Cognition, 20 (1985) 1–23

Headed records: A model for memory and its failures*

JOHN MORTON
MRC Cognitive Development Unit

RICHARD H. HAMMERSLEY
University of Lethbridge

D.A. BEKERIAN
MRC Applied Psychology Unit, Cambridge

Abstract

It is proposed that our memory is made up of individual, unconnected Records, to each of which is attached a Heading. Retrieval of a Record can only be accomplished by addressing the attached Heading, the contents of which cannot itself be retrieved. Each Heading is made up of a mixture of content in more or less literal form and context, the latter including specification of environment and of internal states (e.g. drug states and mood). This view of memory allows an easy account of a number of natural memory phenomena as well as a variety of laboratory findings such as the differences between recall and recognition.

The theory further proposes that Headed Records can neither be deleted nor modified. Data apparently against such a hypothesis can be accounted for in terms of the retrieval process.

1. Introduction

Psychological theorists are gradually coming to terms with the idea that commonly occurring, natural phenomena are within their province (cf. Neisser,

*The three authors claim they are equally responsible for the ideas presented in this paper. We are grateful to Philip Barnard and John Bowers for extensive discussions and to the members of THLUTS at the Applied Psychology Unit, Cambridge, for their encouragement and criticisms, notably Richard Young and Hilary Buxton. We would also like to acknowledge the comments of Z. Pylyshyn on earlier versions of this paper. A substantial part of this work was done while the first two authors were still members of the Applied Psychology Unit. Portions of this paper were first presented at the meeting of the Experimental Psychology Society in Cambridge, July 1979 and at numerous places since.

Reprint requests should be sent to John Morton, MRC Cognitive Development Unit, 17 Gordon Street, London WC1H OAH, United Kingdom.

0010-0277/85/$7.40

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. One such incident recently occurred to two of the authors of this paper. 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 eluded us. We knew that we would be able to recognise his name, if produced, and also that we would have been able to reproduce all the information currently available to us had we previously just been given his name.

Given this phenomenon, we were interested to see how various theories of memory might account for it. We started with an associative network approach, as this framework is one of the most widely used (e.g. Anderson, 1976, 1983; Anderson and Bower, 1973; Raajmakers and Shiffrin, 1980). In the idealised associative network, concepts, such as the concept of a person, are represented as nodes, with nodes that are associated being connected through links. Generally speaking, the links define the nature of the relationship between nodes, e.g. the subject–predicate distinction (Anderson, 1976; Anderson and Bower, 1973). Let us suppose that the name of the person we are trying to recall is Bill Smith, We would have a BILL SMITH node (or a node corresponding to Bill Smith) with all the available information concerning Bill Smith being linked to form some kind of propositional representation. Now, failure to retrieve Bill Smith's name, while at the same time being able to recall all other information concerning Bill Smith, would have to be due to an inability to traverse the links to the BILL SMITH node. However, this seems contradictory to one principle of associative networks: content addressability (e.g. Anderson, 1976; Anderson and Bower, 1973). That is to say, given that any one constituent of a propositional representation can be accessed, the propositional node, and consequently all the other nodes linked to it, should also be accessible. Thus, if we are able to recall where Bill Smith lives, where he works, who he is married to, then, we should, in principle, be able to access the node representing his name. To account for the inability to do so some sort of temporary "blocking" of content addressability would seem to be needed. Alternatively, directionality of links would have to be

specified (Anderson, 1976; 1983) though this would have to be done on an ad hoc basis.

Next we considered schema approaches (Bartlett, 1932; Rumelhart, 1980; Schank, 1980; Schank and Abelson, 1977). Schema models stipulate that there are abstract representations, i.e. schemata, in which all invariant information concerning any particular thing is represented. So we would have a person schema for Bill Smith that would contain all the invariant information about him. This would include his name, personality traits, attitudes, where he lived, he had a family, etc. It is not clear how one would deal with our example within a schema framework. Since someone's name is the quintessentially invariant property, then, given that it was known, it would have to be represented in the schema for that person. From our example, we knew that other invariant information, as well as variant, non-schematic information (e.g. the last talk he had given) were available for recall. This must be taken as evidence that the schema for Bill Smith was accessed. Why, then, were we unable to recall one particular piece of information that would have to be represented in the schema we clearly had access to? What we would have to assume is that within the person-schema for Bill Smith are sub-schemata (e.g. Schank, 1980) one of which contained Bill Smith's name, another containing the name of his wife, and so on. We would further have to assume that access to the sub-schemata was independent and that, at the time in question, the one containing information about Bill Smith's name was temporarily inaccessible. Unfortunately the concept of temporary inaccessibility is without precedent in schema theory and does not seem to be theoretically motivatable (for other recent criticisms of the schema framework see Alba and Hasher, 1983).

At the end of this little theoretical exercise we were somewhat disillusioned. We had tried to explain a simple memory phenomenon and found that, using the two major memory frameworks, we could only do so by means of devices that seemed difficult to motivate theoretically. Note that we are not simply concerned with the temporary inaccessibility of any piece of information in memory. It is the particular piece of information, someone's name, that conerns us. The phenomenon is pervasive with respect to people's names, and there does not appear to be any generally accepted variable, such as frequency of usage, that is plausible as an account of the inaccessibility of names. Were there alternatives, then, to the theoretical approaches we had tried? We felt there were and attempted to specify one possibility.

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 restric-

tions 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 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 and Bower, 1973; Norman and Rumelhart, 1975; Schank, 1980; Schank and 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.

As a first effort, our rather simplistic hypothesis fared reasonably well. Encouraged by this we decided to pursue our notion of memory units and access keys, formalising and specifying them so that some consistent theoretical skeleton could emerge.

2. Outline of the model

2.1. Preliminaries

Before we start to outline the model we must put it in its proper context. We intend to discuss only a small part of our mental apparatus, namely that directly related to the formation and retrieval of units of memory. As an additional restriction, what we mean by memory has to be understood by reference to our overall theoretical position. Thus we are not concerned here with input analysis, input or output lexicons (Clarke and Morton, 1983), object categorisation (Warren and Morton, 1982) or the properties of the Response Buffer (Morton, 1970). Note particularly that object names, unlike proper names, are stored in the output logogen system (Warren and Morton, 1982). Syntactic operations (Morton and Patterson, 1980) and speech or other output organisation processes also lie outside our current interest. In short, we are concerned only with information that has undergone considerable processing. In the case of language this processing would include lexical categorisation together with some parsing as well as the extraction of information concerning the identity of the speaker, tone of voice and so on. In the case of objects, both categorical and stimulus specific information will be available (Warren and Morton, 1982). The same will be true for more complex scenes, as witnessed by performance on recognition memory experiments for pictures (Standing, 1980). We will assume this dissociation between elements of our cognitive apparatus for the sake of convenience, acknowledging that not all would agree (e.g. Jacoby and Witherspoon, 1982).

Our 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 behavior. 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.

2.2. 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. Second, access to a Record is all-or-none; you either retrieve the Record or you don't. Third, 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. Fourth, 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. Evidence in relation to these claims will be presented in Sections 3.1 and 3.2.

Records could contain a variety of types of information ranging 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. In a list learning experiment, one may have a global Record containing sets of items and other Records for individual items (Postman and Underwood, 1973). 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 2.3) 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. Second, the content of the Heading need bear no propositional relationship to the content of its Record. Third, the information in the Heading will determine the discriminability of its related Record in terms of access. Fourth, 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. This evidence is summarised in Section 3.3.

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.

2.3. 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 shall 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 and Bower, 1973; Rumelhart, 1980; Schank, 1980). Evidence in support of this assumption will be described in Section 3.3 and the data indicate that the more recent of two possible Records is retrieved. We conclude first, that once a match is made the search process terminates and second, that the matching process is biased in favour of the more recent Heading. The latter is a common proposal in discussions of forgetting (e.g. Hasher, Altig and Alba, 1981; Martin, 1971; Postman and Underwood, 1973)

and has been formally incorporated into many models of memory (e.g. Anderson, 1976; Anderson and 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 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.

We will 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.

Let us 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).

Let us now take a more complex case of memory search. Suppose we ask a subject "I'd like you to remember the word list I presented to you three months ago". Initially, the set of rules might extract *WORD LIST* as a descriptor since this phrase would be marked as the topic of the utterance. Suppose the resulting Description did not find a matching Heading or that if it did that the retrieved Record was one of a word list presented by a different experimenter the day before. In such cases other retrieval cues have to be used. At this stage a strategic component enters in the formation of a new Description. The options are either using a temporal modifier related to "three months ago" or using environmental cues such as the experimenter herself or the room in which the interview is taking place. The decision between these will be guided by the rules attached to the processes responsible for Description formation. In behavioural terms, the subject might look at the experimenter and then look round the room. The subject would report "I am trying to remember the last time I was here." To translate all of this into the model, we would say that various features fo the environment were being used as candidate descriptors, either individually or in combination, with the retrieval process recycling. For this to be possible we need to have perceptual processes responsible for face analysis and environmental analysis, similar to those required for language inputs. These, too, would be outside the scope of the Headed Records system. In addition we need another set of descriptor rules for selecting plausible candidates.

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 appratus 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.

2.4. 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 influencial 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.

3. Application of the model

3.1. Natural memory phenomena—some evidence on structure

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, 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. Our data with 'reminding' experiments (Bekerian and Morton, note 1) 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 two 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 cannot 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 transfered 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. At one extreme is the work of Jones (1981; in press) in a variety of memory tasks. Jones (in press) has shown that performance in cued recall for word lists is best described by a model in which the memory representation for the word list is allowed to consist of independent fragments. A study by Hammond, Morton, Barnard, Long and Clarke (1984) provides evidence from a different kind of setting. Subjects were required to learn a computer system which included eight editing commands. A scrambled sentence to be edited was displayed on a VDU 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.

Of the 48 subjects used in the study, 44 made the error, in the circumstances described, of specifying the number on the bottom-most line instead of the number attached to the message string they were editing. That is, they had not inferred the correct definition of the argument. It might be expected first, that there would be one trial learning of the correct argument, and second, that the correction would generalise immediately across all eight commands. Of the subjects who made the error, 11 did not repeat it. Of the remaining 33, the majority, 28, did not make the error again when using the command on which it first occurred but did make it on other commands. In fact, the best description of the data was that the correction as to the nature of the message number argument had to be made separately for every 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 be 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.

3.2. Eyewitness testimony

The topic of eyewitness testimony provides an opportunity to discuss theoretical issues that have far-reaching practical application. In this section, we will focus on the evidence surrounding the effects of post-event information (e.g. Loftus, Miller and Burns, 1978) and particularly the claim that such information can alter the contents of our memories.

The basic design used to demonstrate the effects of post-event information conforms to the traditional retroactive interference design: people first view a slide sequence or film; they then receive post-event information about the details of the event that are inconsistent with what was originally seen; finally a memory test is given for the original event.

Using this basic paradigm, Loftus and her co-workers (Loftus, 1975, 1979a and b; Loftus et al., 1978) have argued that memory is modified to conform to post-event information, resulting in recollective errors. In one study subjects were shown a filmed car accident (Loftus, 1975). Later they were told that a barn had appeared in the film. Although the barn had not, in fact, appeared in the film, over 17% of the subjects, when questioned a week later, agreed that they had indeed seen one. This compared with only 3% of agreement by subjects who had not been fed this piece of misinformation (Loftus, 1975). In another experiment subjects were shown slides depicting a different accident. The manipulation in this case involved detail on one of the slides. Thus, some subjects saw a scene with a STOP sign in it and were later told in a questionnaire that the sign had been a YIELD sign. In some conditions in this experiment, as many as 80% of the subjects indicated at the time of testing that they had seen a slide with the incorrect piece of information (Loftus et al., 1978, Expt. 2). Loftus says: "the new 'sign' information was apparently integrated into the subjects' memorial representation of the event, possibly altering that representation." (Loftus, 1979b, p. 368).

Other seemingly convincing data for the alteration of memories comes from experiments on second guesses (Loftus, 1979a). The second guess technique requires subjects first to choose from a number of fixed alternatives and, assuming they are incorrect, then guess among the remaining alternatives. If the correct response is given on second guess trials at a probability greater than chance, this is taken as evidence that some veridical information is available despite the error on the first guess. Loftus (1979a) used this technique on subjects who had seen misleading information before being asked to perform the memory test. Briefly, the results showed that subjects who chose incorrectly on the first guess did not choose correctly on the second guess at a level that was greater than chance.

The Loftus data is superficially convincing with respect to the modifiability of memories. However, an alternative explanation for the misleading effects

of post-event information can be developed within the Headed Records position. This is made possible by our allowing duplication of information in different Records. The argument goes as follows. There is a Headed Record (or Records) that contains information about the original slide sequence. In addition there would be another Headed Record (or Records) representing the event within which the inconsistent post-event information was embedded. In the Loftus et al. experiment, this event was a questionnaire that had to be answered. Both Records would exist in memory, although only one could be retrieved at test. The accessibility of the Record for the slide sequence requires two conditions: first, its Heading must be unique and discriminable from the Heading for the questionnaire Record; second, this unique information (or some subset) must be present at test so that it can be included in the Description formed.

The misleading effects reported by Loftus et al. (1978, Expt. 1) can be explained in terms of the absence of critical information at test. In the Loftus et al. experiment, the pairs of test slides had been arranged in an order that was random with respect to the original slide sequence. This would have important consequences, as we would hold that the sequential nature of the original slide presentation would be represented in the Heading of the Record for the slide sequence. With random presentation at the time of testing, this sequential (or theme-related) information would be missing from the Description. Accordingly, the Heading for the more recent Record (containing the inconsistent information) would be matched, its Record made available and misleading effects observed. However, given a sequential test order, the Description would include information allowing the Heading for the Record of the original information to be matched. No significant effects of the misleading information would then be observed.

This interpretation, which emphasises retrieval failures, has received support in a number of studies (Bekerian and Bowers 1983; Bekerian and Bowers, note 2; Bowers and Bekerian, 1984). Bekerian and Bowers (1983) repeated the first two phases of the Loftus et al. experiment—the slide sequence and the questionnaire. The critical variation came in the final test. In one condition the test slides were randomised as in Loftus et al. In another, the test slides were presented in the same order as in the original sequence. The misleading effects reported in Loftus et al. were replicated only when the test slides were presented in a random order. No misleading effects were found when the original sequence was followed at test. Bekerian and Bowers (note 2) have replicated these findings with the original material being verbal and the post-event information visual, and has further established that the modality of the input must also match that of the original event in order for the sequential manipulation to be effective. This is evidence for surface form in Headings.

Loftus' data on the second-guess technique (Loftus, 1979) can also be explained in terms of a failure to provide adequate information at test thus preventing access to the original Record. In the Loftus experiment the task demands remained the same for both first and second guesses: the subject was simply instructed to make a second guess with no feature of the retrieval environment being altered. The Description for the first guess would access the most recent Headed Record, i.e. that containing the misinformation. On the second guess, the Description would effectively remain the same and would lead to the same Record. Under these conditions the correct information would not be available and so the second guesses would have to be at chance. Following this logic we would expect to change the behaviour on the second guess if the subjects could be induced, by means of experimenter instructions, to change their Description and to focus on features unique to the slide sequence. Data from Bekerian and Mingay (note 3) bear out this prediction.

The other factor contributing to post-event misleading effects is the relative discriminability of the Headings. In the model we claim that only one Record can be accessed for each retrieval cycle. We have also claimed that if two Headings match a Description, then the most recent will tend to be selected. The only way of accessing the more remote memory would be by creating a Description that discriminated the two Headings. Thus, in the experiments under discussion, if the Heading for the original information is not discriminable from that for the questionnaire Record, then the latter will always be accessed. Results from a recent experiment by Christiaansen and Ochalek (1983) can be explained in these terms. These authors used the standard procedure, following Loftus et al., except that the subjects were given a recognition test before the post-event information was introduced. The post-event information was in the form of a written summary of the slides. After some delay, subjects were given a final recognition test. Prior to this final test, half the subjects were warned that some of the information they had read in the summary was incorrect. Christiaansen and Ochalek found that subjects who had not received the warning were subsequently misled by the inconsistent post-event information even if they had been correct on the first recognition test.

We interpret these data in terms of the discriminability of Headings. Subjects who were correct on the first test would detect the inconsistencies in the post-event information. The Heading for the Record of the summary would reflect this. At the final test subjects would not spontaneously create a Description that focused on the features that discriminated the two Headings. However, if the warning was given prior to the final test, this would serve to force the subjects into creating a Description that was biased towards the Heading for the original Record.

In summary, the explanation we have advanced to account for the effects of post-event information is consistent with our assumption that Records are unalterable. We emphasise the role of retrieval difficulties in accessing the original information and have pointed to two underlying factors. First, if the test environment does not reinstate crucial features present at the original encoding, the Description will fail to match the Heading for the original Record. Second, the Headings for the original and post-event Records must be discriminable if the former is to be accessed.

3.3. Recall, recognition and the specificity of memories

One of the most pervasive phenomenon in memory is that performance on recognition tasks is generally superior to performance on recall tasks (see MacCormack, 1972). From our perspective, these data are accounted for by the assumption that a "literal" representation of the stimulus is one constituent in the Heading of an H–R pair. In a recognition test, the stimulus item itself would be used in the formation of a Description. Provided there is a Heading with this item in it, retrieval of the Record should be fairly problem-free. In contrast, recall requires that the Description be formed on the basis of features currently available, like the experimenter or the surroundings. The chances of forming a Description that would fail to match the target Heading would be greater, and the chances of retrieval failure would thus increase.

There are a number of findings indicating that one of the types of information that can be in the Heading is a more-or-less literal representation of the stimulus. When test material is presented in the same form as in the original presentation, performance on recognition tasks is improved. This has been found for manipulations of modality (Kirsner, 1974), manipulations of speaker's voice (Craik and Kirsner, 1974; Geiselman and Bjork, 1980), manipulation of type face (Kirsner, 1973) and surface format (Kolers, 1979). In recognition, the presentation of the literal stimulus accesses the relevant Heading. The Record would contain information, such as situational information, sufficient to perform the task (cf. Underwood, 1969). Next, we can note that performance on recall tasks has been shown to improve when the environmental features at test were identical to those at encoding. Thus Baddeley, Cuscaro, Egstrom, Weltman and Willis (1975) showed "a massive effect of context dependency" when divers were asked to recall passages of prose learned under water. Godden and Baddeley (1975) showed a 30% decrement in recall of word lists learned either on the beach or under about 15 feet of water and recalled in the other environment. The third likely

feature to be found in a Heading is the state of the individual at the time the event was experienced. Internal states, such as mood or intoxication have been shown to affect the recall of laboratory-learned materials as well as autobiographical materials (Bower, 1981; Teasdale, 1983). We take this as evidence that state-dependent effects are the result of the internal environment being represented as a feature of the Heading. However, these effects are restricted to recall. Godden and Baddeley (1980) found no context dependency effect with recognition memory. This dissociation has been reported by other workers such as Smith, Glenberg and Bjork (1978) who varied the experimental room between presentation and test. In addition, Eich (1980), in a survey of the effects of alcohol, concluded that the state of intoxication, while influencing recall, virtually never led to effects when recognition tests were used. Evidence like this suggests the conclusion that a Heading may consist of at least two components, one of which is the interal and/or external state, the other being a literal representation of the nominal stimulus. When the internal environment present at the time of learning or encoding is reinstated at test, search of the relevant Heading is more successful and so the likelihood of retrieval of the appropriate Record is increased.

We can now summarise our conclusions with respect to the structure of Headed Records. Records are independent units that can only be accessed through the associated Headings. The Headings consist of at least two components, themselves possibly subdivided. Of these components, one represents the state (either the internal or external environment, or both of these) and the other represents the item in its stimulus form. The associated Record would include a processed version of the item and situational information. In recall, then, if there were a change of state, neither of the components of the Heading would be available as a retrieval cue. In recognition experiments the stimulus item itself should be sufficient to address the appropriate Heading. In this way the state variable would have no influence.

However robust "recognition-superiority" may be, there are some cases in which levels of recall exceed those of recognition. The classic demonstration of this switch is the encoding specificity phenomenon (Tulving and Thomson, 1973). These authors showed that subjects who had been presented with word pairs such as TRAIN-BLACK, subsequently failed to recognise BLACK as having occurred previously in the experiment when it was presented in isolation or with a highly associated word like WHITE. Nonetheless, they could successfully recall BLACK when given TRAIN as a cue. Tulving and Thomson give an informal and general account of the principles they suppose to underly such phenomena:

> specific encoding operations performed on what is perceived determine what is thought, and what is thought determines what retrieval cues are effective in providing access to what is stored. (Tulving and Thomson, 1973, p. 369).

In our approach, where the relation between the Heading ("what is thought") and the Description (the "retrieval cues"), the spirit of the above quotation is clearly preserved. To be specific, the encoding specificity data could be accounted for quite simply by assuming that the critical H-R has TRAIN in the heading and BLACK, together with some situational indices, in the Record. Remembering that this Record is not structurally linked to any other Record concerning BLACK, a fundamental assumption in our approach, it will be seen that we can easily take on board this set of data which creates great problems for any undifferentiated network model.

4. Conclusions

In the preceding section we have chosen a variety of experimental areas within which to illustrate the use of the Headed Records model. These areas represent only a subset of those to which the model can in principle be applied.

In recent years there has been an increasing attempt to forge a relationship between investigators concerned with the normal processes of memory and those neuropsychologists working on amnesia (e.g. Cermak, 1982). This has not been too successful from the point of view of the neuropsychologists owing to, according to Cermak, "the theoretical base upon which we are operating" (Cermak, 1982, p. 374). In the Cermak volume the theoretical vocabulary is drawn from the literature on short term memory and learning rather than long term memory. The only exception is an attempt to impose the episodic-semantic distinction onto the residual amnesic abilities, but, as Baddeley (1984) has pointed out, this has a strong element of circularity about it. Certainly no-one attempts to use any of the available schema models, nor do associative networks get any more than a passing gesture. In addition there is a tendency to look for a single cause of amnesia, at least on the basis of the papers in the Cermak volume. This contrasts with current work in other clinical fields, for example, in dyslexia where it is accepted that patients have multiple functional lesions (Patterson, 1981; Shallice, 1981).

We believe that the Headed Records model should be able to address the problem of amnesia, and an attempt will be reported elsewhere (Morton and Bekerian, note 4). Our strategy is the same as that used with the dyslexias, namely to identify the requirements posed by each task presented to the patients, translate these requirements into the terms of the model and then

isolate the potential loci that could lead to the observed deficits. If this exercise forces changes in or extensions to the model as we have presented it above, so much the better.

We have tried to present a model that is complex enough to handle the variety of natural and experimenter-produced phenomena of memory. Our focus, relative to other available models, has been on the nature of the processes that service and use the relatively simple and general system of Headed-Records. The generality of the current for formulation will be replaced by more specific proposals concerning the nature of these processes as we give a more and more detailed account of the data. Thus far it has enabled us to describe a variety of phenomena with a common language and it has proved productive in suggesting experiments. Its survival depends upon its continued utility.

References

Alba, J.W. and Hasher, L. (1983) Is memory schematic? *Psychol. Bull., 93*, 203–231.

Anderson, J.R. (1976) *Language, Memory and Thought*. Experimental Psychology Series. Erlbaum.

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

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

Baddeley, A.D. (1984) Neuropsychological evidence and the semantic/episodic distinction. *Behav. Br. Sci., 7*, 238–239.

Baddeley, A.D., Cuscaro, W.J., Egstrom, G.H., Weltman, G. and Willis, M.A. (1975) Cognitive efficiency of divers working in cold water. *Hum. Fact., 17*, 446–454.

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

Bekerian, D.A. and Bowers, J.M. (1983) Eyewitness testimony: Were we misled? *J. exp. Psychol. Learn. Mem. Cog., 9*, 139–145.

Bower, G. (1981) Mood and memory. *Amer. Psychol., 36*, 129–148.

Bowers, J.M. and Bekerian, D.A. (1984) When will post-event information distort eyewitness testimony? *Int. J. appl. Psychol., 69*, 466–472.

Cermak, L.A. (ed.) (1982) *Human Memory and Amnesia*. H. Jersey, Erlbaum.

Christiaansen, R. and Ochalek, K. (1983) Editing misleading information from memory: Evidence for the coexistence of original and post-event information. *Mem. Cog., 11*, 467–475.

Clarke, R. and Morton, J. (1983) Cross modality facilitation in tachistoscopic word recognition. *Q. J. exp. Psychol., 35A*, 79–96.

Craik, F.I.M. and Kirsner, K. (1974) The effect of speaker's voice on word recognition. *Q. J. exp. Psychol., 26*, 274–284.

Eich, J.E. (1980) The cue-dependent nature of state-dependent retrieval. *Mem. Cog., 8*, 157–173.

Geiselman, R.E. and Bjork, R.A. (1980) Primary versus secondary rehearsal in imagined voices: Differential effects on recognition. *Cog. Psychol., 12*, 188–205.

Godden, D.R. and Baddeley, A.D. (1975) Context-dependent memory in two natural environments: On land and underwater. *Brit. J. Psychol., 66*, 325–332.

Godden, D.R. and Baddeley, A.D. (1980) When does context influence recognition memory? *Brit. J. Psychol., 71*, 99–104.

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 HFO71*, IBM UK Laboratories, Hursley Park.

Hammond, N.V., Morton, J., Barnard, P.J., Long, J.B. and Clark, I.A. (1984) Characterising user performance in command-driven dialogue. *Hursley Human Factors Report HFO93*, 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.

Hintzman, D. and Block, R. (1973) Memory for the spacing of repetitions. *J. exp. Psychol., 99*, 70–74.

Huppert, F. and Piercy, M. (1976) Recognition memory in amnesic patients: Effects of temporal context and familiarity of material. *Cortex, 12*, 3–20.

Jacoby, L.L. and Witherspoon, D. (1982) Remembering without awareness. *Can. J. Psychol., 36*, 300–324.

Jones, G.V. (1983) Structure of the recall process. *Phil. Trans. Roy. Soc. Lond. (Ser B), 302*, 373–385. (Reprinted in D.E. Broadbent (ed.) *Functional Aspects of Human Memory*. London, The Royal Society).

Jones, G.V. (in press) Fragment and schema models for recall. *Mem. Cog.*

Kirsner, K. (1973) An analysis of the visual component in recognition memory for verbal stimuli. *Mem. Cog., 1*, 449–453.

Kolers, P.A. (1979) A pattern analyzing basis of recognition. In L.S. Cermak and F.I.M. Craik (eds.) *Level of processing in human memory*. Hillsdale, NJ, Erlbaum.

Loftus, E.F. (1975) Leading questions and the eyewitness report. *Cog. Psychol., 7*, 560–572.

Loftus, E.F. (1979a) *Eyewitness Testimony*. Cambridge, MA, Harvard University Press.

Loftus, E.F. (1979b) Reactions to blatantly contradictory information. *Mem. Cog., 7*, 368–374.

Loftus, E.F., Miller, D.G. and Burns, H.J. (1978) Semantic integration of verbal information into visual memory. *J. exp. Psychol.: Hum. Learn. Mem., 4*, 19–31.

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

McCormack, P. (1972) Recognition memory: How complex a retrieval system. *Can. J. Psychol., 26*, 19–41.

Morton, J. (1970) A functional model for memory. In D.A. Norman (ed.) *Models of Human Memory*. New York, Academic Press, pp. 203–254.

Morton, J. and Patterson, K. (1980) A new attempt at an interpretation, or an attempt at a new interpretation. In M. Coltheart, K. Patterson and J. Marshall (eds.) *Deep Dyslexia*. London, Routledge & Kegal Paul.

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

Norman, D.A. and Bobrow, D.G. (1979) Descriptions: An intermediate stage in memory retrieval. *Cog. Psychol., 11*, 107–123.

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

Patterson, K.E. (1981) Neuropsychological approaches to the study of reading. *Brit. J. Psychol., 72*, 151–174.

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.

Shallice, T. (1981) Neurological impairment of cognitive processes. *Brit. Med. Bull., 37*, 187–192.

Smith, S.M., Glenberg, A.M. and Bjork, R.A. (1978) Environmental context and human memory. *Mem. Cog., 6*, 342–353.

Standing, L. (1973) Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology, 25*, 207–222.
Teasdale, J.D. (1983) Affect and accessibility. In D. Broadbent (ed.) *Functional Aspects of Human Memory*. Cambridge, Proceedings of the Royal Society.
Tulving, E. and Thomson, D.M. (1973) Encoding specificity and retrieval processes in episodic memory. *Psychol. Rev., 80*, 352–373.
Underwood, B.J. (1969) Attributes of memory. *Psychol. Rev., 76*, 559–573.
Warren, C. and Morton, J. (1982) The effects of priming on picture recognition. *Brit. J. Psychol., 73*, 117–129.
Warrington, E. and Weiskrantz, K. (1973) An analysis of short-term and long-term memory defects in man. In J.A. Deutsch (ed.) *The Physiological Basis of Memory*. New York, Academic Press.
Williams, M.D. and Hollan, J.D. (1981) The process of retrieval from very long term memory. *Cog. Sci., 5*, 87–119.
Winocur, G. and Kinsbourne, M. (1978) Contextual cueing as an aid to Korsakoff amnesics. *Neuropsychologia, 16*, 671–682.

Reference notes

1. Bekerian, D.A. and Morton, J. Could you remember if I said ... Unpublished manuscript.
2. Bekerian, D.A. and Bowers, J.M. Retrieving accurate memories. Manuscript submitted for publication.
3. Bekerian, D.A. and Mingay, D.J. Second-guessing when you're wrong. Manuscript submitted for publication.
4. Morton, J. and Bekerian, D.A. Amnesia: a Headed Records analysis. Manuscript in preparation.

Résumé

Les auteurs proposent que notre mémoire soit faite "d'enrégistrements" (records) isolés, non reliés entre eux. A chacun de ces enrégistrements serait attaché une "en tête" (heading). On ne peut retrouver un "enrégistrement" qu'en se référant à cette "en tête" dont les contenus eux-mêmes ne sont pas évocables. Chaque "en tête" est un mélange de contenu dans une forme plus ou moins litérale et de contexte, ce dernier incluant une spécification de l'environnement et des états internes (humeur, modifications dues aux drogues). Cette conception de la mémoire permet de rendre compte de nombreux phénomènes mnémoniques aussi bien que de travaux de laboratoire tels que ceux portant sur les différences entre le rappel et la reconnaissance.

On propose dans cette théorie que les "enrégistrements" à "en tête" ne puissent être ni effacés ni modifiés. Les données allant à l'encontre de cette hypothèse peuvent s'expliquer en termes de processus de recouvrement.