In: ..dies in Long Term Memory
(Eds. Kennedy & Wilkes)
John Wiley & Co., 1974. pp. 89-101.

Chapter 5

Structuring Experience—Some Discussion Points

John Morton

General Background

It is clear that it is no longer profitable to talk about 'Long-Term Memory' as though it were a single system operating on a single principle with a common code. From the kinds of errors (or, more impressively, the lack of errors) I make when reciting well learned speeches or singing a song, I am forced to conclude that they are coded in my memory in some literal (e.g. phonological) form. When I recall the contents of a scientific paper, whatever I do there is little or no literal recall, and the errors reflect the amount of reconstruction and assimilation involved at some stage between the reading and the recall. When I describe the route from Cambridge to Steeple Morden (once travelled, recently) I draw the information largely from visual images and cannot repress the image of the car I was travelling with breaking down on the way. The more familiar route to Oxford, on the other hand, has as its first expression the simple verbal form 'Bedford, Buckingham, Bicester'.

Some groups of items are held in ordered lists which usually reflect a natural or conventional ordering such as the names of the months or the four gospels; other groups, such as Shakespeare's plays or the seven dwarfs, although small, are very labile in order and ease of recall. If we do have lists of such items their organization must be very different from the ordered lists. This reflects the way they are learned and is reflected in the more complex way in which they are retrieved.

Currently there are attempts to formalize a division of long-term memory into two parts. Tulving (1972) proposes a distinction between Semantic Memory and Episodic Memory. Bruner (1969) makes a related distinction between 'memory with record' where specific events are recoverable and 'memory without record'. In the latter case experience is incorporated into some larger structure changing, for example, the rules by which the organism operates, 'but which are virtually inaccessible in memory as specific encounters.' There is, of course, also the possibility of a clear record of an event together with an abstraction based on the event with or without any indication of the relationship between the event and the abstraction. Certain authors have also noted the case where the record of the event is singularly inaccessible but the abstraction plays a considerable role in influencing behaviour. Another example of a division is Mandler's (1967) separation of a buffer memory from long-term storage. Related to this is the division into Lexicon and Topicon suggested by Smith and Claxton (1972). Such divisions incorporate a number of factors including:

Formal (or shared) knowledge *vs* idiosyncratic experience.
Permanent (relatively) storage with limited access *vs* limited storage for currently relevant material with much more free access.
Time independent knowledge *vs* time coded knowledge.
Linguistic information *vs* knowledge of the world.

While suggesting the Semantic–Episodic division, Tulving also stresses the interdependence of the two systems and other workers are confirming what a number of linguists had previously discovered—that a reasonable semantics for a language system has to include knowledge of the world. Thus, Collins and Quillian, who started off with a fairly simple, hierarchial model of semantics with fairly limited methods of operation have had to complicate their descriptions of how people operate in verifying simple propositions (1972). Winograd (1971), Rumelhart, Lindsey and Norman (1972) and Norman (1972) give other evidence favouring the breakdown of sharp divisions.

Local Background

My own interest in this area, apart from the general problem of the language user (Morton, 1968), springs from some problems in applied psychology. The specific issue was the design of computer systems for patients' records using a visual display and a keyboard. As the nurses would be expected to operate the system in addition to their other duties it is not feasible to design a command system involving special syntax and mnemonics. Such systems are suitable for full-time operators, (such as airline booking clerks), but the likelihood of forgetting the correct forms or making a keying error is too high for them to be acceptable to nonprofessionals. Instead, then, the system has to be organized in terms of successive multiple choices. Thus at one stage in the interaction the user might see on the screen the sequence of instructions shown in Figure 5.1.

Suppose the user wishes to enter the information that a patient has recovered sufficiently to have solid foods. She would then press key number 4 and proceed to the next level. If the information to be entered was that the patient had recovered sufficiently to bath himself however the choice is no longer obvious. If the same patient can now also go to the lavatory by himself, can get up for visitors or go into the TV room then the current focus in the users' mind would be Mobility. If, on the other hand, she is in the process of revising the *bathing*

Nursing Procedures

1. Bathing	5. Oral Hygiene
2. Mobility	6. Lavatory
3. Care of Pressure Areas	7. Dressings
4. Feeding	8. Special Treatments
9. Investigations	

Figure 5.1. A display which might occur in the middle of an interaction between human and computer where the latter is the repository of patients' records

status of all the patients on the ward then the natural choice would be key 1. A strict single hierarchy would be very inconvenient in this situation.

In one operating system, at Kings College Hospital in London from which the above example is taken, problems of this kind were reduced by involving the nursing staff in the design of the system. The system nevertheless took some time to evolve. The end product can be described by a directed graph whose organization reflects the organization of the information in the user. What we would like to do is find experimental procedures which will readily reveal at least the major part of these structures.

Another problem is that the internal organization is likely to be different for different groups of people. Thus, in a hospital, nurses will be likely to classify patients with tonsilitis and appendicitis together in contrast to throat cancer and prostate operations since the former need little nursing and the latter pair more intensive nursing. For medical staff, on the other hand, it would be more natural to classify the tonsilectomy and the throat cancer patients together and the prostate and appendix patients together on the basis of the parts of the body concerned.

A more dramatic example of the same principle can be seen in the context of a large store. Given that there are too many departments to display on the screen at the same time then some superordinate headings have to be devised. For inter-departmental communication it seems likely that the most appropriate headings would reflect the physical layout of the store since this is the classification with which people in the store are used to working. For customers however, ordering goods via their computer link, such a structure would not be very useful and headings such as 'food, clothing, hardware, garden, kitchen . . . ' would be appropriate as a first-level classification.

Our practical problem, then, is that of devising methods for discovering the principles of organization in memory. There follows preliminary reports on the results of three enquiries.

Colours

Informants, mainly members of staff of the APU, were asked to 'give as quickly as possible as many names of colours you can think of'. This they did with their eyes closed. The spoken lists were taken down in shorthand, and the informants were stopped when they paused for 20 seconds or more, when they made some comment like 'do you want details/sub-divisions?' or 'I don't know any more' or when they started to produce compound names such as 'Royal Blue' or 'Shocking Pink'. When they had finished they were asked if they knew the 'colours of the rainbow'. Of 33 people asked, 18 were able to recite the seven in the proper order. Of these only five had produced the seven in order in the main task. The other 13 either omitted one item or had some colours in a different order from the spectral one. Of the 15 who could not recite the spectrum colours, only one mentioned them all in the main task. Some of these people knew mnemonics for the spectrum, such as 'Richard of York Gains Battles In Vain', 'Read Only Your Garden Book In Verse' or 'Very Idle Blondes Get You On Retreat'. However, it was clear that they only had access to the colours sequence through the mnemonic and not as an independent list.

The performance of the people who could recite the spectrum must be interpreted with care. Some claimed they were not using the spectrum list in the main task but were using some other strategy, or no strategy at all. A list such as RGBOYVI could be produced by a combination of retrieving individual colour names plus inter-colour associations without reference to the spectral list. In the other cases we could suppose that the list was retrieved from store correctly but output incorrectly owing to the pressure for speed. Still, it is unlikely that the alphabet, the months or the four gospels would be subject to such influences. This was checked by an enquiry in which 16 people were asked to 'recall as quickly as possible, not necessarily in the correct order the books of the New Testament'. Twelve produced 'Matthew, Mark, Luke, John' immediately. The other four blocked (i.e. could think of no New Testament books for about 30 seconds) but continued in the usual way when cued with 'Matthew'. There are a number of reasons why this might be so—e.g. amount of experience, length of list and extent of inter-item associations and classifications between items which are non-adjacent and on the list. All we can say is that the spectral list, when it is known, exerts strong influence on recall of colour names but that it does not dominate.

Turning now to colours outside the spectral list we can find other differences between the groups, which seem to illuminate the recall strategies. The two groups produced about the same average number of non-spectral names, 5·2 for the spectral, (*S*) group and 5·4 for the Non-spectral (*N*) group. They differed, however, in respect of the triple Black–White–Grey, and in the three most frequent (by Thorndike–Lorge) non-spectral colours, Brown, Pink and Purple. The *N* group named an average of 4·0 of these six, the *S* group only 2·2. This difference was tested by calculating for each subject the difference between the number of this set which was mentioned and the number of other non-spectral colours mentioned. The distributions of these differences for the two groups were compared using a Mann–Whitney which gave $z = 3{\cdot}032$ ($P = 0{\cdot}0012$). Some of the *S* group explicitly rejected Black–White–Grey as not being colours (though if one was mentioned it was recorded). However, as performance on Brown, Pink and Purple was equivalent to that on the other names the author is inclined to minimize the influence of such internalized theories. Instead I wish to suggest that a word-frequency effect operated for the *N* group, but the *S* group, being committed to a general strategy of retrieval from LTM, effectively dampened the word-frequency effect. Such a suggestion implies a multi-component model of word retrieval, with colour names being retrieved explicitly from LTM as a result of some strategy or being produced by some more peripheral mechanism, such as the Logogen System (Morton, 1969) as a result of a semantic context 'colour'. The latter method would be expected to show the word frequency effect.

Store Departments

Having claimed blithely that people who work in a store have a model of the store in their heads it was thought useful to present some data backing up the claim. Pennie Ottley persuaded one of the stores in Cambridge to cooperate, and interviewed 29 people who worked there. These interviews took place in the departments where they worked. The informants were presented with a piece of paper and asked:

> 'Please write down as quickly as possible the departments you would expect to find in a large store. Be fairly specific—for example you might divide "sports equipment" into "Indoor Sports" and "Outdoor Sports".'

These instructions were worded to minimize the likelihood of us being accused of having fixed the data by deliberately biassing the informants, and any questions about the detail of the task were answered guardedly. The 'Sports' example was chosen to make sure we got the right level of detail—the store we used had no Sports department. Subsequently 33 female members of the APU subject panel were asked the same questions with the name of the store we had used thrown in as an example.

We expected to find two clear influences in the lists from the store workers. One of them would be the physical layout of the store; the other would be a classificatory influence, with departments related to clothing all grouped together, for example. These two classifications were related, as in most stores, but we expected the physical layout to reveal itself unambiguously.

In fact we were wrong. The only clear influence of the working situation was that 11 of the 29 people named their own departments first. From then on most people looked as though they were going on a random walk. If we analysed the data in complete detail we would probably be able to show that the store workers produce significantly more pairs of names of physically adjacent departments than do the control subjects, but that was not the point. We thought we were demonstrating something obvious. Several people have been at pains to point out that for a task like the one we gave there are many possible retrieval strategies and influences. All we established was that the other factors, whatever they were, were stronger than that of the layout of the store. Semantic groups were apparent, but what is one to make of a sequence like . . . TOYS, FASHIONS, TOOLS, PERFUMES, CAMPING . . . ? The answer is probably 'nothing'.

Although we have not yet run any further tests on our informants, we have no reason to believe that the subjects had no knowledge of the layout. It was just that we asked the wrong question to reveal that knowledge. Yet there were some appreciable differences between the two groups of subjects. One of them was in the number of departments mentioned which were not present in our store. We checked this by pairing each store worker at random with a control subject who had produced the same number of departments in her list plus or minus one. We were able to match 19 pairs in this way. A Wilcoxon test on the number of extra-store departments gave a $T = 21$ ($P < 0{\cdot}005$). The reason was clear—the control subjects used other stores as models: the store-workers in general only use their own, in whatever way.

The most common of the extra-store departments were the 'sports' departments which had been given as an illustration in the instructions to the subjects. It might seem that this is a simple case of priming. However, 11/29 of the store workers and 5/33 of the controls did not mention sports, and none of the subjects mentioned it first. For the control subjects the sports departments were evenly distributed through the list; for the store workers there was a strong tendency for them to be at the end of the list, (13/18 put it in the second half of the list, $P = 0{\cdot}048$; Binomial test). There are many possible reasons for this but we can be sure that models of the process which predict simple priming effects would have problems.

Food

Lists are sometimes more directly revealing of multiple strategies. Thus if you ask people to produce a shopping list for a week's holiday in a country cottage, sequences appear which unambiguously reflect a variety of influences such as:

Dairy goods— . . . eggs, cheese, lard, butter . . .
Desserts— . . . tinned fruit, fresh fruit, jellies, tinned milk, rice pudding . . .
Breakfast— . . . bacon, porridge, cornflakes, milk . . .
Grain products— . . . flour, rice, barley, sago, spaghetti . . .
Cake making— . . . lard, margarine, plain flour, self-raising flour, currants, sultanas . . .
Physical form— . . . tinned soups, tinned meats, tinned fish, tinned fruit . . .

Existence of multiple classifications equally clearly leads to chaining and interleaving of strings. Some examples are given in Table 5.1.

It is not clear what might be made out of these data. Examples can be found to illustrate the influence of virtually any classification one can think of. One exception to this is the contents of the refrigerator, which none of the 22 protocols we have looked at illustrate. No doubt people do know what is in the refrigerator so we must assume that, as with the department store, some kinds

Table 5.1. Multiple classification of food items
Examples

A	cream	dairy (cream–margarine)		
	milk			
	butter			
	lard	cake making (lard–sultanas)		
	margarine			
	plain and S.R. flour			
	currants			
	sultanas			
B	butter	Dairy		Fats
	margarine	Dairy		Fats
	cooking oil			Fats
	eggs	Dairy	B'fast	
	bacon		B'fast	
	cheese	Dairy		
	tea		B'fast	
	coffee		B'fast	
	Marmalade		B'fast	
	jam		B'fast	

of retrieval strategy are less efficient than others in producing responses. It is worth pointing out, however, that when my wife gives me a shopping list it contains logical sequences such as 'cream, orange juice, beef, eggs, spaghetti, apples'. The key to the sequence is the layout of the supermarket round the corner, and the items are generated in the order in which they will be picked off the shelves. This, of course, is a deliberately adopted strategy with a rigorous exclusion from *response* sequence of any out-of-order items. Any model of the list generation process would have to take such factors into account.

It seems likely, then, that a study of shopping lists will tell us more about the nature of the retrieval mechanism than about the organization of the data-base. However Dick Byrne and the author have found another kind of food list that displays different properties: a list of ingredients of a dish. Such lists can readily be generated by people who know how to make the dishes. The problem is discovering how such lists are stored. One likely method is to have a copy of the ingredients as they are listed in a recipe book. One wouldn't expect such a list to be more strictly ordered than the spectrum colours and one would expect priming to be effective. So, if you ask someone what are the ingredients of lemon meringue pie or sherry trifle you might expect the first responses to be 'lemons' and 'sherry' respectively. And you would be wrong. Of 10 subjects asked the ingredients of lemon meringue pie only one said 'lemons' first. The others all began with 'flour', or 'fat'. The reason is quite plain—there is no list of ingredients as such. What happens is that the housewives used as informants retrieve the program for making the dish. They run through this program and extract for verbalizing the ingredients required by the experimenter. They all agree that this is indeed what they do, and this introspective conclusion is supported by two kinds of evidence. First there is a very high correlation (usually unity) between the order of the ingredients and the order in which the ingredients appear when the same people are asked to produce a recipe (at a second interview). Secondly, pauses occur as the lists of ingredients are produced, at the natural gaps in the recipe—corresponding to the pastry, the lemon filling and the meringue. It seems likely that the program for the pastry and meringue sections at least are not stored with the lemon meringue pie recipe but rather have to be retrieved from elsewhere. Listing the ingredients is thus a process which is at least three levels deep.

The list thus turns out to be fairly strictly ordered, but only as a consequence of the ordering of the operations involved in the manufacture. This being so it is not surprising that there was no priming from the name of the dish. The one informant who did give 'lemons' first in her list mentioned them half way through the recipe. When she had finished the recipe she was asked: 'At what point would you actually *grate* the lemons and squeeze them?'. 'Well, I think I'd do it before I began, because I don't like doing things like that in the middle.'

Discussion

The results of the enquiries outlined above lead us to fairly simple conclusions. Basically they show that the kind of organization revealed by an experimental situation is heavily determined by the precise nature of the task. Information concerning a particular topic can emerge in a number of ways and the strategy adopted to start with affects the whole performance. Thus, by the present account, one can start to retrieve colour names either by using 'scientific' knowledge or by relying on a semantic probe into a word store. If the former operates, as it nearly always does when such knowledge is directly available, then the latter method seems to be ruled out. When store workers are asked to list departments they seem not to use what would probably be the most efficient method—that of using special information. Instead they use a mixture of strategies including (we presume) personal experience as well as semantic links. This then prevents any subsequent attempt at a systematic search from becoming effective.

Such accounts make it clear that the form of the organization and the nature of the retrieval mechanism cannot be studied in isolation. In addition, the degrees of freedom open to us in designing a retrieval mechanism are currently greater than those available in describing alternative forms of organization. The likelihood that the retrieval mechanism operates in different ways in different situations means further that unless we adopt the right level of model we will find ourselves with apparently contradictory results. The author's current view is that we will have to conceptualize retrieval in terms of a problem solver.

The appropriate level of description will be to specify the classes of operation that are possible. Restrictions in the operation of the problem solver may take the form of a limited capacity available for current information. Such a limitation may manifest itself in certain kinds of grouping during output (see

Broadbent, this volume) but also, it is to be expected, in other ways. For example we might expect a limitation in the depth to which a search can proceed down some tree-like organization without the speaker losing his place. Alternatively one or more of the current goals of the task might be forgotten. A good example of this occurs when subjects are asked to introspect about an ongoing task—such as generating a menu according to certain restrictions. At certain stages in the proceedings subjects often cease introspecting. We would expect such an event to coincide with some activity which occupies a large number of slots.

Another factor we will eventually have to contend with is the difference between what might be called *directed* and *autonomous* search. The phenomenon of autonomous search will be familiar to most people. If, for ecample, you try to list the contents of your house by adopting a deliberate strategy, such as going from room to room, items will keep on intruding from other rooms. Calling such intrusions 'simple' associations, as is tempting, doesn't really add anything to our understanding of the situation. More dramatic examples will doubtless occur to most people. My most recent example followed a conversation in which parts of the body (separated) were being discussed. Van Gogh's ear was mentioned which led me to remember, from many years ago, that a war with Spain had once been precipitated by the removal of an ear from the head of an English sea captain. His name, however, eluded me and the rest of the company. Three days later, in some of the same company, 'Jenkins', the name we had been searching for, came to my lips apropos of nothing in the conversation. While it would be unreasonable to expect a coherent account of such events at the moment our models should allow for them at least in principle (cf. Shallice, 1972).

Introspection Level

If we are going to understand complex processes we need to have complex data. Mere listings of output items are insufficient. Not only do we need timing information but we also need to have introspections from our subjects. Only by asking them can we discover the way in which they see a test or the particular idiosyncratic strategies used. Collins and Quillian (1972) have recently shown how introspections can cast light on the way in which subjects tackle a seemingly simple task. Two important questions which need to be tackled are the extent to which introspections accurately reflect the underlying processes and the extent to which introspecting interferes with ongoing processing. If the idea of a limited capacity working memory is correct then we might expect the goal INTROSPECT to reduce measurably our information handling capability. Newell and Simon (1969), however, claim that producing protocols makes no difference to performance in cryptarithmetic tasks. Perhaps the nature of the task makes a difference. We need not expect that a difficult task which uses a small data-base, such as cryptarithmetic, would be subject to the same influences as an easy task (in terms of the underlying logic) which uses a large data-base, such as designing a menu.

The question of the accuracy of introspections is more complicated. A group of us at the APU are currently trying to tackle this problem. In the course of developing a program to simulate a housewife producing a menu it became apparent that we had a number of options open to us with regard to what the program output. Suppose we take a very simple program. The data-base consists of lists of dishes under the headings COURSE 1, COURSE 2 and COURSE 3. Under COURSE 1 will be found the list SOUP, OYSTERS, MOULE, SALAD, ROLLMOP, PATE. The COURSE 2 list is STEW, SPAG, OMLET, COQ, SOLE, CURRY and COURSE 3 is ZAB, YOG, ICE, BRIE, CHEDDAR. Each dish is also the name of a list of features. Thus the OYSTERS List includes WHITE, POSH, FISH and MOULE is CHEAP, POSH, FISH, WET.

Goals are entered into the program which then simple-mindedly proceeds to find a meal to satisfy the goals. It does this by taking the first item in COURSE 1 and testing it successively against all the current goals. When a test fails the next item in COURSE 1 is tested and so on until an item is found to satisfy all the goals. Then the same procedure is followed for COURSE 2 and COURSE 3.

The program is written in a local list processing language based on TRAC macros. As the language is unprincipled it is called Simulation Language Under Trac or SLUT. A flow diagram of the program is shown in Figure 5.2. Its mode

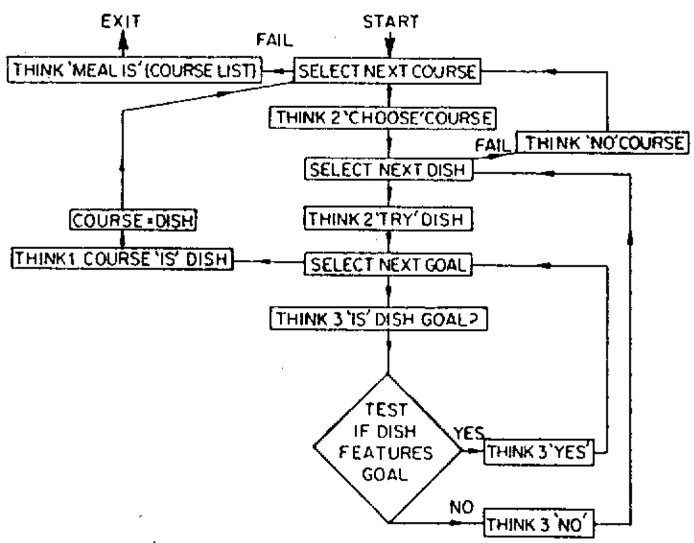

Figure 5.2. Flow diagram of program for selecting a three course meal with given constraints. An 'introspection level' is set before the program is run such that any THINK instruction results in the message attached being printed only if the suffixed digit is lower than the introspection level. Thus the external evidence for the nature of the program's operation differs as a function of the introspection level

of operation is fairly clear from the description. The one unexplained feature is the THINK function. This is an optional PRINT instruction controlled by a numerical suffix on the instruction in the program and an I-LEVEL (introspection level) figure entered by the operator before execution or dynamically. If the THINK suffix is equal to or less than the value of I-LEVEL then there is an output from the program. Output for I-LEVELS of 0 to 3 are given below for GOALS of CHEAP and POSH.

I-level = 0

```
MEAL is MOULE COQ YOG.
```

I-level = 1

```
COURSE 1 is MOULE
COURSE 2 is COQ
COURSE 3 is YOG
MEAL is MOULE COQ YOG.
```

I-level = 2

```
CHOOSE COURSE 1
TRY SOUP
TRY OYSTERS
TRY MOULE
COURSE 1 is MOULE
CHOOSE COURSE 2
TRY STEW
TRY SPAG
TRY OMLET
TRY COQ
COURSE 2 is COQ
....
```

I-level = 3

```
CHOOSE COURSE 1
TRY SOUP
IS SOUP CHEAP? YES
IS SOUP POSH? NO
TRY OYSTERS
IS OYSTERS CHEAP? NO
TRY MOULE
IS MOULE CHEAP? YES
IS MOULE POSH? YES
COURSE 1 IS MOULE
CHOOSE COURSE 2
TRY STEW
....
....
....
```

The particularly instructive aspect of these outputs follows from playing the game of supposing they are separate introspections from four informants and trying to formulate a model for their production. We might deduce from I-level 0 that the whole meal was accessed directly with a CHEAP–POSH search feature (some of the subjects tested by Dick Byrne do in fact claim they do this, substantiating the claim with details of the guests last time the meal was served). Something like Tulving's (1972) Episodic Memory might be brought to bear as an explanatory construct. From I-level 1 we might conclude a direct access to POSH–CHEAP dishes with the courses treated in series. I-level 2 suggests parallel testing of the features for individual dishes in contrast with I-level 3 where the features are clearly tested in series.

When we then learn that all the outputs were produced by the same program we might be surprised. The lesson seems to be that on the basis of a protocol we can eliminate certain kinds of model. Positive conclusions, particularly concerning possible parallel search, can be dangerous.

A useful distinction can be drawn between real-time introspections and *post-hoc* introspections. The examples given above are of the former type. Suppose we ask our subjects what they did in solving the task. Sometimes it is clear that they have no more information available than the experimenter. On other occasions one can get quite complex, but unsubstantiated descriptions of the underlying operations. In an attempt to understand how such descriptions may be arrived at we can imagine that we are clever enough to have as part of the SLUT system a resident program to which the problem-solving section communicated. The former program would contain a number of inferential processes enabling it to draw conclusions about the operation of the latter. These conclusions could then be part of the output. Their validity, however, is simply a function of the appropriateness of the inferential procedures.

Individual humans differ widely in the extent to which they spontaneously introspect (when not explicitly asked to) and in the extent to which they can introspect when asked. We are only beginning to understand the problems of using such data but feel sure that with continued analysis they can be overcome.

Practical Conclusions

The problem of man-computer systems has become lost from view. We have been on a search for techniques which would enable us to understand the principles or organization of our knowledge. This led to considerations of the nature of the data we can collect. Perhaps more important than trying to find *the* best principle of organization is the task of discovering the adaptability of the human operator. Faced with a particular system, what kinds of discrepencies are likely to create serious problems? What kinds of interference are likely between a prior structure and one imposed by a system? How are new structures incorporated into memory and how should they best be presented for learning? And the trouble is that we need the answers now. Systems are being implemented and if we don't know the right way to do it, nobody does.

Acknowledgements

I am grateful to a number of people for contributions of time and thought to the material in the paper. Notably Pennie Ottley, John Fox, Dick Byrne and members of the THLUTS group.

References

Bruner, J. (1969). Modalities of memory. In G. A, Talland and N. C. Waugh (Eds) *The Pathology of Memory*. New York: Academic Press.

Collins, A. and Quillian, M. R. (1972). How to make a language user. In E. Tulving and W. Donaldson, (Eds) *Organization and Memory*. New York: Academic.

Mandler, G. A. (1967). Organization and memory. In K. W. Spence and J. A. Spence (Eds) *The psychology of learning and motivation*, Vol. 1. New York: Academic Press.

Morton, J. (1968). Grammar and computation in language behavior. In J. C. Catford (Ed.) *Studies in Language Behavior* C.R.L.L.B. Progress Report No. VI. University of Michigan.

Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, **76**, 165–178.

Newell, A. and Simon, H. A. (1972). *Human Problem Solving*. Prentice-Hall: Englewood Cliffs.

Norman, D. A. (1972). Memory, knowledge and the answering of questions. CHIP 25, N. C. San Diego.

Rumelhart, D. E., Lindsey, P. H. and Norman, D. A. (1972). A process model for long-term memory. In E. Tulving, and W. Donaldson, (Eds) *Organization and Memory*. New York: Academic Press.

Shallice, T. (1972). Dual function of consciousness. *Psychological Review*, **99**, 383–393.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving and W. Donaldson, (Eds) *Organization and Memory*. New York: Academic Press.

Winograd, T. (1971). Procedures as a representation for data in a computer program for understanding natural language. Project MAC report TR-84. M.I.T.