perception and memory for pictures: single-trial learning of 2560 visual stimuli. *Psychonomic Sci*, **19**, 73–74.

Stroud, T.R.J. 1986, Unpublished final year project, Department of Human Sciences, Loughborough University, UK.

Tulving, E. and Thomson, D.M. 1973, Encoding specificity and retrieval processes in episodic memory. *Psychol Rev*, **80**, 352–373.