

## HOW EPISODIC IS SEMANTIC MEMORY?

Andrew Ortony  
University of Illinois at Urbana-Champaign

At the end of his Philosophical Investigations Wittgenstein (1963) notes that the confusion and barrenness of psychology is to be explained not by the youth of the science but by the conceptual confusions that underlie it. I think one could easily make a case for saying that questions concerning the distinction proposed by Tulving (1972) between episodic and semantic memory generally involve precisely such conceptual confusions. What I propose to do in this paper is first to illustrate such a confusion, and then to suggest a possible way round it.

Addressing the question, "Is there a semantic memory?", Schank (1974, 1975) argues that there is not. Semantic memory, he claims, has been treated by psychologists as being made up of words and relationships between them, and has been proposed as a means of representing both the meanings of sentences and the knowledge required to understand them. This he believes is wrong; what we should actually do is to separate out a language-independent aspect, call it conceptual memory, distinguish it from lexical memory, and identify it with episodic memory because the relationships between concepts in it "could only have been acquired by personal experience." Consequently the distinction between semantic and episodic memory disappears, which is all to the good since semantic memory is an ill-conceived notion.

The first confusion in this position is factual rather than conceptual. It is simply not the case that psychologists who have worked on semantic memory proposed it as an associative network of words. All the workers cited by Schank explicitly deny this in their work. The distinction which he wants to make, while valid enough in itself, is implicit in all existing models, and is quite explicit in the most recent formulation of Quillian's model proposed by Collins and Loftus (in press), where allowance is made for a distinction between lexical and conceptual associations, just as Schank proposes. It is clearly a mistake to confuse knowledge of words with knowledge of word meanings. In all models the role of the dictionary or lexicon or lexical network is to represent knowledge about words which is not primarily semantic, but, to quote Collins and Loftus: "...is organized along lines of phonemic (and to some extent orthographic) similarity."

Schank believes that semantic memory (excluding lexical memory) doesn't exist, and it seems that one of his chief reasons for this belief is that the contents of semantic memory are acquired through personal experience. Now Schank might be right that the semantic/episodic distinction is not worth very much, is not useful, is misleading or whatever, but that the set of associations and other relations represented in a language-independent memory "could only

have been acquired by personal experience" is quite irrelevant, in a sense, probably false, and certainly does not justify the conclusion that Schank draws, namely that "conceptual memory is episodic in nature."

The distinction envisaged by Tulving was not one which could be invalidated by merely claiming a position of extreme empiricism. Tulving expressed the distinction as that between a kind of mental thesaurus, and an autobiographical record of temporally-dated events. Generally speaking memory modellers have not made much use of the distinction. Quillian (1968, 1969), Rumelhart and Norman (1973) and Anderson and Bower (1973), for example, all seem, in fact, to adopt the position advocated by Schank precisely because they make no sharp distinction between representations of sentence meanings and of world knowledge required for understanding. The most notable exception is Winograd (1972) who maintains a separate record of the discourse to help with handling anaphoric references.

It is widely believed that semantic memory is characterized by rigid hierarchical structures giving special prominence to the notion of superset. Once one makes this assumption it follows rather quickly that there are defects in the notion of semantic memory because it cannot handle a variety of problems which appear to originate from a rigidly hierarchically organized knowledge base. It is, I think, in this line of reasoning that the first confusion hides. Yet, all that follows from, for example, the observation that people cannot rattle off the fifty states of the union is that they don't have all of them equally accessible from the concept of state; it says nothing about semantic versus episodic memory. Nor, incidentally, does it show that the superset relation is useless -- rather, it shows that there may be a limited number of readily accessible inverse superset relations, or exemplar relations.

Schank rightly objects that using a rigidly hierarchical system and storing in it "plunger," "hammer" and "saw" indiscriminately under a node for "tool" fails to represent the greater semantic closeness between "hammer" and "saw" than between one of them and "plunger." His solution is to replace one static representation with another, so that while his representation captures that greater proximity between "hammer" and "saw", it fails to represent the similarity of "hammer" and "plunger" that might be relevant in a context involving tools with (relatively) long handles. If we want to capture semantic similarities of this kind we should not attempt to choose one static representation in preference to another. Rather we should be prepared to allow comprehension and retrieval processes to help out. One such candidate is a spreading activation model (c.f. Collins & Loftus, in press; Ortony, 1975); a model which allows activation from various sources (specifically input string and context) to summate and activate concepts which might

not have been activated above threshold by one source alone. If one makes assumptions such as these about the processes that operate on the structures, it can be readily seen that one need not worry too much about which particular static representation one utilizes in memory. Probably there is little to choose between them. Providing we have a way of representing all the information pertinent to a concept that we need, activation levels, associative strengths and criteriality can be invoked by appropriately defined processes to pull out those relationships which dominate in particular contexts. It might well be that the particular structures that represent what people know depend for their relational characteristics not so much on rules for the construction of canonical forms as on the particular details of the individual's experience. Thus a man who consistently uses forks to open cans might be expected to have his concept of fork unusually closely associated with his concept of can-opener (if he has one). Viewed this way, it would seem that one man's plunger might indeed be another man's hammer; a possibility which should not be ruled out by imposing arbitrary structural constraints.

Being flexible and being arbitrary are not the same thing. Sometimes, of course, one can be both. A case in point is Schank's conception of context. On the one hand it seems to be merely a conceptualization, for the only links permitted within a context are those of conceptual dependency plus the causal links. On the other hand "there is a sense in which all of memory can be considered to be one big context." With a notion of context which is so flexible it is difficult to see what purposes it serves. Schank wants to say that "going to a museum in Berlin" and "going to a museum in Boston" is an example of an intra-context association, while "going to a museum in Berlin" and "going to an old hotel (in Berlin?)" is an example of an inter-context association. The reason is that the former pair are simply examples of "going to a museum" episodes whereas the second pair are only related because of "shared properties." But one might as well argue that "hotel going" and "museum going" are within the same context because they both fall within the general paradigm of "things I do when I'm away from home," or "visiting strange or interesting buildings." There simply are no compelling reasons to make the distinction that Schank is searching for on the basis of context-shifts of the kind he describes. The confusion which I think underlies this reasoning is further exemplified by the analysis of the protocols of his daughter, Hana. Schank wants to argue that Hana doesn't jump around between contexts, yet in the context of birds flying Hana jumps to "plane" because, Schank says, planes fly too, but that is just a case of shared properties on Schank's own analysis, so it must, after all, constitute traversing an inter-context link.

Nobody would deny that there are some very serious problems to be dealt with concerning the structure of knowledge and

the processes that operate on them. There are serious problems concerning context. There is no doubt that in different contexts different relationships and semantic connections between concepts are going to be relevant to comprehension and memory (c.f. Anderson & Ortony, 1975). There is also no doubt, as Woods (1975) points out in his excellent paper, that memorial representations are going to have to be both intensional and extensional [2]. Current models have yet to make this distinction although it is quite obvious that one must keep distinct representations for noun phrases derived from predicates, and for names. Not very recently the distinction between Richard Nixon (extension) and The President of the United States (intension) became quite self-evident. Ortony and Anderson (1974) describe an experiment in which subjects showed a marked tendency to make recognition errors depending on whether the noun phrase in a study sentence required an intensional or extensional interpretation.

Considerations such as these are certainly important in modeling memory, but they do not relate to the semantic/episodic distinction in any obvious way. It is important to know which problems go together, and to know how and why (that is the beauty of Woods' paper). I think that a distinction between episodic and semantic memory is also important, not because of tenuous connections to problems of knowledge representation but because of more global considerations concerning the nature of knowledge. Before exploring this in more detail, however, let us clarify what is supposed to be distinguished from what. A distinction between episodic and semantic memory represents a difference not so much between different kinds of memory but between different kinds of knowledge in memory. It is the contents of memory that we distinguish, not the memory itself.

What I am going to suggest may seem to some not to be a distinction between semantic and episodic memory at all. On the other hand I think it captures some, at least, of what Tulving intended and is a useful and important distinction for AI and psychology alike.

The currently fashionable rejection of the notion of an episodic/semantic memory distinction seems to be based on the idea that since everything we know, we know from experience, there is little point in and no room for the distinction. Here lurks another conceptual confusion -- a confusion of knowledge from experience with knowledge of experience. Philosophers have long argued that although all knowledge might arise from experience it is not necessarily all given in experience.

Consider the difference between a personal diary and an encyclopaedia. Typically an encyclopaedia contains information about a great variety of topics -- or concepts. Information in an encyclopaedia constitutes knowledge about the concepts entered in it and consequently

can be used to make inferences involving them. A personal diary, on the other hand, constitutes a different kind of knowledge -- knowledge about the individual whose diary it is. It would be possible, no doubt, to produce an encyclopaedia from a personal diary, but it would take a lot of work, a lot of inference and a lot of assumptions. What I want to suggest is that the knowledge in semantic memory can be thought of as the kind of knowledge which one would find if one were to produce an encyclopaedia from his personal diary. It is knowledge which has been reorganized around concepts from knowledge originally encoded around events. It has the further characteristic that it is not embedded in propositional verbs; it has been freed from their bonds by using inferential heuristics; it contains knowledge which can be represented by such propositions as "All bachelors are unmarried," "Most swans are white," etc. We all have knowledge of this kind, but we couldn't possibly have acquired knowledge about all bachelors or even most swans by personal experience. Of course, there was a personal experience involved when the putative knowledge was acquired, but precisely this is the distinction we must make. Probably I read in a book that all bachelors are unmarried (a dictionary perhaps). So there was a personal experience to the effect that "I read in a dictionary that all bachelors are unmarried." But we cannot identify the knowledge that we learned x under such and such circumstances, with the knowledge x, stripped of source and circumstances. What our teacher tells us, or what we read only licenses us to claim that we know that our teacher told us, or that we read something -- it does not, alone, license us to know, as true, what it was we heard or read. So, while all input can be tagged as to source, semantic memory somehow represents the untagged information, and that untagged information is not a record of personal experience. It is a record of personal experience stripped of whatever makes it personal. It is derived from personal experience but does not represent it. It represents, as it were, knowledge without the guarantee.

Some years ago BBC TV showed a program in which was described the growing of spaghetti on spaghetti trees in the south of Italy. Most of us would probably reject this notion as an absurd fabrication, but not so for all the viewers; some "took it seriously." A proposition describing the personal experience might be something like "I saw a TV program which alleged that spaghetti grows on trees," which we can write as "I saw S'" where S' itself can be represented as "A TV show allege(d) that S" and where S represents the proposition "Spaghetti grows on trees."

The role of episodic memory in my sense is to store a record of the experience in the general form "I experienced S'". My contention is that believers and non-believers of the spaghetti story alike would have stored some such representation, namely a representation of S' tagged with

autobiographical information such as modality of input, relative time and so on. Except perhaps for acidheads and avid Cartesians, most of us are willing to trust our senses most of the time -- we are willing, that is, to infer S' from "I saw S'", but S' still contains source information. Focus now on the knowledge representable by the proposition, S, "Spaghetti grows on trees." If we were to incorporate this piece of knowledge in the data base it could be used in all kinds of inferences about spaghetti (the crop depends on climate, the cost of harvesting it may vary, it is a fruit or vegetable, etc.). However, in general there is no simple way of inferring S from S' and certainly no such inferences about spaghetti can be made from S' alone. The role I want to attribute to semantic memory is that of a repository of knowledge with the characteristics of S, knowledge, that is, which has been stripped of its autobiographical reference, and its source and circumstances of acquisition information too; always the result of inference.

Those of us who realized that the spaghetti sequence was just another skit from the satirical program Monty Python's Flying Circus did not encode something like S in semantic memory -- at most we encoded or represented something like S' which represents source-free information or knowledge about a TV show but does not represent source-free spaghetti information! Those who were duped, on the other hand, stripped S' of its source and circumstances and represented the knowledge expressed by S. Thus knowledge in semantic memory is derived from, but fundamentally different from knowledge in episodic memory. Putative knowledge which is potentially semantic always arrives embedded in episodic knowledge; it can always be represented as being governed by a propositional verb.

I know a lot of things about England and a few things about Lord Nelson. One of the things about Nelson that I know is that he said "England expects every man to do his duty," but that is not one of things I know about England. England's expectations were either never sufficiently relevant or never sufficiently true for me to have freed them from Nelson's claims. The knowledge that I have about Nelson may be in semantic memory because it is free from the bonds of its source -- in forgetting the source I may have assumed it to be true -- but it only gets into semantic memory if it is assumed to be true (enough).

I earlier urged that semantic memory is centered around concepts and that episodic is centered around events. I think there is a sense in which episodic memory is centered around a concept, the concept of "self". If I were to become famous, many of the things in my personal diary might be entered in an encyclopaedia under an entry for me. Personal experiences are all the experiences of an individual and since there has to be a concept of self in memory one might argue that episodic memory is merely the structure of knowledge associated with the concept of

the individual. Such a conclusion I find quite unobjectionable. It suggests that all input and output is mediated through a concept of the self. Input, initially represented there, may be transferred to other concepts under certain conditions thereby freeing it of its autobiographical reference (that doesn't necessarily mean it gets automatically destroyed in its original form). Output can be represented as "I said that..." or by some other performative verb. It is interesting that a characteristic feature of generative semantics is that all sentences are dominated by a verb representing the speech act -- such a view is of course quite a natural consequence of the prominence that I am suggesting be attached to the concept of self.

There are a couple of other reasons which incline me to support this distinction. First there is a mass of experimental research on encoding specificity which strongly suggests the need for some such distinction (see, for example, Watkins & Tulving, 1975). Secondly there is the rather mundane observation that we frequently come across people who have "bad" memories. The problem with such people is not that they have inferior, or diluted knowledge-bases -- they are not lacking in knowledge. Their problem is the maintenance of a record of personal input/output experiences. They tend to absorb what they consider important, they are perfectly able to learn, but they can't remember where they put their car, or to whom they lent a book, or from whom they heard something. If all semantic memory were episodic, they would have to be stupid as well as forgetful!

There remains (as always) a number of unanswered questions if one takes this approach. The most important one, of course, is how and when does information get into semantic memory, and how and when does it get lost from episodic memory. Answers to these questions are beyond the scope of this paper, but some speculations at least about the first might start things off. In order to infer S from S' inference based on some or all of the following considerations will have to be made: Is it relevant? This suggests that not everything we experience is worth recording in semantic memory (what I ate for dinner on some arbitrary day, for example). How reliable is the source? If a three year old child tells me that all men are daddies I may find it relevant to what I need or want to know about the child, but I am not likely to trust him as an information source about universal qualities of men. To what extent does it conflict with what I already believe? The child's claim faces another problem on this score. One must assume that the knowledge base can tolerate some degree of inconsistency, but as Quine (1952) suggests, some things may be much more difficult to give up or revise than others. Relevance or criteriality indices will handle this, together with degrees of truth. Is there independent corroboration? If some putative fact or procedure seems, on the basis of what is already known, very unlikely to be true, independent evidence for it may boost its degree of truth for me.

There are doubtless other considerations (many can probably be gleaned from the belief and attitude literature) but the general point is that an assessment of a degree of belief has to be made. If a threshold is reached the thing has some positive degree of truth and can be stripped of the source and circumstance information (perhaps with some modification), and once this is achieved all is in order to free it from its bound form and to treat it as a bone fide piece of knowledge to be incorporated into the existing structure in semantic memory. Clearly an analysis of speech acts will be required since the rules for freeing the putative knowledge will be different for performatives such as "convince", "promise", "prove," etc.

The problems that we face are difficult and complex. I am confident that they can be solved, but I am equally convinced that to solve them in a satisfactory way is going to require that our models develop from ideas which are both epistemologically and psychologically sound. I have tried in this paper to put forward some arguments and to develop some ideas towards this end.

#### REFERENCES

- Anderson, J.R., & Bower, G.H. Human associative memory. Washington, D.C.: Winston & Sons, 1973.
- Anderson, R.C., & Ortony, A., On putting apples into bottles: A problem of polysemy. Cognitive Psychology, 1975, in press.
- Collins, A.M., & Loftus, E.F., A spreading activation theory of spreading activation. Psychological Review, in press
- Ortony, A., Explorations in memory. Urbana, Ill.: University of Illinois 1975 (mimeo)
- Ortony, A., & Anderson, R.C., Definite descriptions and semantic memory. Urbana, Ill.: University of Illinois 1974 (mimeo)
- Quine, W.V.O., Methods of logic London: Routledge & Kegan Paul, 1952, pp. xi-xvii.
- Rumelhart, D., & Norman, D., Active semantic networks as a model of human memory. in Proceedings of the Third International Conference of Artificial Intelligence, Stanford, 1973, pp. 450-457.
- Schank, R.C., Is there a semantic memory. Castagnola, Switzerland: Istituto per gli Studi Semantici e Cognitivi, 1974 (mimeo)
- Schank, R.C., The structure of episodes in memory. In D.G. Bobrow and A.M. Collins (Eds.), Representation and understanding. New York: Academic Press, 1975.

Tulving, E., Episodic and semantic memory.  
In E. Tulving and W. Donaldson (Eds.),  
Organization of memory. New York:  
Academic Press, 1972, pp. 381-403.

Watkins, M.J., & Tulving, E., Episodic  
memory: When recognition fails. Journal  
of Experimental Psychology: General,  
1975, 104, 5-29.

Winograd, T., Understanding natural  
language, Edinburgh: Edinburgh University  
Press, 1972.

Wittgenstein, L., Philosophical  
investigations, translation by G.E.  
Anscombe, Oxford: Blackwell 1963.

Woods, W.A., What's in a link: Foundations  
for semantic networks. Cambridge, Ma:  
Bolt Beranek and Newman, 1975 (mimeo)