

TOWARD A STATE BASED CONCEPTUAL REPRESENTATION

Nick Cercone and Len Schubert

Department of Computing Science
University of Alberta
Edmonton, Alberta, Canada

Abstract

This paper discusses representations of semantic information derived from those of Schank (1972, 1973), Wilks (1973a, 1973b, 1973c), and Schank et al. (1973). A critical review of their related approaches is given that shows representational inadequacies for language comprehension in their theories. More generally, various problems in representing knowledge are investigated and methods suggested for handling some of them. These problems concern the representation of states, events, actions, and cases, as well as the handling of logical and natural language quantifiers, adverbials, fuzziness, and the meanings of complex concepts.

I. INTRODUCTION

The prima facie appeal of many natural language understanding systems diminishes under close scrutiny. This is due in part to a lack of detailed consideration of many classes of problems inherent in language understanding. Oftentimes an existing theory cannot be augmented to account for a new class of problems. Thus the original theory must be restructured or replaced.

To facilitate the formulation of a general theory of natural language understanding, a representational schema is required that has sufficient expressive power to represent the meaning content of ordinary language. This entails making exactly those predictions and inferences from the representation of an utterance that can be made from the original utterance. The following is a development of some ideas concerning the representation of individual items of "factual knowledge" in a computer, where this knowledge is thought of as being conveyed to the computer in natural language.

The next section enumerates what we believe to be the advantages of two related approaches to the problem of representing natural language in terms of meaning structures. These approaches are those of Schank (1972, 1973), Wilks (1973a, 1973b, 1973c), and Schank et al. (1973). We believe their work exemplifies the correct approach to the problem of representing the conceptual content of natural language utterances.

Subsequent sections outline several difficulties with Schank's and Wilks' theories. These difficulties concern the representation of states, events, actions, and cases, as well as the handling of logical and natural language quantifiers, adverbials, referential opacity, fuzziness, and the meanings of complex concepts.

II. TWO RELATED APPROACHES

The conceptual dependency approach of Schank and the preference semantics approach of Wilks have the following desirable features regarding knowledge representation. These features have not all been articulated by either Schank or Wilks and they would perhaps dispute this characterization of their respective approaches.

(i) Interpretive directness

the meaning structures corresponding to natural language utterances are formed according to simple structural rules. Powerful heuristic criteria, based on the central role of verbs and on preferred semantic categories for the subjects and objects of verbs, guide each choice in the creation of meaning structures. Interpretation of utterances then takes on a "slot and filler" character, rather than requiring extensive trial and error search.

(ii) De-emphasis of syntax

In ordinary discourse it would be absurd not to accept "ungrammatical" constructions such as dangling participles or fanciful locutions such as metaphor. Neither preference semantics nor conceptual dependency impose a syntactic "straightjacket" on admissible utterances. Therefore the abnormal is not excluded as it is in many linguistic systems.

(iii) Emphasis on events

A major part of our interpretative effort in understanding natural language is focused on events, i.e., time-dependent relationships. By contrast, "static" relationships in the world are relatively easy to understand. Therefore the search for fundamental semantic structures and primitives should concentrate on the representation of events.

(iv) Simplicity of primitives

Both Schank and Wilks have shown that there does exist a small, more or less adequate, set of action primitives through which a surprising number of action concepts can be expressed. A minimal set of semantic primitives should be sought that makes it relatively easy to use the meaning representations in a language (and paraphrase) independent way. Relatively few inference mechanisms will then be required.

We will now consider specific difficulties with these two related approaches and suggest possible ways of overcoming some of these difficulties. Perhaps the most important difficulties are those encountered when attempting to represent complex concepts and adverbials. However, other points need to be discussed first. In the following discussion we have in mind Schank's theory more frequently than Wilks'. However, Wilks' theory is quite close to Schank's in its general approach.

III. VAGUENESS AND UNCERTAINTY

Under this first heading we briefly confront two related, very pervasive problems, but offer no definite solutions.

Conceptual vagueness is a seemingly inescapable feature of nearly every ordinary concept, specifically vagueness in the extension of a concept (Is a 1000 foot hill a mountain or the Mississippi-Missouri a river?) and vagueness in the extent of an instance of a concept (where does a mountain or river end?).

Schank demands that his semantic structures express ideas unambiguously and precisely. Yet the concepts which are themselves the building blocks of conceptualizations are "vague" to a greater or lesser extent. Virtually any ordinary concept will serve as an example, e.g. house, mountain, river. Typically a house has rigid walls, a roof, windows, separate rooms, and is suitable for human habitation. Yet we can think of unusual houses where any one or more of these features is lacking. Each of these features contributes to "houseness" without being necessary. Similarly one can show that perceptually related modifiers such as "brown" or dynamic modifiers such as "walking" are vague (where is the boundary between walking and running? Is climbing a steep hill, or a ladder walking? Does a man with crutches walk? Do centipedes walk? Do seals walk? Et cetera). This conceptual vagueness inevitably carries over into conceptualizations. No longer can we say that it either is or is not the case that "John built a house" or "John was walking home"; all we can say is that these are more or less apt descriptions of what John was doing.

It is not clear to us at present how to allow for vagueness in conceptual representations. We expect to adopt Zadeh's (1965) approach, but have not followed this through. Zadeh associates a degree of truth with every proposition and calculates the truth value of compound propositions from

$$\begin{aligned}\tau(\neg p) &= 1 - \tau(p) \\ \tau(p \text{ and } q) &= \min(\tau(p), \tau(q)) \\ \tau(p \text{ or } q) &= \max(\tau(p), \tau(q))\end{aligned}$$

This is the standard truth value algebra of many-valued logic (e.g., Rosser and Turquette, 1952). However, we must be very circumspect in our application of this algebra to compound predicates. For example, the term "small" is hardly appropriate for a baby elephant, yet "small elephant" is entirely appropriate. Thus minimization of truth values for the conjoined terms fails. The reason in this case appears to be that most adjectives are operators on predicates (usually nouns) rather than being predicates themselves. Nevertheless adjective-noun combinations can often be analyzed as conjoined pairs of predications, as Bartsch and Vennemann (1972) have shown. Essentially they substitute comparative predicators for monadic adjectival operators. For example, a "small elephant" is something that is an elephant and is appreciably smaller than the average elephant. It seems that we could reasonably apply minimization of truth values to this conjunction.

Some further remarks on vagueness will be made in the sections on complex concepts and adverbials.

Statements can not only be more or less vague but also more or less credible, depending on the reliability of the sources and principles of inference on which they are based. Again we make no definite proposals for dealing with such credibilities. An obvious first step would be to allow credibility values to be attached to any proposition. But the real problem lies in combining credibilities of propositions, particularly "interdependent" propositions. Also there seems to be a close relationship between "fuzziness" and credibility (as we shall point out again in the section on complex concepts). For example the statement "John is young" is based on the fuzzy concept "young" but also appears to induce a credibility distribution on possible ages of John.

IV. LOGICAL AND NATURAL LANGUAGE QUANTIFIERS

Quantification needs to be included in any representation for natural language understanding for several reasons. First of all, logical and natural language quantifiers appear in discourse, for example "John sent the cards to all of his friends" and "Several of my friends were at the game today". Secondly, the use of quantifiers is required in general knowledge as in "It is always windy near tall buildings". Definite descriptions implicitly make use of quantification as the example "the people of China" shows. Lastly, the definitions of complex concepts require quantification. Any complex action concept like walking has associated with it as part of its definition assertions such as "at all times, some of the limbs of the individual engaged in walking support the individual".

There seem to be three methods in conceptual dependency theory for expressing universal quantification. The first method involves the use of variables assumed to be universally quantified, as in "if one smokes this may cause one to get cancer." Here "one" stands for any person. It is not clear whether a similar approach is envisaged for universal quantification over other (nonhuman) sets. In any case this device is inadequate, as it does not allow for multiple quantification, e.g. "Any politician can fool some of the people all of the time".

A second method for expressing universal quantification is the inference rule. We might have an inference rule in our system that determines "If X is thirsty, infer that X will drink something" where X is a universally quantified object that stands for any person. Thus a machine might easily answer a question like "John is thirsty. Will John drink something?"; however, the existence of this inference rule will not allow a machine to answer questions like "Will John drink something if he is thirsty?" and "Do thirsty people drink liquids?", since no assertion to the effect that someone is in fact thirsty has been made. The problem is that we lack accessibility to a procedurally encoded piece of knowledge as a fact. In other words, knowing how to use a fact does not guarantee knowledge of the fact.

The use of the conceptual tenses timeless and continuing is yet another method for expressing quantification. Schank appears to use timeless to designate habitual actions, for example, "John sells cars". The tense continuing is closely related. It is used in the sense of activity as defined in Evans (1967). However, these special devices do not address the general problem of quantification.

Basically the problem in semantic net representations of quantification is that of indicating the scopes of universal and existential quantifiers, which presents little difficulty in predicate calculus representations. This problem, and others that involve quantified objects of discourse, has been dealt with for semantic network representations by Schubert (1974) in a companion paper. The notation for the graphical representations is fully explained in that paper. Basically, the conventions that are used include: solid loops for proposition nodes and existentially quantified concept nodes; broken loops for universally quantified concept nodes; solid lines to link the parts of a proposition to a proposition node; dotted lines for dependency links joining each existentially quantified node to all universally quantified nodes on which it depends; and broken lines for logical links.

We can classify natural language quantifiers along a spectrum with absolute indicators of set size at one end, and comparative indicators of set size at the other, i.e. those comparing the size of one set to that of another set. The logical quantifier $\exists$ (there exists) belongs to the first of these categories, since $(\exists x)P(x)$ tells us that the set of P's contains at least one member. The quantifier $\forall$ (for all), by virtue of its equivalence to $\neg\exists$, can also be placed in the first category. In the context $(\forall x)(P(x) \Rightarrow Q(x))$ however, where the number of P's is finite, it can alternatively be placed in the comparative category. It tells us that the subsets of P's that are Q's is as large as the set of P's itself. Common absolute quantifiers are "none", "one", "two", "three", ..., "several"; common comparative quantifiers are "all of", "most of", "a small fraction of", "a slight majority of", "one-half of", "two-thirds of", "as many as", "twice as many as", etc. Some quantifiers show both absolute and comparative attributes, especially "some" and "many". For example, in "Many artificial satellites are orbiting the globe" "many" is used absolutely - it appears to imply a cardinality of at least about a dozen. In "Many students attend John's class" "many" is used in the sense "considerably more than attend the average class". This particular use of "many" is discussed quite satisfactorily in Bartsch and Vennemann (1972). They do not appear to be aware of the absolute indicativeness of "many", however, nor of its comparative use in selecting a subset of another set, as in "Many of the world's people are undernourished". Contrast the numerical indication here with that in "many of the apples in the basket were rotten".

In any case, recognizing the absolute/comparative behaviour of quantifiers, we can characterize them systematically by means of predicates on set cardinality and on pairs of set cardinalities as in the following examples. Let "n" denote set size (cardinality).




Fig. 1. "Several children were at the door"

We are regarding "n" as a function from sets onto integers and "several" as a (fuzzy) property of numbers. If instead we regarded "several" itself as a possible value of set size, then it would not be possible to talk about the size of the set, as "n" would be many valued (e.g., a 6-element set might have both size "6" and size "several"). In the next example "many" is expressed in two parts, the first being an absolute indicator of size (about a dozen

or more?), and the second comparing set size to an average set size, as in Bartsch and Vennemann (1972). In the construction "avg #" is regarded as a function on classes of sets whose value is the average cardinality of the sets in a given class of sets.




Fig. 2. "Many students attend John's class. Most of them are graduate students."

To conclude this section we briefly indicate the special notation for time. Moments of time, unlike other concepts, are represented as pairs of parentheses and intervals of time as pairs of square brackets. Propositional links to such time nodes are labelled I and TI respectively. The links may be suppressed altogether by placing times directly alongside the predicate tokens to which they apply. With this notation, complex time relations can be represented quite conveniently. The sentence "While he was in Rome, before he met his murderer, he first sang in La Traviata" is diagrammed in Figure 3. For universally quantified time we use broken parentheses or brackets, in keeping with the general quantification conventions.




Fig. 3. "While he was in Rome, before he met his murderer, he first sang in La Traviata."

V. STATES, EVENTS, ACTIONS, CASES, CAUSES, AND INTENTIONS

Conceptual dependency diagrams and preference semantics templates are capable of expressing four sorts of assertions:

- (1) states - ascribing a modifier to an object or set of objects at some time;
- (2) events - ascribing a change of state to an object or set of objects at some time;
- (3) actions; and (4) causes.

Nonetheless, a sentence such as

The sun was turning red and approaching the Western horizon. <5.1>

raises many questions. In <5.1> the motion of the sun must be done by somebody or something whereas its change of colour cannot be done by somebody or something. Thus, using the <actor-action-object> formalism espoused by both Schank and Wilks, modes of behaviour which are expressed by actions must have actors whereas all other modes of behaviour cannot have actors. In the case of the (apparently) moving sun in sentence <5.1>, one is hard pressed to identify the actor; similarly in the sentence "The breaker was moving toward shore." Consequently we are compelled to regard certain ongoing activities

which intuitively just "happen" as instigated by someone or something.

Just as we are compelled to regard certain ongoing activities as instigated by somebody or something, we are denied the option of regarding certain actions as having an agent as shown in

John was hurting Mary by pulling her hair. <5.2>

In <5.2> the "hurting" not being an action, has no actor whereas in

John was dragging Mary by pulling her hair. <5.3>

the "dragging", insofar as it involves PTRANS'ing does have John as an actor.

We may wonder by what criterion we draw the line between what an actor does and what he causes. In <5.2>, according to Schank, we are to regard the "hurting" as caused by the "pulling" action. But the same is true of PTRANS'ing in <5.3>. Furthermore, even direct bodily action such as moving an arm can be viewed as caused by muscle contraction or, subjectively, as caused by an act of will, either of which again may have antecedent causes.

With respect to the <actor-action-object> formalism, it seems to us that no structural primitives should be associated with actors at all. Instead we will propose a neutral representation in which events are expressed as sequences of states of the participants. The successive states simply express "what happened", without explicit commitment as to "who did it". However, the agent(s) in an event can be identified by supplementary propositions. Thus the notion of an agent can continue to be used to aid interpretation and inference. However, it would be regarded as a rather "fuzzy" higher level concept, understood by the system in terms of the role of a supposed agent within a sequence of causally and teleologically related states. For example, in the sentence "John uprooted the sapling" the term agent would be considered highly applicable to John's role in the event while in the sentence "The avalanche uprooted the tree" its applicability to the role of the avalanche would be considered relatively low. Two factors determining the degree of applicability of the term "agent" to an object in a sequence of states (event) would be the degree to which autonomous control of its behaviour can be ascribed to the object and the extent to which a state of the object can be held exclusively responsible for initiating the sequence of states.

Similarly we propose to separate why something happened (causes, enabling conditions, reasons, explanations, justifications, and the like) from what happened. As with "agents", this does not prevent us from including causal propositions in the representation and relying heavily on them for interpretation and inference. However, time relations and changes of state, not causes, will give coherence to a set of propositions as an "event".

We feel that Schank's instrumental relation between actions can and should be represented in terms of causation and intention. For example, if a system has a conceptualization to the effect that John was PTRANS'ing the ball by PROPEL'ing it, then this conceptualization should also express that the PROPEL'ing was causing the PTRANS'ing. In fact, phrases ostensibly expressing instrumental actions often express no more than causation. An example is the "by" clause in

The effluents were killing the fish by raising the temperature of the water. <5.4>

When there is a difference, it lies in the intimation of purposive causation. In

John woke Mary by blowing his trumpet. <5.5>

purposive causation is expressed, while in

Mary woke up because John was blowing his trumpet. <5.6>

it is not. Sentences <5.5> and <5.6> clearly show that the instrumental relation amounts to a causal relation supplemented by intentional states.

According to Fodor (1972), actions are to be thought of as a proper subclass of events. Let us determine whether this is the case for Schank's notion of an action, and what might justify the special status of actions as opposed to events. According to Schank, an action is something a nominal can be said to be doing at some moment (this is not a quote, but an interpretation of Schank's definition). A study of his proposed inferences shows that in itself an action does not express a definite change in a situation; rather it expresses existence of a situation which tends to produce change, and all actual changes must be inferred. Formulas for actions in Wilks' theory are analogous although they are not described as explicitly as Schank's primitive action concepts. Actions, then, express modes of behaviour which promote but do not guarantee the occurrence of events. For example, the actions PTRANS, INGEST, MOVE do not express changes in location; instead those changes are primary inferences given that an actor is PTRANS'ing, INGEST'ing, or MOVE'ing something. Syntactically, the relationship between an event, say a change in location, and the action, say PTRANS, whose primary inference is that event, corresponds quite closely to the relationship between verbs and their participles respectively. For example to say that John was PTRANS'ing himself with the result that his location changed is quite analogous to saying that he was going somewhere with the result that he went there.* In any case the term "action" is now seen to be quite misleading, since it normally connotes the occurrence of definite events, rather than the existence of a "dynamic" situation which tends to generate events.

Thus Schank's actions (contrary to the connotation of the term) correspond more closely to states than to events! To say that A is PTRANS'ing B is merely to express a momentary truth about the system in which A and B participate, not a change in that system (which remains to be inferred). This view is compatible with the observation that many common modifiers express subtle blends of "passive" and "dynamic" attributes. The examples below bring to mind conceptual images that illustrate a gradually increasing emphasis on dynamics.

blue sky	burning candle
bright sun	blazing fire
glowing (or luminous) candle	billowing smoke

Schank's actions, and, as far as we can determine, Wilks', are "dynamic states," or "activities," or "modes of behaviour" which mediate changes in certain attributes. Thus PTRANS and MOVE mediate changes in location, INGEST and EXPEL mediate changes in containment relationships, and PTRANS mediates changes in awareness.

We believe that the recognition that "actions" in Schank's sense are essentially states rather than events is important, since it leads to a uniform view of all (true) events as sequences of states. In this view the need for identifying "actors" of events does not arise, nor is it necessary to delineate the spurious boundary between "passive" and "dynamic" states.

We now illustrate our representation of states and events. Nothing new needs to be added to the network notation already used in Figures 1-3. We regard any condition which can hold momentarily (blue, moving, running, etc.) as a state. Accordingly, any atomic proposition which is based on a time-dependent predicate is a state proposition. Figure 4 shows two concurrent state propositions: something (the redness of the sun) was increasing throughout some time interval and something else (the distance between the sun and the horizon) was decreasing throughout the same time interval.




Fig. 4. "The sun was getting redder and approaching the horizon."

Actually there are two additional state propositions, concerned with the existence of unique values of redness and distance at all moments of time within the time interval of interest; these have not been made explicit since they are clearly trivial and would clutter the diagram gratuitously.

Events involve a change in state as "the last leaf fell from the tree" illustrates. The definitive characteristic of state changes is the following: if a system has property A at time t₁, and property B at time t₂, then A → B is a change of state if and only if A and B are mutually exclusive properties.* In fact a state attribute such as colour which can assume various values can consistently be defined as a set of mutually exclusive properties, each member of the set being regarded as a value of the attribute. This admits both qualitative attributes such as colour as well as quantitative attributes such as location. Figure 5 shows a simple event involving a single change of state of a "system" with one component (Mary). The time relation "then" implies immediate succession of the two time intervals. Our representation of one of Schank's standard sentences is shown in Figure 6.




Fig. 5. "Mary fell asleep" Fig. 6. "John hit Mary"

An explanatory paraphrase is the following. "Some unknown mode of behaviour of John caused some object to move quickly toward Mary. Subsequently the object reached Mary and exerted a force on her." Note that we have a state and an event here, viz. John's unknown state and the event of the object moving toward Mary and striking her. In accordance with our earlier remarks about causation, the causal connections between John's state and the ensuing event does not make John's state part of that event. Only exclusive and successive states of a particular system of objects form events. A natural inference in Figure 5 would be that John intentionally hit Mary, i.e. that the missing state of John is that he was trying to bring about the event in question. We would represent "trying" by the state predicate "x has active goal y at time t". An additional example is shown in Figure 7.




Fig. 7. "The cow jumped over the moon."

The explanatory paraphrase goes as follows. The cow was somewhere on the ground, propelling itself towards a location above the moon; then it was moving toward that location; then it was at that location; then it was moving towards a place of destination on the ground, such that the moon is between the place of departure and place of destination; then it was at the place of destination. Note that "moving towards" could have been represented in terms of "distance decreasing" as in Figure 4.

An important consequence of our very broad conception of states is that new complex states (modes of behaviour) can be defined in terms of events involving primitive or already defined states. The time of occurrence of these events can extend some distance backward and forward from the moment at which the new state is defined to hold. For example "walking" is defined in terms of successive states of motion and displacement of the walker's feet and body over a "period of observation" encompassing (say) two steps, since an instantaneous "snapshot" of a person is insufficient for deciding whether or not that person is walking (although it may of course supply enough cues to prompt the inference that the person is walking). A tentative definition of "walking" will be given below.

Complex dynamic states (modes of behaviour) such as walking, running, dancing, tumbling, flickering, etc., can be constructed in terms of more elementary states. The constructions are necessarily as complex as the states they describe. Complexity can result from the intricate coordination of several simultaneous activities (e.g. "rolling" expresses rotation and translation at coordinated rates), or from complex time dependencies (e.g. flickering), or from both (e.g. "walking" or even "building a snowman").

At this point the reader may wonder whether we propose to make any use of "cases" in our approach to representation. The answer is yes in the sense that we intend to exploit fully the semantic "preferences" that any given predicate induces on its arguments. For example the predicate "moving" prefers a physical object as its first argument and a physical location as its second argument; the predicate "has-active-goal" prefers a sentient being as its first argument and a state proposition as its second argument. Furthermore, there are broad similarities between the argument preferences of different predicates. For example, several predicates prefer animate objects in certain argument positions. We certainly can (and sometimes do) acknowledge such similarities and give a rough indication of the sort of preferences involved by using suggestive argument markers such as ANIM, THING, PLACE, DIRECTION, etc., instead of noncommittal markers such as A, B, C, ... However, we do not think that these markers can be chosen so that they express not merely similar but identical argument roles and semantic preferences, no matter in which predicate they occur. This view is supported by Bartsch and Vennemann (1972):

"... 'case' is entirely a surface category and not, as Fillmore (1968) suggests, a category of universal semantics. Semantic representations are based on propositions, which consist of a relation (n-ary predicate with $n \geq 0$) with a finite number of arguments filled either with constants or with bound variables. The 'meaning' of an argument as argument is entirely determined by its relation. Therefore, no two arguments have precisely the same meaning, as arguments. Thus, if the meaning of an argument as argument is called a case, then there are as many cases as there are arguments, and this number, if it is finite at all, is a very large one."

Thus semantic cases, while certainly useful heuristically in finding or inferring arguments of predicates have no universal or primitive status.

VI. COMPLEX CONCEPTS

According to Schank's dictionary, if X walks to Z (where X is human and Z is a location) then X PTRANS's X by X MOVEing the feet of X in the direction of Z. A minor criticism of this formula is that it rules out walking on one's hands and knees, or walking on one's hands (admittedly a rare skill). More importantly, the formula admits running, skipping, hopping, jogging, shuffling, and even skating. Presumably, then, the dictionary entry is not intended to capture the full meaning of "walking" as we seem to understand it, but only those aspects which are most essential to language understanding and (immediate) inference.

Similarly Wilks' formulas are also incomplete.* for example, it is correct to say that DRINK implies ((*ANI SUBJ)((*FLOW STUFF) OBJE)((*ANI IN)((THIS (*ANI (THAU PART))) TO) (BE CAUSE)))) but not the converse (which could mean someone was receiving an enema). So again a selection of linguistically important features has apparently been made.

We feel that it is important to formulate more complete meaning representations for two reasons. First, we believe that somewhat more information will be required for adequate comprehension of "ordinary" discourse. Secondly, much more information will surely be required to match the human ability to describe concepts and reason about them. For example, suppose we ask a reasonably articulate person to describe human "walking" in as much detail as possible. We might elicit at least the following information: Each foot of the walker repeatedly leaves the ground, moves freely in the walking direction for a distance comparable to the length of the walker's legs (while staying close to the ground), then is set down again, and remains in position on the ground, supporting the walker, while the other foot goes through a similar motion. The repetition rate is about one repetition per second. The legs remain more or less extended. The body remains more or less erect and is carried forward at a fairly constant rate.

Further details could be added about flexing actions of feet, knees, and hips, the slight up-and-down motion of the body, typical arm motion, and forces exerted on the ground. Figure 8 shows a network which describes "walking" (regarded as a state predicate with three arguments besides time) along the above lines. A few propositions have been omitted so as not to clutter the diagram. These are that each foot is also above the ground (and close to it) while moving, that each foot is also supporting x while stationary; that the duration of each of the unlabeled time intervals [] is approximately half a second; and that the speed of motion of the walker's body is approximately constant. There is no difficulty in adding these state propositions, except that the last requires "moving" to have an additional argument, namely the speed of motion. Note that [ti] is the "time interval of observation" of the walker, and that it contains t, the time at which x is said to be walking. Thus "walking" is defined by behaviour in the temporal vicinity of the moment of predication, specifically about two seconds of motion allowing about three or four steps.

A limitation of our representation of "walking" is that it is not applicable to unusual modes of walking (e.g. on hands and knees) or to animals. One question this raises is how many "kinds" of walking should be represented separately. Also, is there a representation which expresses the common features of all kinds of walking? We have attempted such a representation in Figure 9. The representation is based on the following characteristics of walking in general:

- (1) it is done using limbs that are a subset of the limbs of the individual involved in the walking;
- (2) the number of limbs involved is greater than or equal to two;
- (3) at all times some of the limbs used for walking are in nonsliding contact with the walking surface (this is not the same as saying some of the limbs are in contact with the surface at all times);

- (4) each limb used for walking is stationary on the walking surface at some time and subsequently is moving for some time; and
 (5) the individual as a whole is in motion in the walking direction.




Fig. 8. "Person x walking at time t in direction d on ground g."




Fig. 9. "x walking at time t in direction d on surface g"

The interesting feature of our representation is the use of quantification to describe the role of any number of legs in the walking. Note that without quantification describing the locomotion of say, a millipede would be very tiresome.

A serious flaw in our representation of "walking", and one for which we have no systematic remedy, is that we have ignored the "fuzziness" of many of the meaning components. For example, it seems necessary to put some constraints on the length of stride (lest a human walker be allowed to move forward in millimeter increments), yet to give an exact distance would be absurd.

An important consequence of conceptual fuzziness, with regard to complex concept definition, is that we can no longer draw a sharp boundary between extracting the meaning of an utterance and making probable inferences on the basis of the derived meaning structure. This is because we only find the probable meaning of an utterance. For example, the utterance "John built the house" probably means that he built a large, rigid-walled enclosure with a roof, separate rooms, etc.; but none of this is certain. The utterance "John was laughing" probably means that he was producing a series of voiced sounds by staggered exhalation of air, and that his facial expression was merry; but he might have laughed silently, or his facial expression might have been derisive or even hostile. If we try to reduce semantic uncertainty by excluding from the "meaning" of a term all but its absolutely minimal content, and ascribe everything else to inference, we find that this doesn't work too well. Translate under c3 "go in the case of "house" all that would remain would be a "partial enclosure" - which accommodates a fenced-off field, a shipping crate, or a jacket.

In the case of "laughing" we would perhaps be left with "spasmodic breathing and intent to convey amusement," which could suggest that John is asthmatic and dancing a jig.

Finally we wish to point out that many (perhaps most) concepts can be understood in different kinds of ways. For example, in "John was listening to the incessant chirping of the crickets," is "chirping" understood simply by its correspondence to a particular auditory sensation, or is it understood as a rapidly fluctuating, more or less uniformly high-pitched sound, or even as a complex variation of air pressure with time? Minsky's (1974) recent work on "frame systems" strongly suggests that the kind of understanding of a concept we use at a given time is extremely task and context dependent. This certainly casts doubt on the one-concept-one-formula approach to language understanding.

VII. ADVERBIALS

Our purpose in this section is to draw attention to the major problem of representing adverbial meanings. In particular we consider Bartsch and Vennemann's (1972) approach, which seems promising but will be seen to have serious defects.

Schank (1972) diagrams adverbs as action modifiers without further analysis. Apparently he has not concerned himself with their meanings so far. In the case of many adverbs (as in the case of many adjectives) this neglect is probably justified, since most of the meaning content derives from perceptual processes. For example, in the sentence "Mary walked gracefully" it is difficult to paraphrase "gracefully" in more elementary terms. Essentially we know gracefulness when we see it. Thus perceptual understanding needs to be supplemented only by a few additional facts for language comprehension purposes, such as the fact that graceful motion is generally pleasing, is more or less the opposite of awkward motion, is smooth and well-coordinated, and the like. Other adverbial modifiers, however, clearly require systematic analysis; "quickly" is a good example. This term appears to say something about the speed of an action or activity, comparing it to some standard. An adequate meaning representation for "quickly" should spell this out precisely.

Bartsch and Vennemann (1972) suggest that adverbial modifiers operate on verb meanings in the same manner that adjectival modifiers operate on noun meanings. One problem in this approach is best illustrated by the following example.

John owns a large car. <7.1>
 John is running quickly. <7.2>

Whereas large in <7.1> has as a reference set the set of cars, and John's car is large in relation to the "average" for that set, "running quickly" cannot be analyzed so easily. If the analogy were perfect then the reference set operated on by "quickly" would be the set of "runnings" (whatever that means); but clearly this set of runnings must be further restricted to the set of runnings John is capable of performing. Thus "quickly" appears to operate not on "running" alone, but on "John running". As further examples consider <7.3> and <7.4>.

The cheetah is running quickly. <7.3>

The ant is running quickly. <7.4>

Clearly "quickly" here operates on "running ant" and "running cheetah" respectively.

Thus the nature of the "runner" is being used to narrow the reference set to which we apply a measure function. In <7.2> "quickly" modifies running with respect to John's runnings, or, if we don't know John, at least to human runnings (assuming that John is human). In <7.3> and <7.4> the measure function is applied to the runnings of cheetahs and ants respectively. Unfortunately factors other than the identity or category of the runner can also affect the meaning of "quickly", as shown by <7.5> to <7.9>.

John is running quickly on his hands and knees. <7.5>

John is running quickly on the moon. <7.6>


John is running quickly in Chile. <7.7>

The cheetah is running quickly in the dense forest. <7.8>

The cheetah is running quickly on the plain. <7.9>

The effect of locale on the meaning of "quickly" is seen in the contrast between <7.6> and <7.7> and between <7.8> and <7.9>.

Thus it appears to us that the context which determines the meaning of an adverbial modifier cannot be circumscribed once and for all. In general, adverbials must be allowed to interact with any specific and general knowledge available about the participants in (and setting of) an action. Thus the approach of Zadeh (1972) to the treatment of adverbial "hedges" lacks generality since he needs to specify the (weighted) components of each fuzzy term on which a hedge may operate once and for all. In our semantic network, we would represent <7.3> without the adverb as follows.



- (a) $(\exists x) \{ (\forall y) [\text{cheetah}(y) \ \& \ ?(y)] \Rightarrow x=y \} \ \& \ \text{running}(x)$
 (b) $(\exists x) [\text{running}(x) \ \& \ \text{cheetah}(x) \ \& \ ?(x)] \ \& \ (\forall y) [\text{cheetah}(y) \ \& \ ?(y) \Rightarrow x=y]$

Fig. 10. "The cheetah is running."

(a) and (b) are based on alternative (but equivalent) representations of definite descriptions. Figure 11 then shows an attempt to represent the adverbial construction in <7.8> in keeping with Bartsch and Vennemann's (1972) general approach but taking into account the above considerations. In the representation we show the explicit relationship between the speed of the

cheetah's running as compared to the average speed of a reference set of runnings (by cheetahs in dense forests).




Fig. 11. "The cheetah is running quickly in the dense forest"

Unfortunately the representation suffers from a fatal flaw: the set of cheetahs running in dense forests, required for comparison, may well be empty (if not, replace "dense forest" with "deep snow"). The "reference set" therefore, if it exists at all, is not of this world but of some imaginary world which is our conception of how hard cheetahs would find the going if they were to run through forests (or snow). It remains to be seen whether Bartsch and Vennemann's approach can be modified to overcome these difficulties.

VIII. SUMMARY

We have taken a critical look at what we consider to be the best approaches to a complete language understanding theory to date, namely Schank's conceptual dependency theory and Wilks' preference semantics. In doing so, we exposed many problems and suggested plausible methods for dealing with some of them. In particular, the representation of quantification and of time, the basic representation in terms of states and events, the definition of complex concepts (most importantly action concepts), the handling of adverbial modifiers, as well as the use of time in representations has been examined.

In a companion paper by Schubert (1974), notational difficulties are eliminated for such necessary constructions as logical connectives, descriptions, quantifiers, propositional attitudes and other modalities, as well as other constructions not previously incorporated into semantic nets. This notation appears to be adequate for the "state based" conceptual representation here proposed.

Obviously much work remains to be done. We hope our exposition and partial solution of several major problems will help give direction to this work.

REFERENCES

- Bartsch, R., and Vennemann, T. (1972). *Semantic Structures*, Athenäum Verlag, Frankfurt, Germany.
 Evans, C. (1967). "States, Activities, and Performances", *Australasian Journal of Philosophy*, 45, pp 293-308.
 Fodor, J. (1972). "Troubles About Actions", in *Semantics of Natural Language*, D. Davidson and G. Harman (eds), D. Reidel Publishing Company, Boston, Massachusetts.
 Katz, J. and Fodor, J. (1964). "The Structure of Semantic Theory", in *The Structure of Language*, J. Fodor and J. Katz (eds), Prentice-Hall, Englewood Cliffs, New Jersey.
 Lakoff, G. (1972). "Linguistics and Natural Logic", in *Semantics of Natural Language*, D. Davidson and G. Harman (eds), D. Reidel Publishing Company, Boston, Massachusetts.
 McCarthy, J. and Hayes, P. (1969). "Some Philosophical Problems from the Standpoint of Artificial Intelligence", in *Machine Intelligence*, B. Meltzer and D. Michie (eds), American Elsevier Publishing Company, New York, New York.
 Minsky, M. (1974). "A Framework for Representing Knowledge", *Artificial Intelligence Memo No. 306*, A. I. Laboratory, MIT, Cambridge, Massachusetts.
 Montague, R. (1970). "The Proper Treatment of Quantification in Ordinary English", *Department of Philosophy, University of California, Los Angeles*.

- California.
- Montague, R. (1972). "Pragmatics and Intensional Logic", in Semantics of Natural Language, D. Davidson and G. Harman (eds), D. Reidel Publishing Company, Boston, Massachusetts.
- Quine, W., (1960). Word and Object, MIT Press, Cambridge, Massachusetts.
- Rosser, J., and Turquette, A. (1952). Many-valued Logics, Amsterdam, North Holland.
- Schank, R. (1972). "Conceptual Dependency: A Theory of Natural Language Understanding", *Cognitive Psychology*, v 3, pp 552-631.
- Schank, R. (1973). "Identification of Conceptualizations Underlying Natural Language", in Computer Models of Thought and Language, R. Schank and K. Colby (eds), W. H. Freeman and Company, San Francisco, California.
- Schank, R., Goldman, N., Rieger, C., and Riesbeck, C. (1973). "Margie: Memory, Analysis, Response Generation, and Inference on English", Third International Joint Conference on Artificial Intelligence, SRI, Menlo Park, California, pp 255-261.
- Schubert, L. (1974). "On the Expressive Adequacy of Semantic Networks", TR74-18, Department of Computing Science, University of Alberta, Edmonton, Alberta.
- Wilks, Y. (1973a). "Preference Semantics", Stanford AI Project, Memo AIM-206, Stanford University, Stanford, California.
- Wilks, Y. (1973b). "Natural Language Inference", Stanford AI Project, Memo AIM-211, Stanford University, Stanford, California.
- Wilks, Y. (1973c). "Understanding Without Proofs", Third International Joint Conference on Artificial Intelligence, SRI, Menlo Park, California, pp 270-277.
- Winston, P. (1970). "Learning Structural Descriptions from Examples", PhD Thesis, MIT, MAC-TR-76, Cambridge, Massachusetts.
- Zadeh, L. (1965). "Fuzzy Sets", *Information and Control*, 8, pp 338-353.
- Zadeh, L. (1972). "A Fuzzy-Set-Theoretic Interpretation of Linguistic Hedges", *Journal of Cybernetics*, v 2, n 3, pp 4-34.

FOOTNOTES

- ¹ Cf. the selectional restrictions of Katz (1964).
- ² For a more complete description see below.
- ³ This something includes natural forces in a vague unspecified sense.
- ⁴ Unlike Schank we do not regard "he was going" and "he went" as equivalent; we claim that "he went there," unlike "he was going there," affirms that he did arrive at his destination, and that it is decidedly odd to say "he went there but didn't get there."
- ⁵ For example, A=solid, B=liquid; A=round, B=rectangular.
- ⁶ Cf. Lakoff's (1972) lexical decomposition trees.
- ⁷ It would seem that the human interpretative process does not proceed on the basis of the minimal conceptualization formed by embedding the minimal content of terms into the conceptualization. Rather the original term itself suggests what we could infer in addition to the minimal content. This seems like a reasonable view on the basis of efficiency considerations as well. It should be much simpler to insert probable inferences in a semantic structure by direct reference to the word definitions instead of analyzing the minimal representation and then looking for applicable inference rules.
- ⁸ An opaque context is one which does not allow substitution of referentially equivalent expressions or does not allow existential quantification. This is a well known problem discussed at length by Quine (1960) and in Artificial Intelligence by McCarthy and Hayes (1969) and Moore (1973).
- ⁹ See Partee (1972) for details.