

A MODEL FOR KNOWLEDGE AND ITS APPLICATION TO
DISCOURSE ANALYSIS

BRIAN PHILLIPS
University of Illinois at Chicago Circle
60680

TABLE OF CONTENTS

Introduction	4
A model for knowledge	4
Nodes	6
Arcs	7
Paradigmatic relations	8
Variety	8
Instance	10
Typical	11
Manifestation	11
Syntagmatic relations	16
Discursive relations	18
The metalingual relation	20
Status relations	23
Negation	25
Inheritance	26
Episodic and systemic memory	28
Quantification	29
Processes in the network	31
Path-tracing	33
Pattern-matching	37
Discourse analysis	44
The structure of coherent discourse	44
Anaphora	44
Spatial, temporal, and causal cohesion	45
Thematicity	46
The role of the encyclopedia	52
Anaphora	53
Spatial, temporal, and causal cohesion	54
Thematicity	56
Implementation	57
Processes	58
Paradigmatic path tracing	58
Causal connectivity condition	58
Discovering general and specific propositions	58
Metalingual decomposition	60
Metalingual abstraction	60
Inference of omitted discursive relations	62
The system	62

Analysis of some stories	63
Common patterns	65
A boat capsizing	67
A fall into water	70
Embedded themes	72
Discussion	76
Acknowledgements	79
References	79

INTRODUCTION

An important contribution of natural language processing has been to direct attention to the structure of language at the discourse level, which has led to a greater awareness of the role of "meaning" in language for "A text is best regarded as a SEMANTIC unit; a unit not of form but of meaning" (Halliday & Hasan, 1976, p. 2). This being so, discourse analysis will deepen our understanding of meaning and vice-versa.

In this paper I present a model of meaning strongly influenced by Hays (1969a, 1969b, 1970, 1973) and show how it is able to capture the organization of discourse. In particular I seek to define the organization of coherent discourse and to show how knowledge is used to infer a coherent structure when, as usually is the case, the surface form is elliptic. The hypotheses are used to build an automatic system to test the coherence of discourse.

A MODEL FOR KNOWLEDGE

The philosophic stance is taken that our knowledge of a concept is the meaning of that concept: "someone who knows what tiger means . . . is required to know that stereotypical tigers are striped" (Putnam, 1975, p. 249).

Many models of knowledge have been developed for use in computational environments. Some are for restricted domains (Black, 1968;

Bobrow, 1968; Colby, 1973; Raphael, 1968; Winograd, 1971; etc.) The present model is more in the tradition of Klein, Oakley, Suurballe, and Ziesemer (1972), Quillian (1969), Rumelhart, Lindsay, and Norman (1972), Schank (1975a), Shapiro (1971), Simmons (1970), and Wilks (1972), where no particular context is prescribed. It will be apparent that at many points the present model draws upon these earlier systems. Some of the differences between systems are probably differences in notation. However no system is at a stage of constancy or completeness that makes it worthwhile to devote much effort to establishing the equivalences. Although it would be possible to present only the parts that I believe to be novel, giving the whole system in a common notation will ease the task of the reader.

The model, hereafter called the encyclopedia, endeavors to be consistent with available psychological and linguistic views of the structure of language and thought, for any automated language system must closely imitate the workings of human cognition to be successful (Collins & Quillian, 1972).

The encyclopedia encodes common knowledge of the world which may differ from scientifically accurate descriptions. Putnam (1975) calls such knowledge "stereotypical":

The fact that a feature . . . is included in the stereotype associated with a word X does not mean that it is an analytic truth that all Xs have that feature, nor that most Xs have that feature, nor that all normal Xs have that feature, nor that some Xs have that feature. . . . Discovering that our stereotype is based on nonnormal or unrepresentative members of a natural kind is not discovering a logical contradiction. . . . [but] The fact is that

we could hardly communicate if most of our stereotypes weren't pretty accurate as far as they go. (pp. 250-251)

The encyclopedia is schematized and implemented as a directed graph; in current parlance it is a network model. Nodes characterize concepts and arcs relations between concepts. The most general statement to make about the model is that relations and concept types are the necessary system primitives; some concepts may be primitive, but the model does not depend on the existence of primitive concepts.

Discussion will cover the nodes and relations of the model. Attention will also be given to network processes.

No psychological validity is claimed for the content of any of the structures shown; the claim extends only to the relational structure. Questions of content must be answered empirically.

Nodes

There are four types of nodes: event, entity, attribute, and modality. The first three correspond to simple verb, simple noun, and simple modifier, respectively. The fourth type of node is novel. Its role in the system will become clear after a description of arcs. For the meantime it will have to suffice to say that it is used in the spatio-temporal causal, belief, and hierarchic organization of knowledge. Its ancestor in linguistic theory is "modal" in the modal/proposition dichotomy of Fillmore (1969). Schubert (1976) has predicate nodes that are similar in motivation, but different in use from modality nodes.

Nodes of the encyclopedia are not labeled (Collins & Quillian, 1972). An arc, termed name, points from a node into a dictionary of print names. For clarity nodes in diagrams will be annotated, but this should not be taken as representing the implementation, which is as shown in Figure 1.




Figure 1 Labeling nodes

In all the following figures • is an event, entity or attribute node. Annotations on these nodes are enclosed in //, <>, and [], respectively. Modality nodes appear as □ and are never annotated.

Arcs

Five types of arcs are used in the network: paradigmatic arcs are taxonomic, syntagmatic arcs form propositions, discursive arcs link

propositions, the metalingual arc is used to associate a concept with a story in network form that defines the concept, and status arcs characterize beliefs and desires.

Paradigmatic Relations

Variety. A readily observable aspect of human behavior is the existence of folk taxonomies. These have been studied in detail by ethnosemanticists in order to discover their cognitive significance and structure:

Man is by nature a classifying animal. His continued existence depends on his ability to recognize similarities and differences between objects and events in his physical universe and to make known these similarities and differences linguistically. Indeed, the very development of the human mind seems to have been closely related to the perception of discontinuities in nature. In view of this, the study of folk taxonomic systems, which have received a great deal of interest in recent years, has a high significance in interpreting the logical processes going on in our minds, as well as in understanding the application and utility of the taxonomic systems themselves. (Raven, Berlin, & Breedlove, 1971, p. 1210)

For example, mammal, bird, and reptile might be classified as kinds of vertebrates. In the network, the relation is termed variety (abbreviated to VAR in the figures). Figure 2 diagrams this knowledge.




Figure 2 Paradigmatic organization

Varietal nodes are seen as representing concepts at a categorical level, hence variety is the category-subcategory relation. Berlin, Breedlove, and Raven (1968) show the existence of covert categories in folk taxonomies, i.e., nodes having scientific, but not folk, names, say "vertebrate" in Figure 2. These categories are revealed by memory, classification, and other experiments. This is counter to the view of

Conklin (1962) for whom concepts must have monolexemic labels. Covert categories enable Raven et al. (1971) to show a degree of uniformity in taxonomies: about five hierarchic levels with seldom more than five hundred items under one node. Berlin et al. (1968) claim that items in a folk taxonomy form non-intersecting categories, i.e., the structure is strictly tree-like. This view is not held here, for a typewriter can be classified both as a machine and as a writing instrument. Consequently, varietal structures are not restricted to being tree-like. Loops, however, do not seem possible. Nor is it necessary that a node have a name.

Instance. Logic since Aristotle has distinguished between category (or type) and a specific member of a category (token). This membership relation is termed instance (IST). For example, "William Proxmire" is an instance of "person", Figure 2. Most instances are not named, taking their name from their varietal parent, but a major exception is people, e.g., "Peter", "Aunt Sally", Figure 1.

Any path through the paradigmatic organization of knowledge which follows only arcs having the same directionality (termed a paradigmatic path) contains at most one instance arc. Traversing this arc represents a cognitive transition from thinking about categorical concepts to thinking about particular concepts, e.g., from thinking about man to thinking about Abraham Lincoln, or from blueness to the blue of your car.

Rumelhart et al. (1972) use an ISA relation that covers both variety and instance, e.g., ISA(Luigi s,tavern) and ISA(tavern,establishment). The present feeling is that a distinction does exist; hence the two relations of the encyclopedia.

Typical. A third condition of knowledge needs to be represented. Concepts have both universal and occasional properties. For example, "birds eat worms" is an occasional, not universal, fact about birds as some never do, but even those that do are sometimes found eating fruit, fish, or even not eating, without the proposition being necessarily false. However, "birds have wings" is expected to be true at all times for all birds; it is a pathological situation if a counter-example is found. To represent the arguments of occasional predications, the typical (TYP) arc is used. Thus the "bird" in "birds eat worms" is as in Figure 2. It is also possible to use the typical relation to attach occasional properties to members of categories, i.e., to instances. In Figure 2 is shown the representation of "William Proxmire makes foreign policy statements" where this is a statement of an occasional habit rather than a record a specific act. No position is taken on how noteworthy knowledge is recognized as such in the development of the encyclopedia.

Manifestation. The final paradigmatic relation is manifestation (MAN). This corresponds to the phenomenon of object constancy: An object may undergo change in space and time, but it is still perceived

as the same object. For example, William Proxmire before and after his hair transplant is still William Proxmire. Also an object may participate in many different actions but still preserve its identity, e.g., Albert Einstein playing a violin and Albert Einstein writing on a blackboard remains Albert Einstein. In the system each different situation involves a distinct node. To a node defined by an instance are linked, by manifestation arcs, nodes that correspond to an object in its different guises. Manifestations of "William Proxmire" are shown in Figure 2. Manifestations do not usually have names different from that of their parent instance; a rare exception to this is the Evening Star and the Morning Star which are both Venus at different times of the day.

Manifestations of varietal and typical concepts are also possible. The latter are used for properties that are true of the concept but only at some point or period of time, for example, "vertebrates are born", Figure 2. For typical arcs this notion is redundant as typical embodies spatial and temporal indeterminacy. However manifestation does have a use with the typical arc in representing coreference. Suppose it can happen that a person can trip causing him to be hurt. The "person" in the encoded event is a typicalised "person", but it must be the same person that trips that is hurt. Figure 2 shows the use of manifestation relations to indicate this identity. More will be said later about the formal representation of the causal relation

indicated in Figure 2. If only typical arcs were used, the interpretation would be that anyone tripping could cause literally anyone to be hurt. Multiple manifestations can also be used with variety if coreference needs to be marked.

[This next paragraph is almost certain not to make sense until the reader has completed as far as, and including, the section "Inheritance", and so he may choose to leave it and return later.

Other systems, Quillian (1969), Rumelhart et al. (1972), and Schank (1975a), do not use manifestation but capture object constancy by having one and the same node for a participant in all of its propositions. This is a viable alternative. Nevertheless, information on the relative standing of the appearances of the participant has to be representable. If a single node were used in the encyclopedia, the differentiation could be made on the modality nodes of the propositions. This route was not taken as it is more convenient, for example, to let the nature of the inheritance be determined completely in paradigmatic organization, rather than in a mixture of paradigmatic and discursive structures. For even without manifestation, the varietal structure will require the process of inheritance.

Of the four arcs, variety, typical, and manifestation can be iterated; instance cannot. Figure 2 and 3 illustrate iterative arrangements of variety and manifestation. That typical also has this property is seen from considering that "While dreaming, some people

talk or sleep-walk". None of these propositions are universally true, but only of arbitrary people. Figure 2 contains this situation. The above examples present only paradigms of entities, but events and attributes also exhibit this kind of organization.

If paradigmatic structure is a loopless directed graph then there will be origin nodes, that is, nodes without entering arcs. Can anything be said about the number or kinds of concepts associated with origin nodes? It is speculated that entities can be divided into domains of being each of which has its own paradigm. Possible domains are thing, soul, role, time, etc. Thus to represent Ford as President of the USA the structure in Figure 3 would be used. Figure 3 also shows how the totality of John Brown (JB) and his fragments, as in "John Brown's body lies mouldering in the grave but his soul goes marching on" can be represented.




Figure 3 Domains of being

To date scholars have only studied entity paradigms in detail. Little investigation of attribute or event paradigms has taken place. It is hard to intuitively discern the hierarchical ordering of these concepts, i.e., to know which concepts imply others. Red, yellow, etc., are obviously varieties of color, but does having mass imply having color?--but many gases and glass have mass but are colorless. Or does having color imply having mass?--but red light, blue jokes, etc. Or are they quite independent attributes that happen to have a

large intersection in their domain of applicability? These are all open questions in the taxonomy of attributes. The event paradigm is also open to much speculation.

Syntagmatic Relations

Syntagmatic relations connect nodes from different paradigms (with one exception). Relations of participation, similar to Fillmore's (1969) case relations, connect entities and events. A relation of application (APL) links attributes to events or to entities. A relation of part-whole (P-W) connects a unity to its components. A syntagmatically related structure is termed a proposition.

Four relations of participation are distinguished: agent (AGT), instrumental (INS), objective (OBJ), and experiencer (EXP). The role characterized by each is derived from dichotomies animate/inanimate and causal/non-causal, as given in Table 1 (Fillmore, 1969).

	Animate	Inanimate
Causal	AGENT	INSTRUMENTAL
Non-causal	EXPERIENCER	OBJECTIVE

Table 1 Relations of Participation

Thus "Angry Bill ferociously hit Fred with the handle of an axe" is diagrammed as in Figure 4.




Figure 4 Syntagmatic organization

The set of case relations does not include locative and temporal relations. Sentence adverbials (Chomsky, 1965) are not part of syntagmatic structure, but of the contextual structure, which is here represented on modality nodes. Bound adverbials are part of syntagmatic structure, e.g., "ferociously" above, and are related to the event node by a relation of application.

A part-whole relation is used in Figure 4 to show the relation of "handle" to "axe". This relation differs from other syntagmatic relations in that it connects nodes of the same type, e.g., two entities. A case can be made for this relation to be considered a paradigmatic relation; for the present it has been put in with the syntagmatic mainly because it is not used by the process of inheritance, of which more later.

Discursive Relations

Propositions do not occur in isolation. They are tied together in cognition in a number of ways. The spatial, temporal, and causal connections are characterized by discursive arcs. Intuitively these are relations between whole propositions and it is to fail to capture this feeling if, say, a cause relation directly links two event nodes. Modality nodes are used to represent situations in which the whole proposition is involved. Though schematically linked to an event node, conceptually the modality belongs to the whole proposition. Discursive arcs relate the modalities of propositions. Thus "Mary slapped John because he chased her" is represented as in Figure 5.




Figure 5 Discursive organization

The one causal relation, cause, admits of no finer distinction. Others (Schank, 1975a; Halliday & Hasan, 1975) distinguish three kinds of causation: reason, result, and purpose. The single cause relation of the encyclopedia models the first two directly. Purpose (or enabling) causation is seen as separable into cause together with a desire for the consequent. For example, a cup may fall causing it to break. The fall could be accidental or it could be deliberate with the purpose of breaking the cup. The same causal relation exists between the actions in both cases, but the analysis of the purposive situation will involve "desire".

Time arcs do permit subdivision. A proposition may be simultaneous (SML) with another proposition: "Fred washed the car while John chased Mary", Figure 5. A sequential ordering of propositions is also found, characterized by a sequence (SEQ) relation. The suggestions made here for the organization of space are only a working set for which little justification can be offered: location (LOC)--a neutral statement of position, contact--in physical contact, and near, far, above, below, left, right, etc., which are self-explanatory. Figure 5 represents the location of "Fred washed the car" as being "garage". Since this work was completed Sondheimer (1977) has proposed an analysis of space and time.

The Metalingual Relation

Speech acts do not make use only of forms having physical reference, e.g., table, John, blue. A most important aspect of language behavior is abstraction. Human social, scientific, and intellectual development is dependent on the ability to create and control abstract concepts. A quick appraisal of this paragraph reveals many such concepts: language, behavior, social, etc. A system that seriously hopes to approach human capabilities must have a corresponding ability.

One part of modeling abstraction is representation; but what is to be represented? Abstraction involves knowing a situation in which the abstract term applies and replacing the situational description by the abstract term. An example is "tragedy". The scene to which it is applicable is, say, "Someone does a good act that results in his death". This definiens is encoded in Figure 6. "Tragedy" names a single node.




Figure 6 Metalingual organization

The general propositions of the definiens are conjoined using a modality node linked to the modalities of the propositions by part-whole relations. In general there may be any number of levels of modalities related by part-whole. To complete the association of the definiendum with its definiens, a metalingual (MTL) arc links the former to the appropriate modality in the latter, Figure 6. If any situation matches the definiens, then the abstract term is appropriate. The process of matching will be discussed later.

The definiendum can also be any concept, the choice is idiosyncratic; there is no reason why this device cannot be used with apparently non-abstract concepts, for example, a dog could be "man's best friend" for some, in contrast with a non-abstract definition of "canine

animal". Non-abstract definitions have the form "genus-specificata". In the encyclopedia, the representation is made up from a node related by variety to the genus (animal), to which are attached the properties in the specificata (canine).

Rumelhart and Ortony (1976) use a relation similar to metalingual, ISWHEN, but do not show how participants are equated in the definiendum and definiens, nor the processes that use such definitions.

The metalingual arc is used in another context. Some propositions contain embedded propositions. For uniformity it is desirable to restrict participation in propositions to entity nodes. Thus the matrix proposition has an unnamed participant in objective or instrumental role and this node is defined by a metalingual arc to the modality of the contained proposition. For example

(1) Peter believes Fred chased the cat.

is represented as in Figure 7.




Figure 7 Embedding propositions

Status Relations

Knowledge in an encyclopedia is a model of the beliefs of one person. Nevertheless the knowledge is not all of the same status. In addition to containing the person's beliefs, it includes representation of beliefs about his own desires and of his beliefs about the beliefs and desires of others. His personal beliefs and desires interpret, control, and direct his personal activities. The knowledge about others is the basis for interacting and communicating with them. For example, a conversation with a child about the structure of matter is quite different from one with a nuclear physicist because of different conceptions about their levels of knowledge and hence what can be taken for granted. One has knowledge about individuals, e.g., your brother, Nelson Rockefeller, etc., and about groups, e.g., politicians, sports writers, Russians, etc.

A distinction can be made between subconscious and conscious knowledge. The former is, for example, the knowledge of language underlying its use or (2).

(2) The Sun circles the Earth.

Conscious knowledge is learnt or communicated knowledge, e.g., what one has been taught about the solar system. There is no reason for the two kinds of knowledge to be in accord regarding the same entities. One has learned, for example, that the Earth circles the Sun.

Subconscious beliefs of self are unmarked in the encyclopedia. Subconscious beliefs of another are indicated by a believe arc between a node representing the believer and a modality node covering the network representation of the content of the beliefs. The subconscious belief of (2) by "people" is given in Figure 8.




Figure 8 Knowledge status

Conscious beliefs are represented as propositions embedded within an event "believe". An example is given in (1) with its representation in Figure 7.

It is not only propositions that have belief status, but also simple concepts, e.g., ghosts. To accomodate this information, the placing of modality nodes is generalized. Previously only propositions were associated with modalities; now any node can have its own modality. On this modality information about a concept's existence and belief status can be represented, as in Figure 8 for "Fred Smith".

It is unlikely that each node or proposition is immediately linked to its believer. Using part-whole relations and modality nodes, domains of belief, which may intersect, can be created as in Figure 8 for "Hugo".

Hendrix (1975) partitions semantic networks to delimit domains of belief; here the same effect is gained through the use of modality structures.

The desires of people are situations that they would like to exist. The content of these goals can be represented by a modality covering (complexes of) propositions or single concepts, e.g., peace. If the goals are subconscious, a desire relation links the desirer to the modality. For conscious states, the modality is part of a metalingually defined objective of an event "desire". In modeling behavior, these goals provide the situations that other behavioral actions are intended to contribute towards achieving.

Negation

Negation is a property that is marked on a modality. The most common site for negative marking is a propositional modality. Thus Figure 9 contains the proposition "I do not like tomatoes".




Figure 9 Negation

When some other constituent of a sentence is negated, say using strong stress, this is marked on a corresponding modality, so "John did not hit Mary" is encoded as in Figure 9.

It is not anticipated that negation is a common feature in knowledge, for "A person sometimes learns a negative fact when it contradicts something that might be inferred by mistake or that is true for a similar concept. But, most negative facts are never learned" (Collins & Quillian, 1972, p. 319).

Inheritance

A node will inherit properties from nodes higher in its paradigmatic path. Quillian (1969) used superset relations for the same purpose. In Figure 2, B inherits the properties of A, C those of B, D those of E, and E those of D. Inheritance is transitive, thus E inherits properties from A, B, C, and D. This permits parsimonious

representation of properties: A property need only appear at the ancestor of concepts having the property. Inheritance is inhibited only if the inheritable property is contradicted on a lower node. For example, although the property "fly" may be associated with "bird", it is prevented from being inherited by "penguin" by having explicitly "penguin not fly".

The generality of inheritance depends on the form of representation of the property at the ancestor node. Properties that are universally true at all times, e.g., birds have wings, are attached directly to a varietal node and are obligatorily true of all descendents. If at any time a bird without wings were reported, it would be cause for further explanation. Some other properties are always true but only at intermittent times, e.g., people eat, whose representation involves the manifestation relation. It is not odd that a person can be seen not eating, but if you watched long enough, it would be fully expected to observe this behavior sometime. Finally there are occasional properties that make use of the typical arc in their representation. These properties are not universal, being merely noteworthy recollections about a concept, e.g., The French are rude. It would well be possible to have a complete history of an example of the concept and not witness the property without being disturbed by its absence.

Episodic and Systemic Memory

Tulving (1972) distinguishes episodic from semantic memory. The former "receives and stores information about temporally dated episodes or events, and temporal-spatial relations between events" (p. 385). The latter is

knowledge a person possesses about words and other verbal symbols, their meaning and referents, about relations among them, and about rules, formulas, and algorithms for the manipulation of these symbols, concepts, and relations. Semantic memory does not register perceptible properties of inputs, but rather cognitive referents of input signals. (p. 386)

Abelson (1975) distinguishes episodic from propositional memory, and Woods (1975) contrasts intensions with extensions along similar lines. The term I prefer, following Hays (1978), is systemic rather than semantic, propositional, or intensional.

The localization in space and time of knowledge is represented in the encyclopedia by spatial and temporal organization of propositions using the appropriate discursive relations. A proper subpart of episodic memory is contained in paradigmatic organization. Manifestations of instances (remember there are also manifestations of varietal and typical nodes, so it must be thus stated) represent spatio-temporally localized information about members of categories. Consequently knowledge represented on manifestations of instances, or their manifestations, is in episodic memory. This is only part of episodic memory as categorical knowledge can also be present. For example, in "Jung changed our view of dreams", the reference is to the categorical

notion of dreams, not to any specific ones. Nor is it sufficient for a proposition to have a non-categorical participant to be in episodic memory for "Prior to the Revolution, Russian peasants were feudal serfs" contains categorical participants, yet is episodic. The total extent of episodic memory is ultimately decided through spatial and temporal relation of discursive organization, not by paradigmatic structure.

Quantification

Paradigmatic arcs have the capability of capturing the essence of quantification, including scope. To illustrate the facility, consider (3) and (4) which are equivalent to the formulae (5) and (6) Figure 10 encodes (3) and (4).

- (3) There is a book that is read by every scholar.
- (4) Every chorister knows a song.
- (5) $\exists x \forall y [(book(x) \ \& \ scholar(y)) \supset read(y,x)]$
- (6) $\forall y \exists x [(song(x) \ \& \ chorister(y)) \supset know(y,x)]$




Figure 10 Quantification

If for a given chorister in (4) it is necessary to determine the song he knows, i.e., to evaluate the Skolem function, the information is present as a predication of that individual and should be retrieved using his name, say "George" and "Drink to me only with thine eyes" -in Figure 10.

It is also possible to give distinct representation to unquantified statements, such as (7), as in Figure 10.

(7) A person likes candy.

Paradigmatic arcs are here achieving representational power equivalent to the partitioning of networks by Hendrix (1975).

The above is a systemic rendition of "all". The quantification can also be characterized episodically by every manifestation of a concept having the property. Interpreting "all" (Woods, 1975), could call upon either systemic or episodic facts. A question containing a universal quantifier may be answered by either examining a varietal node (Are all mail-boxes blue?), or by examining every manifestation (Do all mail-boxes stand at street corners?).

It should be noted that "all" requires that the predication be true only at some time, e.g., All people die; it does not require continuity in time, e.g., All birds have wings. Thus universal quantification is also true if the predication is found for a manifestation of the varietal node, or is found for every instance of the concept.

Processes in the Network

The model for knowledge described above is only part of a system to model cognitive behavior. Thought is simulated by processing knowledge. Different aspects of behavior correspond to different processes, but with one and the same encyclopedia common to all. A system for discourse analysis requires processes that use the encyclopedia to find patterns of organization in a discourse. It would be possible to describe solely the requirements of discourse analysis, but greater overall insight is gained through a preliminary general examination and classification of cognitive processes. Once this is accomplished,

discourse analysis is seen not to be a unique process but as composed of more basic general ones. Simulation of many aspects of cognitive behavior can be performed by complexes of these general processes: discourse analysis is just one such complex.

Processes can be classified in various ways: functionally, by complexity, or by the class of relation involved.

The function of some processes is external; they deal with input and output. Some internal processes find relations between new information and knowledge already in the encyclopedia, others investigate the validity of new knowledge, etc.

Processes are of two type of complexity, either path-tracing or pattern-matching. The dichotomy is justified by showing that there are tasks that can only be done by pattern-matching. This topic is considered in detail later.

Of the infinite number of possible ordered sets of arcs, only some define significant paths in the network. An example of a relevant set of arcs is the arcs of a paradigmatic path; this defines possible inheritances. Other significant sets are causal chains, which are represented by a string of cause arcs between modalities. This suggests that processes that use the same kind of relations or identical relations are significant.

A functional classification of processes does not give a deeper understanding of cognitive processes. However, classification by

complexity and by kind of arc is revealing. Path-tracing and pattern-matching differ in power. For the former, subpaths can be defined by the kind of arc found in the subpath. Henceforth processes in the network will be described according as they are path-tracers or pattern-matchers.

Path-tracing

Path-tracing processes try to establish paths between nodes along arcs of the network. Quillian (1969) established this methodology for semantic nets. A particularly common type of path is the paradigmatic path. In Figure 3 there is a paradigmatic path between "Ford (as President)" and "thing", but not between "rock" and "soul". The definition of a paradigmatic path is valid for entities, events, and attributes.

Any paradigmatic path in the network will conform to the structure shown in Figure 11 where * indicates any number of occurrences including zero of the marked relation.




Figure 11 Paradigmatic paths

The structure follows directly from the iterativity of variety, manifestation, and typical arcs and their possible relative orientations. Strings of arc labels representing paths through the tree are obviously regular expressions, i.e., the strings are sentences of a type 3 language. Paradigmatic path-tracing can thus be characterized by a finite state automaton (Hopcroft & Ullman, 1969).

Any process that can be characterized by a finite state automaton is formally termed a path-tracing process in the system.

One such process is testing the applicability of an attribute to an entity, e.g., whether "fresh fish" or "round smoke" is acceptable when the relationship is not explicitly in the encyclopedia. Assuming the named entry points to the encyclopedia are at variety or instance nodes, an entity F_1 (e.g., horse) can inherit properties from an entity E_2 (e.g., animal) if there is a path between F_1 and E_2 of the form $(\overline{IST}) \overline{VAR}^*$, where $\overline{X}$ indicates a relation that is the converse of X and $()$ indicate an optional arc. Properties may be attached to E_2 either directly or with typical and/or manifestation arcs, i.e., the path from E_2 to the node F_3 in the representation of the property has the form $TYP^* MAN^*$. Thus the path from E_1 to E_3 has the form $(\overline{IST}) \overline{VAR}^* TYP^* MAN^*$. Analogously, an attribute A_1 can apply to an entity if there is a similar path to an attribute that is encoded as applying to the entity. Thus if there is a path

$$(8) \quad \langle E_1 \rangle (\overline{IST}) \overline{VAR}^* TYP^* MAN^* \overline{APL} \overline{MAN}^* \overline{TYP}^* VAR^* (IST) [A_1]$$

then A_1 can feasibly apply to E_1 . That is to say, the explicit encoding of "emotional animal" would make it reasonable to infer "sad horse". The path (8) is composed of paradigmatic paths linked by a single application arc. Each segment is a regular expression. As type 3 languages are closed under concatenation (Hopcroft & Ullman, 1969, theorem 3.8), it follows that (8) is also a regular expression and that attribute applicability testing is a path-tracing process.

Propositions in a discourse should be consistent with encyclopedic knowledge. Consistency is established by finding a proposition in the encyclopedia that is a generalization of the discourse proposition, e.g., given the discourse proposition

(9) Marv gobbled the caviar.

and finding the generalization

(10) People eat food.

A novel statement, e.g., "Harry munched the spider", which is not consistent with (10) (assuming "spider" is not a variety of "food"), would evoke a demand for further explanation, or similar. Consistency judgment can be formulated as a complex of path-tracing processes. In the network form of (9), "Marv" is the agent and "caviar" is the objective of "gobble". Figure 12 encodes (10).




Figure 12 Consistency judgment

The words in the discourse proposition provide entry points into the network of Figure 12 through the dictionary and converse name relations. From "gobble", node 1, paths along paradigmatic arcs are traversed to locate nodes from which "gobble" could inherit properties, e.g., node 2. Next from the entries for "Marv" (A), and "caviar", (B), analogous paths are followed, reaching C and D, respectively (among other nodes). From C and D arcs corresponding to the participatory relations of "Marv" and of "caviar" to "gobble", i.e., agent and objective, respectively, are followed. If all paths intersect at a single node, e.g., node 2, then the proposition containing the intersection is the general proposition sought. Each path from an entry point to an intersection can be characterized by a regular expression. There are

only four case relations, which sets a finite upper bound to the number of paths to be followed. Hence this process is also a path-tracing process.

Locating existing knowledge, propositions that are already explicitly in the encyclopedia, is effectively identical to the consistency testing process above, but with downward paradigmatic paths being followed instead of upward ones. Thus given "Oswald assassinated Kennedy" and the network of Figure 13,




Figure 13 Finding known propositions

paths can be traced from node 1 to node 2, from node A to node B to node 2, and from node C to node D to node 2. The common intersection is in the known proposition.

Pattern-Matching

Pattern-matching is used in processes where two configurations of nodes and arcs must match. One such process is with metalingually defined terms. If a discourse configuration matches a metalingual

definition, then the part of the discourse so matched may be replaced by the term. Figure 14 contains representations of (a) "Fred ate some cake that made him sick" and (b) the definition of "poison": "Someone ingests something that makes him ill".




Figure 14 Pattern-matching

If the latter matches the former, then "poison" describes the discourse situation. Earlier, a path-tracing process was used to establish consistency between a general and a specific proposition. The same process can be used to pair propositions of the discourse and the definition. However, there is an aspect of complexes of propositions that prevents path-tracing from being a complete solution. If the complex contains coreferential items, as "poison" does, this coreferentiality

must be examined; if it were not errors could result. For example, consider a discourse containing "John's eating the worms made Fred sick". Each proposition matches part of the definition of "poison", but it should not be taken as an act of poisoning. The coreferentiality condition prevents a match. As, in general, there can be any number of coreferential participants in a complex of propositions, it is not possible to define a regular expression to characterize the coreferentiality test. This can be shown by considering a definition of an abstract term that contains n coreferential concepts. There is in general no bound on n as the definition can contain any number of propositions. If a complex of discourse propositions is to match the definition then there must first be a unique corresponding proposition in the definition for each discourse proposition. This can be done using the path-tracing process described above. But over and beyond this, the coreference condition must be satisfied. For each manifestation of the coreferential concept in the definition there must be a corresponding manifestation of one and always the same concept in the discourse. Also the syntagmatic role of corresponding manifestations in their respective propositions must be the same. The acceptance condition involves pair-wise counting. This is equivalent to accepting strings of the form $a^n b^n$, which are not sentences in a type 3 language (Hopcroft & Ullman, 1969). This demonstrates that processes that compare complexes of propositions containing coreferential items are not, in general, path-tracing processes.

A process characterized by a device more powerful than a finite automaton is formally designated a pattern-matching process.

Paraphrasing discourse using metalingually defined terms is another pattern-matching process. Metalingual definitions can be recursively embedded. For example, "buy" may be defined in terms of "give", which in turn may be defined in terms of "have". Recursion is not a property of regular languages, hence this process is not a path-tracing process.

Matching discourse configurations against definitions, called abstraction, is an extension of the process that substantiates discourse propositions by seeking generalized propositions in the encyclopedia, discussed earlier. The components of a definition are generalized propositions and hence the substantiation process will find them if they correspond to part of the discourse. Schematically, two discourse propositions DP_1 and DP_2 , may match generalized propositions GP_1 and GP_2 and GP_3 , GP_4 , and GP_5 , respectively, as shown in Figure 15.




Figure 15 Abstraction

This is the normal output when judging consistency. Propositions of a definition are under a conjoining modality, to which the metalingual arc points. If it is found that some of the general propositions are part of definitions, i.e., GP_2 and GP_3 in MLD, then these definitions are examined to see if all the conditions for their use are satisfied, i.e., coreference and contextual (e.g., cause arcs) conditions. For example, in "poison", Figure 14, the coreference of the agent of "eat" and the applicand of "ill". If a definition is satisfied, then the part of the discourse matching the definiens can be paraphrased.

The definitional nets so far presented are not adequate for paraphrasing, but must be augmented to include the roles of entities of the definiens with respect to the definiendum. This is done with manifestation arcs. A network definition of "buy" (in "A buys thing from B for money") is given in Figure 16. The verbalization is "A gives money to B and B gives thing to A".

"Jane gave a bicycle to John and John gave money to Jane" to be generated. "Money" was unexpressed in the original, but is present in the definition, and appears in the paraphrase. The process fills empty slots by the appropriate concept from the definition, in this case "money". The agent, experiencer, and objective slots are filled in the source statement and are transferred to the paraphrase.

Another abstract term can point to the same definitional network, say "sell" in the case of Figure 16. The net then has all the information for paraphrases between the two abstract terms as well as for decomposition and abstraction.

There is no productive relationship between the roles of the same participant at different levels of abstraction. Case relations represent only the causal/animate perception of participation in an event. More detailed descriptions of the roles of participants can only be given in context. For example, "money" is perceived as instrumental in "buy", but at the next level of decomposition, it is in an objective role in "give".

The outputs of both abstraction and decomposition are structurally indistinguishable from any other proposition in the encyclopedia and therefore can again be subject to either of the processes. As pattern-matching is a recursive process this ability for output of the process to be accepted as input is essential.

The distinction between path-tracing and pattern-matching processes may be psychologically significant. Inhelder and Piaget (1964) find that prepuberty children cannot use logical equations such as $\sim(A \wedge B) \supset \sim A \vee \sim B$. The equations involve coreference and hence their application requires a pattern-matching process. It could be speculated that this more powerful process only appears at maturation.

DISCOURSE ANALYSIS

The Structure of Coherent Discourse

In this section the hypothesis concerning the kinds of organization present in coherent discourse is outlined. A fuller description can be found elsewhere (Phillips, 1975). The role of the encyclopedia in discourse is then exemplified.

A discourse is judged coherent if its constituent propositions are connected. Various types of cohesive links are observed in discourse: anaphoric, spatial, temporal, causal, and thematic. I will formally describe the structure of a well-formed discourse in terms of these connectives.

Anaphora

A discourse has reference to objects. Coherence is given by repetition of the reference. Two kinds of anaphora can be distinguished. The first is marked by the presence of a proform (or by repetition of the form): [It is usual for coherent discourse to ex-

hibit several kinds of cohesive links. Thus the examples invariably contain more than the one specifically being illustrated.]

(10) Henry travels too much. He is getting a foreign accent.

Antecedents may be nominal, verbal, or clausal. The second kind of anaphora has a dependent ^{that} is an abstract term for the antecedent, for example

(11) John put the car into "reverse" instead of "drive".
The mistake cost him \$300 to repair.

"Mistake" in (11) is an abstract characterization of the gear selection expressed in the first sentence. Nagao and Tsujii (1976) address this issue.

A conventional way to label the recurring characters in discourse is as "dramatis personae". However, cohesion can result not only from multiple appearances of people (10), but of any concept, as in (11).

Spatial, Temporal, and Causal Cohesion

Space, time, and cause give coherence to a set of clauses or sentences.

(12) The King was in the counting house, counting out his money.
The Queen was in the parlor, eating bread and honey.

The actions in (12) are set in different rooms, but of the same "palace".

(13) After Richard talked to the reporter, he went to lunch.

The temporal sequence of events in (13) is expressed by "after".

(14) John eats garlic. Martha avoids him.

To non-aficionados garlic is known only for its aroma, detection of which causes evasive action.

Cause, illustrated in (14), is an important discourse connective (Schank, 1975b). The importance is perhaps ethnocentric; in other cultures different positions may have to be taken, for example. a teleological world view (White, 1975).

The causal chain of propositions in discourse is termed its plot structure.

Thematicity

Coherent discourse is expected to have a theme, to have a topic.

For example

- (15) DF drowned today in MB reservoir after rescuing his son who had fallen into the water while on a fishing trip.

is a news story from the New York Times with a theme that I will call "tragedy". In this section I wish to justify the claim that a thematic structure condition is universal by examining different examples and analyses of general discourse for evidence.

The notion of theme is much used but not often defined with clarity. It is variously stated to be "The subject of discourse . . . a topic" (Oxford English Dictionary); "the playwright's point of view towards his material" (Mabley, 1972, p. 14), etc. In Abelson (1973) there is a list of themes (admitted to be neither fixed nor exhaustive): admiration, devotion, appreciation, cooperation, love, alienation, betrayal, victory, dominance, rebellion, mutual

antagonism, opposition, and conflict. Occasionally one finds overt comment on the lack of a theme: "The thing that puzzled me most about The Last Remake of Beau Geste was its lack of a point of view" (Barry Took, "Cinema", Punch, December 7, 1977). Equally infrequently one can find a succinct amplification of the structure of a theme: "On the other hand, the suspension of disbelief is what thrillers are about." (Sheridan Morley, "Theatre", Punch, November 19, 1975).

A theme may be explicitly stated in discourse. In technical writing it is quite usual to express a complete definition, definiendum-definiens: Kuhn (1962) defines "paradigm" as an "achievement" that is "sufficiently unprecedented to attract an enduring group of adherents away from competing scientific activity . . . [and] sufficiently open ended to leave all sorts of problems for the redefined group of practioners to resolve" (p. 10). Much of the rest of the book then discusses paradigms as models for scientific revolutions.

If a discourse has an implicit theme, it has to be inferred by the reader. An author, therefore, should use themes that are known to the reader. One possibilty is that there is only a finite number of themes. But lacking evidence for this position, I will hypothesise that the number of themes may be unlimited in the same way that the vocabulary of a language is open. A reader may not know a word that is used by an author; in a similar fashion he may not recognize a theme.

There are studies that indicate the existence of abstract themes in language. In folk-tales, Propp (1968) analyses a reader's expectations about the structure of the tale. Propp starts by comparing the following events from different tales:

1. A tsar gives an eagle to a hero. The eagle carries the hero away to another kingdom.
2. A princess gives Ivan a ring. Young men appearing from out of the ring carry Ivan into another kingdom.

Propp infers that "a tale often attributes identical actions to various personages. This makes possible the study of the folk tale according to the functions of the dramatis personae" (p. 20). Folk tales are analysed in terms of functions. The above examples are described as containing two functions: "Acquisition of a magical agent" and "Transference to a designated place". An example of Propp's analysis is

(22)	<u>ACTION</u>	<u>FUNCTION</u>
	A tsar, three daughters.	INITIAL SITUATION
	The daughters go walking,	ABSENTATION
	overstay in the garden.	VIOLATION
	A dragon kidnaps them.	VILLAINY
	A call for aid.	MEDIATION
	Quest of the three heroes.	CONSENT TO COUNTERACTION
		DEPARTURE
	Three battles with the	
	dragon.	STRUGGLE
		VICTORY
	Rescue of the maidens.	INITIAL MISFORTUNE
		LIQUIDATED
	Return.	RETURN
	Wedding.	WEDDING

(p. 128)

Functions correspond to metalingually defined concepts of the encyclopedia. Propp shows that this genre of discourse can be analysed as an ordered string of abstract concepts.

Linde (1974) finds that there is a prescribed pattern in verbal descriptions of apartments. Only two discourse strategies are used by her subjects to express the spatial structures, and of these, one is considerably more frequent than the other:

There are at least two logical possibilities for . . . [the overall description of apartment layouts] . . . the speaker may describe a map of the apartment, or he may describe a tour of it. Examples of each are the following:

I'd say it's laid out in a huge square pattern, broken down into four units. If you were looking down at this apartment from a height, it would be like . . . like I said before, a huge square with two lines drawn through the center to make four smaller squares. Now on the ends . . . uh . . . in the two boxes facing out on the street you have the living room and a bedroom. In between these two boxes you have a bathroom. Now between the next two boxes, facing the courtyard you have a small foyer and then two boxes, one of which is a bedroom and the other of which is a kitchen and a small foyer a . . . a little beyond that.

Well you walk in the door and there's a kitchen and then off the kitchen is one bedroom. As you go straight in from the doorway through the kitchen you go into the living room. And then to the left of the living room are two bedrooms. The two bedrooms are on the same side of the building and the living room and the kitchen are on the same side of the building.

Both of these descriptions are reasonable answers to the question "Would you describe the layout of your apartment?" Our intuition certainly informs us that both speakers have fulfilled the task that was proposed to them. What our intuitions do not tell us is that descriptions like [the first] are extremely rare, while descriptions like [the second] are extremely common. Of 72 apartment descriptions, only 3 are of the form of a map . . . while 69 are the form of a tour (pp. 8-9)

The tour may be a composition of separate episodic events of moving between rooms of the apartment. The plan is more obviously systemic,

involving spatial (left, right, etc.) and componential (part-whole) organization.

Longacre (1968) notes that in a given language there is a finite number of discourse types which can never be mixed or confused. Discourse from various Philippine languages suggest four contrasting discourse prose genres:

Narrative: recounts some sort of story
Procedural: tells how to do something.
Expository: any sort of explanatory essay.
Hortatory: attempts to influence or to change conduct.

Narrative discourse is composed of the following tagmemes:

+APERTURE +EPISODE +DENOUEMENT +ANTI-DENOUEMENT +CLOSURE +FINIS

APERTURE provides temporal and spatial setting and introduces some of the principal dramatis personae. CLOSURE gives final commentary on the main participants, "they lived happily ever after". Nuclear tagmemes EPISODE, DENOUEMENT, and ANTI-DENOUEMENT show a great variety of exponence . . . typically any paragraph type may be an exponent - plus embedded discourse of the PROCEDURAL or EXPOSITORY genre.

A correspondence can be informally recognized between some of Propp's functions and Longacre's tagmemes. For example, between "Initial Situation" and "Aperture", and "Reward" and "Closure". For Propp the peak of the discourse is in the function "Initial Misfortune Liquidated", and for Longacre it is in the tagmeme "Anti-Denouement".

The idea of a hierarchic organization of tagmemes mentioned by Longacre, above, is paralleled in Lakoff's (1972) transformational generative model that uses Propp's set of functions. A phrase structure component generates a "deep structure". For example, the tale of (22) may be represented by the tree structure of Figure 17.




Figure 17 Textual deep structure

The conclusion is that there are prescribed patterns in all genres of discourse; I term these patterns "themes". I do not offer a complete inventory of themes; their discovery is a matter of empirical investigation.

Any extended discourse is unlikely to be organized according to a single theme. I hypothesise that a coherent discourse is characterized by a single rooted tree of themes, as schematized in Figure 17. All themes must be proper subthemes of the matrix theme. A text with an overlapping thematic structure is incoherent:

(23) Eating fish made John sick. He caught measles last May.

shown schematically in Figure 18.




Figure 18 Incoherent thematic structure

An important point to conclude this section. The inferred connections may not correspond with those intended by the author. This is another problem. Here I only address the analysis of a story by a reader. If he connects it in the manner described above, then it is coherent for him.

The Role of the Encyclopedia

Not all of discourse structure is overtly stated; discourse is highly elliptic. In (13) the discourse connective "after" is present to mark a temporal sequence, but in (14) there is no realization of the causal relation between the two propositions. Normally one assumes that a discourse is coherent; hence (12) is most acceptable if the rooms are taken within the same habitation. Evidently a reader must infer omitted structure. The inferences are made from his cognitive store of world knowledge.

There is much discussion at present about inference as part of understanding. To make inferences is easy; the problem is to make the right ones. It helps to have a goal. It is suggested that discourse can be said to have been understood when it has been judged coherent, as defined above.

In the next sections are presented the role of the encyclopedia in determining and representing the dimensions of coherence spelled out above.

Anaphora

If the dependent is a proform then part of understanding is to determine the correct antecedent. There are syntactic constraints (Langacker, 1969) which serve to narrow down choices for antecedents and to give an order of preference. Winograd (1971) also established an ordering for the choice of antecedents. Nash-Webber (1976) used lambda abstraction to establish possible antecedents. The chosen antecedent, when substituted for the proform, must produce a meaningful proposition that is coherent in context. A meaningful proposition is one that has a counterpart in the encyclopedia. Wilks (1975) discusses a method of finding the most semantically acceptable antecedent. In encyclopedic terms, the counterpart may be the self-same proposition, or, more likely, a systemic proposition. The process of finding such a proposition has been described earlier. If no generalization is found, the input proposition is not consistent with encyclopedic knowledge.

Abstract terms can be defined by complexes of general propositions, each having sufficient conceptual content to define situations in which they apply. For example, a definition of "mistake" must be such that it applies to part of the first sentence in (11). The process of abstraction needed here was presented above.

Spatial, Temporal, and Causal Cohesion

To infer omitted spatio-temporal and causal relations, i.e., the discursive relations of the encyclopedia, it is also necessary to locate general propositions. Systemic memory, of course, includes these relations. Schematically, Figure 19, from a discourse proposition P_1 we can locate P_2 , by the means already described. P_2 may have a discursive relation R to another systemic proposition P_3 . A proposition P_4 , a particularized version of P_3 , and the relation R , between P_2 and P_3 , can be added to the discourse. Often P_4 will be a proposition already stated in the discourse, then merely the relation need be inferred to augment the plot structure.




Figure 19 Inference of discursive structure

It may may, however, be necessary to infer a chain of propositions to link the propositions of the original discourse. Intuitively there must be a limit on the number of propositions that can be inferred in a sensible path, but at present no insight can be offered.

To exemplify the process in greater detail, let us consider some of the knowledge that is used in the analysis of (15): "In water and not able to act causes drowning". In Figure 20 the network form of this knowledge is presented.




Figure 20 Example of causal inference

From the discourse propositions "DF in water" and "DF cannot act", paradigmatic structure enables the systemic propositions A and B to be found. There is a coreferentiality condition that must be tested in the manner described earlier. The discourse propositions pass the

test, so the complex represented by the modality C exists in the discourse. The discursive relation cause can be followed from C to D. The latter is a plausible inference, and in fact, a specific equivalent of D is one of the original propositions of the discourse, i.e., "DF drown". The concepts of the systemic propositions are linked to the rest of the encyclopedia by typical arcs. It is so because of the nature of the knowledge is such that it is only something that could happen in the given circumstances.

The indications from the testing of Thorndyke (1976) are that inferences are a psychological reality in understanding natural language texts.

Thematicity

In the present system, a thematic concept is defined structurally, it is anything having a metalingual definition. A theme is therefore a complex of generalized propositions. The process of detecting the applicability of abstract terms, and hence of finding themes, is abstraction, described above. Abstraction is a recursive process and is thus one way to capture the embedding of themes hypothesised to exist in discourse.

The paraphrases elicited by Mandler, Johnson, and DeForest (1976) and Rumelhart (1976) show that subjects do create descriptions of texts that vary in abstractness in accord with the hierarchy of themes proposed here.

IMPLEMENTATION

In this section the implementation of the structures and processes presented above is described. The original program was written in SNOBOL for a CDC6400.

In a complete system there should, of course, be a parser. For the present this does not exist; the system only embodies the cognitive component. This means that the overall organization is not as it would be in a complete text analysis system, where interaction between the syntactic and semantic components is essential (Woods, 1971; Schank, 1975a; Winograd, 1971; Erdman, 1975). The justification for this omission is that for the present I am seeking to establish only the nature of the semantic organization of a coherent discourse. Once this structure has been identified it will provide the goal for a complete system.

Input to the system is accordingly in a cognitive form that retains the logical ellipsis of the surface form.

Most of the processing is performed by "Normalizer" which infers omitted logical and thematic structure. A judgement of coherence is then a simple task: if the discourse is not logically connected or does not have a single theme, then it is incoherent; otherwise the matrix theme indicates the topic of the discourse.

Processes

A component of all processes is a breadth first path-tracing routine, called "Ripple". A search path is defined by a sequence of arc types. A path does not explicitly state whether an arc can be repeated. The network is assumed to be syntactically well-formed and this controls repetitions. An arc can be marked as obligatory; otherwise it is optional. A goal of the search can be defined. This may be a particular node, or a node marked with a specified "activation tag", i.e., a node reached by another path, when seeking an intersection.

Paradigmatic path-tracing

Paradigmatic path-tracing is implemented by Ripple with a path sequence VAR IST TYP MAN (see Figure 11). A converse paradigmatic path is $\overline{\text{MAN}} \overline{\text{TYP}} \overline{\text{IST}} \overline{\text{VAR}}$. The properties associated with a varietal concept may be found by Ripple with a path TYP MAN starting from the concept.

Causal connectivity condition

This process uses Ripple with cause as the path definition. It also has to include P-W and $\overline{\text{P-W}}$ to be able to reach from and to conjuncts.

Discovery of general and specific propositions

All propositions of a discourse must match general propositions in the encyclopedia. The procedure is to make cyclic calls of Ripple. The first is from the modality node of the discourse proposition. Each node reached, other than the modality, initiates another search in the

encyclopedia. For example, given the discourse and encyclopedia of Figure 12, the process is as follows: from "gobble" node 1, a converse paradigmatic path plus a typical arc plus manifestation is followed to, for example, node 2 in the encyclopedia. Rippling from "gobble" in the discourse gives nodes "Marv" and "caviar". The syntagmatic arcs traversed are noted. From "Marv", node A, a converse paradigmatic path plus typical plus manifestation plus converse agent is followed, with a goal of a node activated from the prior discourse node, i.e., "gobble". [Not all of these arcs have to be present, they are optional except for the syntagmatic arc.] Node 2 satisfies this goal. From "caviar" node B, a converse paradigmatic path plus typical plus manifestation plus converse objective is followed with a goal of a node activated from the prior discourse node "gobble". Again node 2 satisfies the goal. Thus the proposition at 2 is a generalization of the discourse proposition.

The condition on an acceptable generalized match is that it must contain all the syntagmatic information of the discourse proposition; the generalization may contain more information but it cannot contain less. Separate searches are made for syntagmatic structure and for spatio-temporal information on the modality of a proposition.

It is only necessary to change the path description from that used in Figure 12 from converse paradigmatic path to paradigmatic path for the routine to locate more specific propositions.

Metalingual decomposition

The search for general propositions also flags nodes that have metalingual definitions. New propositions having the structure of the definiens are made by copying the definiens but with node names drawn from the proposition that is being paraphrased. These new propositions are then considered as part of the discourse. To make the copy, the breadth first search routine is used to pass through the definiens. For each node and arc in the definition, an equivalent structure is created. The end of scanning a proposition of the definiens is marked by reaching a typical arc, i.e., at the point at which the definition is linked into paradigmatic organization. If a participant in the generalized proposition matches a participant in the discourse proposition then this participant fits into the corresponding role slot in the definiendum, otherwise the concept in the definiendum is used. For example, given "Peter buys a bicycle from Jane" and the definition of "buy" as in Figure 16. In locating the systemic definiendum, the correspondences of "Peter" to "A", etc., are also found. When copying reaches the node that matched "Peter", this name is inserted into the paraphrase. There is no correspondence for the instrument, so "money" is the inserted from the definiendum.

Metalingual abstraction

In searching for general propositions, some may be found that are components of metalingual definitions. These have modalities that are pointed to by a part-whole arc.

The process that tests coreference and contextual requirements uses Ripple to traverse in parallel the candidate discourse propositions and those of the definiens. Typical arcs in the definiens limit the search. Each node of the definiens is compared with equivalent nodes in the discourse propositions at each step. A proposition is rejected if a node has no equivalent or it does not possess all the properties of the nodes of the definiens, including arcs to nodes matched in the previous step, i.e., if nodes X (systemic) and Y (discourse) were taken as corresponding nodes at one step, then if on the next step a node of the definition has an arc to X then the discourse must have the same arc to Y. Only those propositions that match the definiens will not have been rejected and can be rewritten using the abstract term. The process can be illustrated using the definition of "poison" given in Figure 14b, and its application to the discourse in Figure 14a. The crux of the test is at that node of the definition having the two manifestations arcs emanating from it. If the discourse proposition did not have two manifestations it would be rejected. This is how "John's eating the worms made Fred sick" is eliminated. Or if it does have two manifestations, they must point to nodes that were satisfied on the previous step of the comparison. Thus if one of the manifestations pointed into another proposition the test would fail.

Inference of omitted discursive relations

A search along discursive arcs in systemic memory from counterparts of discourse propositions may lead to a proposition that is flagged as a generalization of another discourse proposition. If this is so then the discursive arc may be added between the discourse propositions. If the proposition reached is not flagged then it and the discursive arc are copied, and added to the discourse. The copying routine was given in the discussion of decomposition.

The System

The flow chart of the analysis is shown in Figure 21. The meanings of the annotations are:

OLDINFO has a discourse proposition as its argument. It finds systemic equivalents. It calls a routine SPACETIME to compare spatio-temporal contexts. SPACETIME is also called during the search for general propositions when a non-event node is found with an attached modality. If OLDINFO is presented with a modality that has only part-whole relations to other nodes, it does nothing.

LOGCON has a systemic proposition as its argument. It succeeds if it finds a link to a general proposition corresponding to a proposition of the discourse (including propositions added by inference). It also generates INTERLIST, a list of causal inferences from propositions of the discourse.

MLST is a list of nodes found to have metalingual definitions.

CONJLIST is a list of conjoined propositions. When a discourse proposition is matched against the encyclopedia, it sees if the encyclopedic proposition is a constituent of another modality. A CONJUNCTION TEST routine uses CONJLIST to locate discourse propositions that can be grouped.

TRANSFORM has two modes. In one it is used to decompose propositions that contain a metalingually defined concept. A second mode creates causally inferred propositions.




Figure 21 Flow chart of the system

ANALYSIS OF SOME STORIES

I want to show that abstract patterns are quite general, that all linguistic behavior is based on such patterns. Obviously such a claim

must be substantiated by the discovery of such patterns. A number of stories of drowning were used to test this hypothesis. The second claim of proper embedding of themes was also tested by a more complex drowning story.

In the examples a refined hypothesis of discourse connectedness is used. One habit in discourse is to set the stage (Propp's "Initial Situation", Longacre's "Aperture"). In terms of the model this aspect should be recognizable by the occurrence of space and time relations. We find "today", "in MB reservoir", "On October 11, 1974", "DF of Queens", etc. A greater structural complexity of expression is to be expected elsewhere in the stories (see Longacre's comments on nuclear tagmemes, above). Longacre (1972) includes in the natural outline of a discourse, recognition of a peak within the discourse. Various surface markings for the peak are given: tense change, extra long sentences, rhetorical underlinings, etc. Taking an ethnocentric view of the world (cf. White's teleological commune), it is suggested that in the underlying form, the peak will lie within causally related propositions. It is thus expected to find the theme within the causal structure and so I focus on this organization. This would be inappropriate if the stories were descriptions of the kind elicited by Linde, above.

Common patterns

Short factual accounts of drownings were elicited from freshmen in Linguistics and English. The instructions given sought only to define a topic and an approximate length: "Write a drowning story that, for example, you might expect to find as a column filler in the New York Times." A sample is

(Story 1) The body of Horatio Smith was found last night in the Niagara River. He was drowned when his boat overturned on the river.

The hypothesis formed is that an acceptable drowning story must give the following information:

- (a) Why the victim was in the water.
- (b) Why the victim was not able to save himself.

The rationales for these requisites are:

- (c) A person is not usually found in water, and therefore some explanation of this location is expected.
- (d) By an instinct of self-preservation, one would expect the victim to try to extricate himself from his predicament. The story should say why he couldn't.

Figure 22 shows the cognitive form of this requirement.




Figure 22 The drowning theme

The empty modalities indicate that a matching story must have something that stand in a causal relationship to the other propositions, i.e., explain why they happened (what caused them). If not originally explicit, this information must be recoverable through encyclopedic knowledge.

One way in which the content of systemic memory may be substantiated is by examination of negative propositions in the stories; writers presumably only need to negate normal expectancies, and this is what generalizations are. Unfortunately there are few negatives in the stories. Stories 31 and 5 indicate the assumption of swimming ability. Stories 12, 32, and 42 show that the success of rescue attempts is anticipated.

The stories elicited fell into several categories. The two analysed here are a boat capsizing and a fall into water. In spite of surface difference it will be shown that at an abstract level the stories conform to the thematic pattern. The analysis of one story from each category is presented.

Stories not analysed include such happenings as a person eating too much then going for a swim; Jesus freaks trying to walk on water; and water-skiers having accidents while watching bikini-clad occupants of passing boats.

A boat capsizing

(Story 1) The body of Horatio Smith was found last night in the Niagara River. He was drowned when his boat overturned in the river.

(Story 2) Eggbert Willis, 56, of Bayside, drowned this morning after the boat he was rowing overturned near Devil's Cove.

(Story 3) The body of John Smith, 58, was discovered today at the foot of West Ferry Street. He was reported missing four days ago by his wife after he failed to return from a boating trip. His boat had capsized. Death was due to drowning.

(Story 31) A small sailboat was afloat on a calm peaceful lake when suddenly the mast of the vessel struck some cables overhead and the boat capsized. The two men aboard drowned, one because he was hit by the boat and rendered unconscious, the other didn't know how to swim.

Story 1 is analysed. In it, some of the causal and thematic structure is absent and is reconstructed using the following knowledge:

- (i) If a person is in a boat and the boat overturns, this may cause him to be injured and to be in the water.
- (ii) If injured a person may not be able to "act".
- (iii) If a person is in water and cannot "act" then he may drown. Figure 18 shows the encyclopedic form of this piece of knowledge.

The nodes appearing in the encyclopedic entries for all of these facts exemplified in this section are linked to varietal nodes by typical arcs. This is so because the facts are not obligatory on some concept but are something that may happen to some examples of this category some of the time.

In the analysis of the collection of drowning stories, only the parts of the story that are relevant to the drowning are considered. For example, in Story 1, only the second sentence is processed; the first deals with an event after the death and consequently is excluded.

The result of the analysis is shown in Figure 23. The original propositions of the discourse are:

Boat contains Horatio Smith
Boat overturns on Niagara River.
Horatio Smith drowns.

The antecedent conditions expressed in a general form in (i) match the specific situation in Story 1. Thus it is inferred that Horatio Smith is in the water and that he is injured. From (ii) it follows that he is not able to act. By being in the water and not able to act, he can drown, a fact stated explicitly in the story, as shown in Figure 23.

connected structure. The subsequent backup then takes the correct path.

A fall into water

The second category of drowning story requires the addition of the knowledge that "If a person falls, he may injure himself". Ten stories in this category are listed below.

(Story 5) Early this morning, James R. Smith, age 7, was found in a swimming pool near his home. Investigators say the boy stumbled into the pool in the darkness early this morning whilst looking for his pet kitten. Unable to swim, the boy drowned.

(Story 7) At the home of Mrs. John Smith on Elmwood Avenue, a boy, Mark, age 15, drowned in his pool. The boy was with two other friends. They were performing water stunts when Mark fell and smashed his head on the bottom of the pool.

(Story 12) Yesterday afternoon, the life of a Buffalo youth was taken when he slipped on rocks at a local quarry. The failure of attempted rescues resulted in the drowning of Michael Smith, age 7, of 29 Oak Street, Buffalo.

(Story 19) A 12 year-old boy was found drowned in Ellicott Creek. Sources say the boy ran away from home and fell accidentally into the water.

(Story 26) A 10-year old boy died last night when he fell into a small pond. His friends say he was chasing his parakeet which had escaped from its cage, when the incident occurred.

(Story 32) Steve Smith, of Hickstown, drowned today while sailing on Glasslyke Lake. Mr. Smith, who was knocked overboard when struck on the head by a seagull, perished before help could reach him. His son Edgar's attempts to save his life proved futile.

(Story 37) An unidentified man was seen by several persons falling into the Niagara River at the foot of Ferry Street. He was later pulled from the water and pronounced dead at the scene. The cause of death was drowning.

(Story 38) Today, the world's greatest swimmer died. John Whale was preparing to take a bath when he tripped and fell into the bath. Cause of death was drowning.

(Story 40) On October 11th, 1974, an unidentified man drowned in his bathtub at the Hotel Sheraton. The drowning was due to the fact that he fell into the tub in trying to make himself sober.

(Story 42) An 11 year-old boy drowned today after falling into the canal where he and his friends were playing. The two other boys, both eleven, tried to save their companion but were unable to do so.

(Story 4) A body was found early yesterday at the foot of the Mango River, near Clubsport. The body is believed to be that of Jose Gepasto. It seems as if Mr. Gepasto's car made a wrong turn on the highway and plunged into the water.

Story 4 is analysed. Note that it does not explicitly mention the motion of the person, only that of the car. Understanding the story requires that it be known that:

- (a) If a person is "contained" by something that falls, then he also falls.
- (b) If a person is "contained" by something that is in contact with something (e.g., water), then the person is in contact with the something (water) too.
(But not if the something is a submarine!)

Further it is not given that Jose Gepasto drowns. This can be inferred, but the inference chain is open ended. The analysis continued making causal inferences and conjunctive groupings, some of which led to the discovery of the theme. Only when the system ran out of logical and conjunctive possibilities, did it make the connectedness test. Figure 24 shows the network developed in the analysis.




Figure 24 A fall into water

Embedded themes

Story 22, in fact taken from the New York Times, looks like a drowning story, and as shall be shown, does contain this theme. However, it contains more. The claim is made that it is a "tragedy".

(Story 22) DF, 43 years old, of Queens, drowned today in MB resevoir after rescuing his son D, who had fallen into the water while on a fishing trip at TF, near here, the police said.

This theme is defined in Figure 6 as a situation in which "Someone does a good act (e.g., rescue) and dies (e.g., drowns)". It will be seen that the tragedy is a proper subtheme of the drowning theme. Thus,

though the story may be said to have two themes, one is part of the other, and by our hypothesis, the discourse is still coherent. At each step the encyclopedic knowledge used in the inference and an outline of the inferred nodes are indicated. Figure 25 shows an outline of the evolved structure, where the original discourse propositions are shown by ● and inferred propositions by ○.




Figure 25 Embedded themes

- Step 0. Initial state. (Nodes 1, 2, 3, 4).
- Step 1. Fall causes injury. (Node 5).
- Step 2. Injury causes inability to act. (Node 6).
- Step 3. In water and not able to act causes rescue. (Node 7 and a link to node 3).
- Step 4. To rescue someone who is in the water, get into

- the water. (Nodes 8, 9).
- Step 5. Acting causes weariness. (Node 10).
- Step 6. Weariness causes inability to act. (Node 11).
- Step 7. In water and not able to act causes drowning.
(Node 12 and a link to node 4).

Note that the antecedent condition in Step 3 is the same as in Step 7. Both resultant situations are possible and are noted. The system can select either. However, the wrong choice does not lead to a connected structure and backup to the alternative has to be made.

After Step 7 the discourse has an inferred causal structure connecting all the original propositions.

The theme "tragedy" fits, the rescue is a (partial) cause of the demise. Rescue is a variety of act and good can apply to it and drown is a variety of die. The drowning theme is also present.

Although the drowning theme is not defined in terms of the tragedy, it can be seen that one is properly embedded in the other. The process that performed the analysis is at present incomplete because the notion of embedding is not well understood for the highly structured network. The process used the transitivity of cause and the conjoining of propositions. Thus the tragedy encompasses propositions 3, 10, 11, and 4 and the drowning 8, 9, 10, 11, and 4. The transitivity of cause lets the chain 3, 10, 11 be equivalent to the chain 10 and 11.

A postmortem on this example reveals a serious flaw. As can be seen the rescue is the cause of the father being in the water. The

analysis has failed to distinguish desire for an action, a goal, from the execution of the action. A more satisfying analysis would include some of the mechanisms to be found in the robot planner of Furugori (1975). Step 2 should be seen as setting up the conditions for the son to drown, which is an event that should be prevented. This provides a goal for the subsequent activities. One way to prevent someone from drowning is to save him, this is a subgoal that would directly achieve the goal. If you want to rescue someone who is in the water, then it may be necessary to get into the water. With this subgoal included, the goal can be achieved, and the analysis resumes at Step 5. Figure 26 shows this preferred analysis of this fragment of the story. This would not change the relative status of the themes.




Figure 26 Improved story analysis

DISCUSSION

Much of the representative power of the encyclopedia is unused in the system for discourse analysis and therefore remains to be tested and evaluated. There is also not at present a parsing system to effect transduction from surface to encyclopedic form. The methodology is first to try to establish an adequate conceptual representation which provides a goal for a complete system. Although one example demonstrated the embedding of themes, it did not exhibit recursive abstraction. Further examination of themes in discourse should overcome this.

There are two aspects of the encyclopedia, paradigmatic and metalingual organization, that set this model apart from any other current system. Discussion will be directed to comparative comments on these aspects.

It is evident that the present system makes much use of paradigmatic organization. Yet Schank (1975a) seeks to minimize the need for this kind of knowledge. His conclusion arises from the observation that people do not make responses based on paradigmatic associations, but rather on episodic associations. This is not telling evidence against the existence of such structure, rather it may say something about the cognitive process of free association. In Schank's system there is no need for paradigms at the level of conceptual representation as words are transformed into conceptual primitives by the conceptual parser. The parser thus contains knowledge that is functionally equivalent to paradigmatic structure.

The question then to ask is whether having a single level of representation of concepts, the primitives, is the most beneficial for conceptual processing. I would claim not for two reasons.

Firstly there is the presence of thematic structure in discourse. Metalingual organization enables the content of a text to be thematically described at many levels of preciseness. It is possible to be quite superficial, or by decomposition to become more and more detailed, or vice-versa using abstraction. The depth of analysis can be determined by the requirements of understanding a given text: essentially abstracting or decomposing until causal links are established over the text. It is not apparent that definitions of themes can be controlled in Conceptual Dependency Theory, in that if stated in terms of primitives, each abstract term could become extremely large. In contrast, the encyclopedia can define themes in terms of lesser themes.

Secondly paradigmatic structure enables comparatively small chunks of knowledge, say involving a single causal relation, to be retrieved and pieced together to complete the underlying form of discourse. Rather than attempting to patch general knowledge, Schank and Abelson (1975) have introduced "Scripts", large preformed knowledge structures whose function is to limit the possible inferences in understanding. This has a danger of essentially idiomatizing understanding, with a consequent difficulty in handling deviant situations. And as Wilks

(1976) points out, one problem with Scripts is that they are invoked in their entirety by word association. Thus it is suggested that, for example, "I bought some beer from the supermarket, drove home, and drank it while watching a football game on television" would evoke a multitude of Scripts by the presence of such words as "supermarket", "drive", "football", etc. Hence the desired reduction of possible inferences is not achieved. Paradigmatic organization enables recognition of higher level structures, including propositions that are part of metalingually defined concepts. Partially recognized abstract terms can be used to predict their completion. The encyclopedia thus has general, productive bottom-up and top-down capabilities.

Even though an abstract definition should be activated by the appearance of an appropriate word in the text, the structure will not in general be large, and so not produce an overwhelming number of extraneous active nodes.

On the other hand it is certainly advantageous to have a multitude of overt themes in analysing discourse. Searches can be initiated from these terms as well as from the more specific discourse propositions. To illustrate this, consider an exhaustive undirected search for a goal n states distant in a space where each state is linked to m other states. The number of nodes activated will be of the order m^n . If the goal is known then a bidirectional search will reduce the exponent to $n/2$. But a more significant reduction takes place if there are

stated intermediate subgoals, i.e., themes in the hierarchy, say g of them, with an average separation of n/g , the exponent becomes $n/(2g)$.

ACKNOWLEDGEMENTS

I would like to thank David Hays for his continual support and guidance throughout this work. Ray Bennett and Randy Walser helped in trying to make this discourse coherent, the faults that remain are mine.

REFERENCES

- Abelson, R. P. The structure of belief systems. In R. C. Schank & K. M. Colby (Eds.), Computer Models of Thought and Language. San Francisco: W. H. Freeman, 1973.
- Abelson, R. P. Concepts for representing mundane reality in plans. In D. G. Bobrow & A. Collins (Eds.), Representation and Understanding. Studies in Cognitive Science. New York: Academic Press, 1975.
- Berlin, B., Breedlove, D. E., & Raven, P. H. Covert categories in folk taxonomies. American Anthropologist, 1968, 70, 290-299.
- Black, F. A deductive question-answering system. In M. Minsky (Ed.), Semantic Information Processing. Cambridge: MIT Press, 1968.
- Bobrow, D. G. Natural language input for a computer problem-solving system. In M. Minsky (Ed.), Semantic Information Processing. Cambridge: MIT Press, 1968.

- Chomsky, N. Aspects of a Theory of Syntax. Cambridge: MIT Press, 1965.
- Colby, K. M. Simulation of belief systems. In R. C. Schank & K. M. Colby (Eds.), Computer Models of Thought and Language. San Francisco: W. H. Freeman, 1973.
- Collins, A., & Quillian, M. R. How to make a language user. In E. Tulving & W. Donaldson (Eds.), Organization of Memory. New York: Academic Press, 1972.
- Conklin, H. C. Lexicographical treatment of folk taxonomies. In F. W. Householder & S. Saporta (Eds.), Problems in Lexicography (Publication No. 21). Bloomington: Indiana Research Center in Anthropology, Folklore, and Linguistics, 1962.
- Erdman, L. D. Overview of the Hearsay speech understanding research (Computer Sciences Review 1974-1975). Pittsburgh: Carnegie-Mellon University, 1975.
- Fillmore, C. J. Toward a modern theory of case. In D. A. Reibel & S. A. Schane (Eds.), Modern Studies in English: Readings in Transformational Grammar. Englewood Cliffs, NJ: Prentice-Hall, 1969.
- Furugori, T. A memory model and simulation of memory processes for driving a car (Technical Report No. 77). Buffalo: State University of New York, Department of Computer Science, 1975.
- Halliday, M. A. K., & Hasan, R. Cohesion in English. London: Longman, 1975.

- Hays, D. G. Linguistic foundations for a theory of content analysis. In G. Gertner, B. O. R. Holsti, K. Krippendorff, W. J. Paisely, & P. J. Stone (Eds.), The Analysis of Communication Content. New York: Wiley, 1969. (a)
- Hays, D. G. Applied Computational Linguistics. In G. E. Perren & J. L. Trim (Eds.), Applications of Linguistics. Cambridge: Cambridge University Press, 1969. (b)
- Hays, D. G. Linguistic problems of denotation. In M. Bierwisch & K. E. Heidolph (Eds.), Progress in Linguistics. The Hague: Mouton, 1970.
- Hays, D. G. Types of processes on cognitive networks. Paper presented at the International Conference on Computational Linguistics, Pisa, Italy, 1973.
- Hays, D. G. Cognitive Structures. Book in preparation, 1978.
- Hendrix, G. G. Expanding the utility of semantic networks through partitioning. Menlo Park, CA: Stanford Research Institute, 1975.
- Hopcroft, J. E., & Ullman, J. D. Formal Languages and their Relation to Automata. Reading, MA: Addison-Wesley, 1969.
- Inhelder, B., & Piaget, J. The Early Growth of Logic in the Child. New York: W. W. Norton, 1964.
- Klein, S., Oakley, J. D., Suurballe, D. J., & Ziesemer, R. A. A program for generating reports on the status and history of stochastically modifiable semantic models of arbitrary universes. Statistical Methods in Linguistics, 1972, 8, 64-93.

- Kuhn, T. S. The Structure of Scientific Revolutions. Chicago: University of Chicago Press, 1962.
- Lakoff, G. Structural complexity in fairy tales. The Study of Man, 1972, 1, 128-150.
- Langacker, R. On pronominalization and the chain of command. In D. A. Reibel & S. A. Schane (Eds.), Modern Studies in English: Readings in Transformational Grammar. Englewood Cliffs, NJ: Prentice-Hall, 1969.
- Linde, C. The linguistic encoding of spatial information. Unpublished doctoral dissertation, Columbia University, 1974.
- Longacre, R. Discourse, Paragraph, and Sentence Structure in Selected Philippine Languages. Santa Ana, CA: Summer Institute of Linguistics, 1968.
- Longacre, R. Hierarchy and Universality of Discourse Constituents in New Guinea Languages. Washington, D.C.: Georgetown University Press, 1972.
- Mabley, E. Dramatic Construction. Philadelphia: Chilton, 1972.
- Mandler, J. M., Johnson, N. S., & DeForest, M. A structural analysis of stories and their recall: From "Once upon a time" to "Happily ever after" (Technical Report No. 57). La Jolla: Center for Human Information Processing, University of California at San Diego, 1976.
- Nagao, M., & Tsujii, J. Analysis of Japanese sentences. American Journal of Computational Linguistics, 1976, Microfiche 41.

- Nash-Webber, E. L. Semantic interpretation revisited (Technical Report No. 3335). Cambridge: Bolt, Beranek, & Newman, 1976.
- Phillips, B. Topic analysis. Unpublished doctoral dissertation, State University of New York at Buffalo, 1975.
- Propp, V. Morphology of the Folktale. Austin: The University of Texas Press, 1968.
- Putnam, H. The meaning of "meaning". Mind, Language, and Reality: Philosophical Papers (Vol. 2). Cambridge: Cambridge University Press, 1975.
- Quillian, M. R. The teachable language comprehender. Comm. ACM, 1969, 12, 459-475.
- Raphael, B. SIR: Semantic information retrieval. In M. Minsky (Ed.), Semantic Information Processing. Cambridge: MIT Press, 1968.
- Raven, P. H., Berlin, B., & Breedlove, D. E. The origins of taxonomy. Science, 1971, 174, 1210-1213.
- Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. A process model for long term memory. In E. Tulving & W. Donaldson (Eds.), Organization and Memory. New York: Academic Press, 1972.
- Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory (Technical Report No. 55). La Jolla: Center for Human Information Processing, University of California at San Diego, 1976.
- Rumelhart, D. E. Understanding and summarizing brief stories (Technical Report No. 58). La Jolla: Center for Human Information Processing, University of California at San Diego, 1976.

- Schank, R. C. Conceptual Information Processing. Amsterdam: North-Holland, 1975. (a)
- Schank, R. C. The structure of episodes in memory. In D. G. Bobrow & A. Collins (Eds.), Representation and Understanding: Studies in Cognitive Science. New York: Academic Press, 1975. (b)
- Schank, R. C., & Abelson, R. P. Scripts, plans, and knowledge. Paper presented at the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, USSR, 1975.
- Schubert, L. K. Extending the expressive power of semantic networks. Artificial Intelligence, 1976 7, 163-198.
- Shapiro, S. C. The MIND system: A data structure for semantic information processing (Report R-837-PR). Santa Monica, CA: RAND Corporation, 1971.
- Simmons, R. F. Some semantic structures for representing English meanings (Technical Report NL-1). Austin: Computer-Aided Instruction Laboratory, The University of Texas, 1970.
- Sondheimer, N. K. Spatial reference and semantic nets. American Journal of Computational Linguistics, 1977, Microfiche 71.
- Thomson, P. W. The role of inferences in discourse comprehension. Journal of Verbal Learning and Verbal Behavior, 1976, 15, 437-446.
- Tulving, E. Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), Organization of Memory. New York: Academic Press, 1972.

- White, M. Abstract definition in the cognitive network: The metaphysical terminology of a contemporary millenarian community.
Unpublished doctoral dissertation, State University of New York at Buffalo, 1975.
- Wilks, Y. Grammar, Meaning, and the Machine Analysis of Language.
London: Routledge, 1972.
- Wilks, Y. A preferential, pattern-seeking, semantics for Natural Language inference. Artificial Intelligence, 1975, 6, 53-74.
- Wilks Y. Frames, scripts, stories, and fantasies. Paper presented at the International Conference on Computational Linguistics, Ottawa, Canada, 1976.
- Winograd, T. Procedures as a representation for data in a computer program for understanding natural language (Report MAC-TR-84).
Cambridge: Artificial Intelligence Laboratory, MIT, 1971.
- Woods, W. A. What s in a link: Foundations for semantic networks. In D. G. Bobrow & A. Collins (Eds.), Representation and Understanding: Studies in Cognitive Science. New York: Academic Press, 1975.