COMPUTATIONAL INTELLIGENCE, III
N. Cercone & F. Gardin (Editors), G. Valle (Co-Editor)
Elsevier Science Publishers B.V. (North-Holland), 1991

Cognitive Perspectives on Knowledge Storage and Use

John Morton

MRC Cognitive Development Unit
17 Gordon Street
LONDON WC1H 0AH

The Headed Records model of human memory assumes that memory consists of independent records which can be accessed only by searching headings which are attached. Some evidence in favour of the model is described and the general computational advantages of such a system are discussed. The relation with case-based representations is noted.

I. INTRODUCTION

Psychological theorists are gradually coming to terms with the idea that commonly occurring, natural phenomena are within their provence (c.f. Neisser, 1976). This is not to say that models of memory should be judged solely on the basis of how well they can accommodate natural memory phenomena. However, we ought, by now, to be suspicious of any theory of memory that cannot at least contemplate an explanation for memory experiences that occur outside standard laboratory situations. We begin this paper, consequently, with a simple and naturalistic memory phenomenon which we will term "What was his name?".

Most people admit to having experienced the situation of feeling able to recount virtually everything they know about a particular individual except their name. Such incidents occur often to me. In one notable case, someone's research was being discussed. The main results were familiar to the people involved. We knew where the man worked, where he lived, the name of his wife and the last time he had given a talk at the Applied Psychology Unit. But, the person's name, lets call him BILL SMITH, eluded us. We knew that we would be able to recognise his name, if produced, and also that we would have able to reproduce all the information currently available to us had we previously just been given his name.

Given this phenomenon, my colleagues and I were interested to see how various theories of memory might account for it. We looked at the associative network approach (e.g. Anderson, 1976, 1983; Anderson & Bower, 1973;

Raajmaker & Shiffrin, 1980) and schema approaches (Bartlett, 1932; Rumelhart, 1980; Schank, 1980; Schank & Abelson, 1977) and found that, using these two major memory frameworks, we could only account for this everyday event by means of devices that seemed difficult to motivate theoretically (see Morton, Hammersley & Bekerian, 1985, and Morton & Bekerian, 1986). We attempted to specify an alternative theoretical approaches.

We first supposed, following Norman and Bobrow (1979), that memory consists of discrete units each containing information relevant to an "event", an event being, for example, a person or a personal experience. Information contained in a memory unit could take any number of forms, with no restrictions being placed on the way information is represented, on the amount being represented or on the number of memory units that could contain the same nominal information. Attached to each of these memory units would be some kind of access key. The function of this access key, we suggested, is singular: it enables the retrieval of the memory unit and nothing more. Only when the particular access key is used can the memory unit, and the information contained therein, be retrieved. As with the memory unit, we felt that any type of information could be contained in the access key. However, two features would distinguish it from the memory unit. First, the contents of the access key would be in a different form to that of the memory unit, e.g. represented by a different code. Second, the contents of the access key would not be retrievable, possibly due to the code they were in.

With this preliminary, and admittedly sketchy model, we returned to our example of Bill Smith. We assumed that all retrievable information about Bill Smith would be in the memory unit. The name "Bill Smith", would be in the access key for this unit but not in the unit itself. Since we were able to access the memory unit without the name, there must also be other elements to the access key the use of which allowed this access. The access key, therefore, did not have to be matched completely to allow access to the linked memory unit. Since the name, "Bill Smith", was in the key and not in the memory unit, it could not be recalled, in spite of our being able to recall all the other information about Bill Smith.

The only thing that immediately sprang to mind as a problem was this: How was it that we could ever remember Bill Smith's name if it was only in the access key and hence, by definition, not available for recall. Of course, this would only be only problematical if we assumed that a concept, such as the name Bill Smith, was represented only once in memory, as others do (e.g. Anderson, 1976; Anderson & Bower, 1973; Norman & Rumelhart, 1979; Schank,1980; Schank & Abelson, 1977). However, we had already assumed that the same nominal information could be represented both in memory units and in access keys, although the format would be different. Thus, given the possibility of multiple representations, the problem of the eventual recall of Bill Smith's name was resolved: the name Bill Smith, would be represented in a different memory unit from the one we had retrieved. In the original incident Bill Smith's name was found by first finding his wife's name and then consciously using that as a retrieval cue.

II. OUTLINE OF THE MODEL

IIa. Preliminaries.

The model is designed for retrieving answers to explicit questions, as a means of recruiting knowledge in the course of any cognitive operation, as a means of interpreting our environment and as a means of guiding our behaviour. These latter assumptions we share with a number of other theorists (Norman and Bobrow, 1979; Schank, 1980). We feel that our approach will force a discussion of how memory is influenced by the particular demands of the task and the way the individual perceives these demands (e.g. memory strategies). Ultimately, this will require a discussion of the processes through which memory operates. In other approaches, discussion of processes has largely been left implicit or driven simply by the need to get a simulation off the ground.

Remember that efficiency is not a primary concern. It isn't necessary that our memory system is the best solution for the problems it currently has to face. There are a number of inter-related reasons for this. The first is that things can selected against in evolution only when an alternative is in competition. Thus we could be stuck with a system which is good for a phylogenetically older and simpler organism. For example, it could be that we have a memory system which is maximally effective for moving from particular kinds of stimuli - smells, scenes, individuals (faces and other cues) - to appropriate action. Language and our complex conceptual system came later. Where the mankind line branched off from the ape line we did not need a different kind of memory to run the new systems: language & theory of mind. The new systems represented such an enormous leap that a memory better fitted for a creature using language and other symbols was not necessary.

IIB Formalising the Structure of Memory

We have already introduced the suggestion that there are two basic structures of memory, the memory unit and the access key. We will borrow the term *Record* from Norman and Bobrow (1979) to refer to the memory unit. Records are the unit of storage for recallable information in memory. There are a number of

features of Records that we would like to make explicit. First, Records are discrete; that is, are independent of one another. No connections link Records that happen to be related in terms of their content. In this respect we deviate significantly from network models and from schema models. Secondly, access to a Record is all-or-none; you either retrieve the Record or you don't. Thirdly, Records have no restrictions on the amount of information they contain, nor on the format for information. Any apparent limitations can be discussed more efficiently in terms of the demands set at the time of storage, the conditions prevailing at the time of retrieval or the interaction of these two. Fourthly, there is duplication of information in Records. The same event or any of its constituents can be represented in multiple Records, and the format for the event's representation may or may not vary. In this respect, too, we differ from other current viewpoints. Finally we shall maintain that once information is represented in a Record, it is not subject to alteration either by modification or addition of new information.

Records could contain a variety of types of information from lists or propositions to attributes of situational frequency or recency (Underwood, 1969) and can be at any level of abstraction. So, when one sees a film, one could have a Record containing a precis of the plot while other Records may contain lower level information such as the physical details of one of the characters. Whatever the information may be, we would want to maintain that Records will contain only processed information, unlike their access keys.

For the access keys we have chosen the term *Headings*. As stated earlier, Headings form the means through which Records are accessed and are made up of a number of distinct elements. Remembering involves first of all searching the Headings with some information (see section IIc) until a match is found. Only if a Heading is matched will the contents of its Record be accessible for recall.

We hold the following to be characteristic of Headings. First, the format of information contained in Headings is different from that in Records. Secondly, the content of the Heading need bear no propositional relationship to the content of its Record. Thirdly, the information in the Heading will determine the discriminability of its related Record in terms of access. Fourthly, the contents of Headings, as with Records, are not subject to change once they have been laid down. Finally, the contents of a Heading are inaccessible and can never be subject to recall.

Candidate information that might be contained in Headings are literal representations of the nominal event, including the environmental features surrounding the event, and internal states existing at the time an event is experienced. We suggest this because of the data indicating that reinstating such information at retrieval facilitates and enhances memory performance.

In this section we have focused on the basic structure of memory. However, one important feature of the Headed Records (HR) model is that it encourages specification of the processes of retrieval. It is to these that we turn in the next section.

IIc Processes of Memory

We assume that the method of access to the HR system is the same regardless of the circumstances of the retrieval. We have already stated our claim that only the Headings can be searched. The search process, then, involves looking for a match between some information and a Heading. The information that is used for the search we term the *Description*, following Norman and Bobrow (1979). The Description is formed from currently available information from external sources (such as an explicit question), internal sources (a Record that has just been retrieved) and a Task Specification. The Task Specification contains a list of the current goals. We will say more about the formation of Descriptions below. For the moment we need only consider that a Description can comprise a number of independent fields. These will include some environmental information and some internal state variables as well as lexical, propositional or other content.

The nature of the match required between the Description and a Heading will be a function of the type of information in the Description. If the task is to find the definition of a word or information on a named individual then a precise match may be required at least for the verbal part of the Description. We assume that the Headings are searched in parallel. On many occasions there will be more than one Heading that matches the Description. However, we require that only one Record be retrieved at a time, an assumption we share with a number of other memory theorists (Anderson, 1976; Anderson & Bower, 1973; Rumelhart, 1980; Schank, 1980). Evidence in support of this assumption is summarised in Morton, Hammersley & Bekerian (1985). The data indicate that the more recent of two possible Records is retrieved. We conclude firstly that once a match is made the search process terminates and secondly that the matching process is biassed in favour of the more recent Heading. The latter is a common proposal in discussions of forgetting (e.g. Hasher et al, 1981; Martin, 1971; Postman & Underwood, 1973) and has been formally incorporated into many models of memory (e.g. Anderson, 1976; Anderson & Bower, 1973).

In summary, then, the first stage of the retrieval process involves forming a Description which is then used to search the Headings. If more than one Heading matches the Description

then the more recent is selected.

There are two possible outcomes of a search cycle; either it is successful, and a match is found between the Description and a Heading, or no match is found. When a Heading has been matched, its Record will be retrieved and made available for further processing. Some Records will contain prescriptions for actions, and these will just be run off. More generally the Record will be evaluated in the light of the goals represented in the Task Specification. If the current goal is to answer a query, the Record will have to be searched for the required information. To give two simple examples: if one is going to phone home, the retrieval process will come up with a Record with only one entry which would be sent on to the motor systems without any evaluation (and, if one had recently moved a mistake could occur); if you need the address of a theatre, the Record that is retrieved may need checking in order to verify that it is indeed the theatre one intends to go to. Indeed, the addresses of a number of theatres could be in the same Record and the particular address would need to be selected.

There is, of course, no guarantee that the retrieved Record will contain the information that is sought. The Record of theatres, to continue the previous example, may be incomplete or for the wrong city. In such cases, or in the case that no Record had been retrieved, there are two options: either the search is continued or it is abandoned. If the search is to be continued then a new Description will have to be formed since searching again with the same Description would result in the same outcome as before. Thus, there has to be a list of criteria upon which a new Description can be based. The process responsible for creating Descriptions, then, will have REDESCRIPTION as one of its modes of operation and FAILURE as one of its outcomes. In the case that a new Description is formed then the whole retrieval cycle will be repeated. The progress of the retrieval cycle has been traced by Williams and Hollan (1981). In the task they set their subjects, recalling the names of the members of their high school classmates, much of the cycling process appeared to be conscious. Deliberate decisions were made, for example, concerning what might be a likely strategy. However, we see no reason for tying the retrieval cycle to consciousness.

To recap, the phases of the retrieval cycle are:

1. The formation of a Description
2. Search of the Headings for a match to the Description
3. Retrieval of the associated Record
4. Evaluation of the retrieved Record.

Let me make the operation of these phases more explicit by working through some typical cases. In the course of doing this it will become clear that certain supplementary processes and classes of information will also be required.

We can first take the simple case where an explicit question has been asked such as "What is Bill Smith's address?" The speech input would be converted into an appropriate code by perceptual and linguistic processes. We have already specified that these processes are outside the range of our current interests, being autonomous structures that do not rely on the Headed Records system for their operation. Two things now have to happen: the query has to be converted into a Task Specification and, in addition, a Description has to be formed. In the case of language inputs these two procedures will have something in common. Thus the language processes will mark "Bill Smith" as a proper noun and "Bill Smith's address" as a topic noun phrase. The latter will form part of the Task Specification and the former would form part of the Description, since one of the likely rules for Description formation would be to use proper nouns as Descriptions. If there were a Record headed BILL SMITH it would be retrieved. It would then have to be examined in the light of the Task Specification to determine whether it actually contains Bill's address. This analysis shows that the following processing components are required:

1. A set of rules for extracting goals and plausible descriptors (like proper nouns) from an input string.
2. A store for holding the retrieved Record.
3. A process for examining the retrieved Record.
4. A Task Specification used to guide the evaluation of the retrieved Record.

Note here the rule-governed nature of the Description formation. The set of rules used for obtaining descriptor terms and goals will, itself, be in a Record, access to which would be a default condition for the control processes in the case that there is linguistic input to the system. It is not the case, for example, that all the contents of a retrieved Record are instantly and automatically used as a Description.

The nature of the evaluation processes does not seem to pose any special problems. A number of alternative proposals are already in the literature. If, for example, it were convenient to represent the internal structure of Records in terms of networks then the problem of examining a Record would be subject to any one of a number of current solutions (e.g. Anderson, 1976; Hayes-Roth and Hayes-Roth, 1977).

It should be clear that the processes responsible for forming Descriptions must be closely related to those responsible for creating Headings. At least they must both have access to the same guiding principles. We merely note, as an interesting developmental problem, that these principles will have to be learned by the child and that if the principles change then retrieval of Records laid down before the change would be very problematical.

We can now move to the operation of the system in interpreting situational experience; being in a restaurant, for example. The principles of operation are the same. In the course of our development we will have built up a large number of routines to guide our behaviour in different circumstances. Among these routines will be Records containing information about those features of the environment that can be used as Descriptions to access appropriate Records. The basic cycle involves first of all the default retrieval of the control Record used to select a salient feature of the environment as a suitable descriptor. Next a Description is formed and the search process leads to the retrieval of a referent Record which is used to interpret the environment and guide our actions.

The evaluation process will be operating continuously during the last of these stages to ensure that the referent Record continues to be appropriate. Such a process is not peculiar to the present model but would also be required in some forms of schema theory. Other forms of schema theory avoid this problem, at least in principle, since appropriate schemata would be switched in, relatively passively, by being directly stimulated by the changing features of the environment (see Rumelhart, 1980).

A switch of referent record may be necessitated by a change in the environment (if, for example, someone starts a brawl in the restaurant), a change in the demands imposed in the same environment (e.g. starting to discuss business with one's dinner companions after having ordered dinner) or some change requiring a general problem solving routine to be retrieved (as when one's spouse arrives unexpectedly in the same restaurant). In all these cases the sequence of events would be:

1. Detection of the inadequacy of the referent Record by the evaluation routine
2. Formation of a new Description
3. Retrieval of a different referent Record.

We have assumed that only one referent Record can be used at a time. This is the simplest assumption. If we assumed that more than one Record could be used at a time in this way, then the processing apparatus would have to be more complex. In either case we could not expect a very direct relationship between the dynamics of behaviour and the underlying representation. The processes that mediate between the representations in the Records and actual behaviour will have the effect of smoothing over the underlying joints, much as the underlying structure of an utterance is disguised by the time it becomes speech.

We have already pointed to the need for representation of internal states and environmental features both in the Headings of the Records of those memories and in the Descriptions used to retrieve them. The simplest conceptualisation for our purposes is to imagine a set of registers in which these states are noted. These will be updated automatically. In addition to indices for mood and physiological state we would expect registers concerned with where we are (which country, city, general locale), what day it is, what time of day, what date (for some people) and who one is talking to. All these, at least, are features of our instantaneous being. Whether or not they all participate in Headings and Descriptions is an empirical matter to be tested by experiments using the state dependent learning methodology. Inasmuch as such indices do participate in Headings and Descriptions they will participate in all Headings and Descriptions and this will be an automatic function.

IId. Summary of the Model

- Headings include state information and literal features of the stimulus and the environment. Headings cannot be retrieved.
- Records contain processed information and can be accessed only if their linked Headings have been matched.
- Retrieving information from the Headed Records system involves the formation of a Description that contains information which in principle can match that found in a Heading. The selection of information for the Description will be influenced by the goals currently demanded by the external environment and any other information that might be relevant to those goals (e.g. a related Record that has been retrieved). Rules for constructing the Description are also used in the construction of Headings. However, the actual strategic selection of information is assumed to be contained in control Records that are summoned whenever there is a demand for information.
- The search process involves a match between the Description and a Heading. The nature of the match will be influenced by the interaction between current goals and the nature of the information contained in the Description (resulting from choosing a strategy as how to achieve these goals). The search process either results in a successful match or in a failure. In the case of a match, the Record linked to the matched Heading will be made available. At this point, goal-related and strategy-related information will be influential in evaluating the appropriateness of the retrieved Record, and the selection of specific information that either satisfies goals or allows for alternative strategies to be put into operation, as in the case of reformulating the Description.
- A Description will be reformulated in the situation where there has been a failure to match a Heading. Reformulation will also occur when the information retrieved is relevant to the goals but is insufficient. In the case of a failure, the Description must be changed if another Record is to be retrieved. The other option is to terminate search and to discard the unsuccessful goals from the buffer.

III. APPLICATION OF THE MODEL

IIIA. Natural Memory Phenomena

One of the reasons for the development of the Headed Records model was the difficulty of talking about natural memory phenomena in the context of existing models. There are a number of memory problems that are sufficiently pervasive to justify the status of data. One of these was discussed in the introduction. The following seems to have the same status.

Don't You Remember? - An experience reportedly shared by everyone (at least everyone from a sample of over 300 people) is that of being reminded of an incident but completely failing to recall it. The cues given may be overwhelming - for example, when and where the incident occurred, who was concerned, some of the events - but there is still no recall. These cues may be more than sufficient to describe the event uniquely but fail completely to trigger it. The cue that finally triggers the memory may be no more significant than any other cue, and may even be incidental. In Headed Records terms, this cue must be in the Heading of the Record for that particular incident and the other, seemingly more central information could only be in the Record. The theory does not require that all the components of a memory can be used for access neither does it predict which elements will be so used. The point we wish to make at the moment is that the registration of incidents in our lives is a complex and variable affair. If we are to begin to understand it we will have to have a framework within which the necessary processes can be anatomised and the appropriate questions asked. Unpublished data with 'reminding' experiments indicates that locations and other people are usually sufficient cues as are single unique features. There are, however, considerable individual differences. One of our subjects, for example, could only recall some incidents when cued by the actions that preceded them.

The memory phenomena we have described encapsulate the structural distinction that is central to the Headed Records system. They illustrate that there are two kinds of information. There is information, the Record, that is present in memory but can not be used as a means of access and there is information, the Heading, that cannot be recalled but can be used as a means of access.

Fragmentary Knowledge - In the two memory frameworks that are most commonly used, associative networks and schema, it is natural to think in terms of the progressive and inevitable accumulation of related knowledge. Within the Headed Records model such is not the case. To start with, it is a central assumption that Records are never changed. When new information arrives it is interpreted by reference to an existing Record. As a consequence of this interaction a new Record will be created. Certainly, relevant information from the existing Record will be transferred to the new Record, but the existing Record will remain as it was. To give a simple example, suppose that we learned that BILL SMITH had moved to a new address. It would be sufficient to set up a new Record which included that address with just enough information from an old Record to identify BILL SMITH, but without all the information one had on all the places he used to live at. The old information would still be present in memory but would not be directly accessed when thinking about the new place.

Evidence in favour of the fragmentation of memory comes from a number of sources. A study by Hammond, Morton, Barnard, Long and Clark (1984) provides strong evidence. Subjects were required to learn a computer system which included eight editing commands. A scrambled sentence to be edited was displayed on a V.D.U. screen. Using an editing command resulted in the creation of a modified sentence which was displayed below the previous one. Each of these messages had a reference number attached to it that had to be entered when any command was used. Such variables are known technically as "arguments". The argument in question referred to the message being changed but, in the normal course of events, was also the number attached to the bottom-most message. The subjects were not given an explicit definition of this argument but had to infer it from examples. If the subjects made an editing error a command would have to be issued with a message number that no longer corresponded to that attached to the bottom-most message. Most subjects failed to use the argument corresponding to the message to be corrected. They then had to correct that error. While subjects were invariably correct subsequently on commands on which they had made this argument error, there seemed to be no generalisation from command to command.

One can only conclude that there was a separate Record set up for each command, each having to be corrected individually. In the usual associative network representation of the required knowledge, or in a schema model, there would be a single node specifying the general argument message number to which would attached its definition. Any correction to this definition should immediately generalise to all commands. Only in a model which allows for fragmentary knowledge is the observed pattern of errors intelligible. Hammond, Morton, Maclean, Barnard and Long (1982) give further examples of fragmentary learning of computer systems.

IIIB AI Relevance.

What might the advantages of such a memory be?!

1. Restricted content addressing, while increasing the likelihood of search failure, speeds up search time

   a. what you search with is restricted to the set of descriptors.
   b. where you search is restricted to headings.
   c. post-search evaluation - fewer records will have to be examined in detail. This will make more difference with a serial search than with a parallel one but, in general, will reduce the number of items retrieved.

2. Having context specificity in the Heading stops the system trying (inappropriately) to apply the solution to a related problem in a (significantly) different context. The most extreme example of this for humans would be the way in which we use information obtained by observing others. If we just learn by trial and error - even if the choice of trial is based on as complex a cognitive computation as we could imagine, it is still going to be a long and painful process to find the solution path. If, on the other hand, we are allowed to observe others solving the same problems, specifically, watching adults - then we get good clues as to which directions to move in. However, we cant just go ahead and apply the solutions which others have succeeded with, because our effector systems are different and because the state of the world, and its reaction, depends on who we are. So, for example, we are not as strong as our father, and so cannot climb, carry or turn or otherwise manipulate in the crucial step; or we are not as large or brave or with sufficient status.

SO: having a SELF/OTHER distinction in the headings, allows us the best of both worlds. We can distinguish two extreme states: one of them being when the situation requires a fast reaction. We do not want to access an OTHER headed record when we have to act in a hurry. We stick to something with a successful outcome for ourselves. Given more time we can relax that constraint in the Description, have a look at other people's solutions and try to adapt them for our own capability.

3. Keeping all past records allows you to get out of local minima.

4. Use of the HR architecture has the advantage of allowing you to make the most of any kind of structure. A record can be a schema or a network of any other kind, an ordered list, an unordered list: it can be a compiled output, a high level program fragment, a text, instructions for interpretation of the current input or anything else that seems to be appropriate. What is principled is the storage and handling.

5. HR allows (as a matter of principle) a mixture of script and instance-based interpretations/action. In fact, each event record will be a mixture of the two. A decision which has to be made, both experimentally and computationally, is whether or not we want to have "pure" scripts or mops or what have you as well.

6. A realistic parallel search (followed by serial evaluation).

7. The Heading-Description system configures itself for the particular environment. Therefore, by the time that language has established itself in the picture, plus other factors, we can no longer address HR which were laid down in the first few years. Thus, infantile amnesia - well established as a real phenomena - doesn't need a Freudian explanation. So the same device can be trained to operate perfectly well in any domain or combination of domains. The more there are and the wider the significant context dependencies the longer it takes to act and the higher the risk of error. Secondly, learning (of which development is a sub-set) takes place every time a record is accessed to deal with the current situation and a new record laid down with (hopefully) improvements.

IIIC. Multiple Representations

One consequence of a Headed Records system is that there will be multiple representations both of knowledge and of skills. This will occur for a number of reasons. First of all, procedures corresponding to mental or physical skills will be represented in records. Whenever such a developing skill is used and is changed as a result of this practice then there will be a new record of the improved form of the procedure. However, the old form will remain, since there is no overwriting in the system. The new form will be used next time the procedure is called, following a recency principle, but the old form would be re-used if, for example, the context of use of the two forms had been different and the context for the old form was reinstated. Another reason for preferring to have multiple representations is that it allows the agent to experiment with alternative representations without destroying the original representation. The issue of re-representation has been extensively studied in child development by Karmiloff-Smith (1986, see also Campbell, this volume for computational advantages).

IV. Case Based Reasoning

The Headed Records theory stands in relation to the schema theory of memory as case-based knowledge representation stands to rule-based representation. The role of case-based representations is discussed by Faltings (this volume) in the following way: "By storing knowledge as precise cases, we not only solve

the problems of complexity and non-compositionality, but also that of context-dependence by generalizing with respect to a *particular* problem." (original italics). Oppacher and Deugo (this volume) claim that traditional case-based reasoning has several shortcomings. They give as an example the possibility that newly added cases may redundantly overlap, but this would not be a problem unless there is storage pressure. Oppacher and Deugo further remark that "cases developed a long time ago may have outlived their usefulness in the current environment." If the conditions under which cases are invoked include the existing environmental specification, as with the Headed Records model, then, unless storage is at a premium, this would not be an issue. And there is always the possibility that old conditions would arise again, requiring the use of old cases.

REFERENCES

Anderson, J.R. (1976). *Language, Memory and Thought*. Hillsdale, NJ: Erlbaum.

Anderson, J.R. (1983). *Architecture of Cognition*. Cambridge, MA: Harvard University Press.

Anderson, J.R. and Bower, G. (1973). *Human Associative Memory*. Washington, D.C.: Winston.

Bartlett, F.C. (1932). *Remembering*. London: Cambridge University Press.

Hammond, N.V., Morton, J., Barnard, P.J., Long, J. and Clark, I.A. (1984). Characterising user performance in command-driven dialogue. Hursley Human Factors Report HF093, IBM UK Laboratories, Hursley Park.

Hammond, N.V., Morton, J., MacLean, A., Barnard, P.J. and Long, J. (1982). Knowledge fragments and users' models of systems. Hursley Human Factors Report HF071, IBM UK Laboratories, Hursley Park.

Hasher, L., Attig, M. and Alba, J. (1981). I knew it all along: Or did I? J. *verb. Learn. verb. Behav.*, **20**, 86-96.

Hayes-Roth, B. and Hayes-Roth, F. (1977). The prominence of lexical information in memory representations of meaning. *J. verb. Learn. verb. Behav*, **16**, 119-136.

Martin, E. (1971). Verbal learning theory and independent retrieval phenomena. *Psychol. Rev.*, **78**, 314-332.

Morton, J., & Bekerian, D.A. (1986). Three ways of looking at memory. In N.E.Sharkey (ed.) *Advances in Cognitive Science 1*. Chichester: Ellis Horwood.

Morton, J., Hammersley, R.H., & Bekerian, D.A. (1985). Headed records: A model for memory and its failures. *Cognition*, **20**, 1-23.

Neisser, U. (1976). *Cognition and Reality*. San Francisco: Freeman.

Norman D.A. & Bobrow, D.G. (1979) Descriptions: an intermediate stage in memory retrieval. *Cognitive Psychology*, **11**, 107-123.

Norman, D.A. and Rumelhart, D.E. (eds.) (1975). *Explorations in Cognition*. San Francisco: Greene Man.

Postman, L. and Underwood, B. (1973). Critical issues in interference theory. *Mem. Cog.*, **1**, 19-40.

Raajmakers, J. and Shiffrin, R. (1980). SAM: A theory of probabilistic search of associative memory. In G. Bower (ed.) *The Psychology of Learning and Motivation*, Vol. 14. New York: Academic Press.

Rumelhart, D. (1980). The building of blocks of cognition. In R. Spiro, B. Bruce and W. Brewer (eds.) *Theoretical Issues in Reading Comprehension*. Hillsdale, NJ: Erlbaum.

Schank, R.C. (1980). Language and memory. *Cog. Sci.*, **4**, 243-284.

Schank, R.C. and Abelson, R.P. (1977). *Scripts Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale, NJ: Erlbaum.

Underwood, B.J. (1969). Attributes of memory. *Psychol. Rev.*, **76**, 559-573.

Williams, M.D., & Hollan, J.D. (1981). The process of retrieval from very long memory. *Cognitive Science*, **5**, 87-119.