

D. L. Waltz

University of Illinois  
Urbana IL

By choosing to analyze examples of language which are unconventional and cannot always be grasped instantly by people, I believe that otherwise hidden processes active in language understanding can be exposed for examination. This paper addresses specifically the problems of how much knowledge is attached to a word, how much meaning is contained in syntactic structures, and ultimately what kinds of processing mechanisms are necessary to understand things which may be said in English.

I would like to make several main points: first, that while processing is generally organized around verbs, the target structure, or structure that a speaker wishes to transmit to the hearer, may very well not be organized around verbs; second, that the part of speech of a word and the syntactic structure of a sentence have little in general to do with their meanings; third, that words may in general have large framelike structures attached to them which guide processing; and fourth, that language understanding, especially of novel language like poetry and metaphors, often involves the forcing of agreement between two or more frames to form a target structure. Finally I will say a few words about relevant current AI programs and about how easily they might be extended to handle such linguistic phenomena.

Unfortunately, this paper points out many more problems than it solves, and does not offer any more than a sketch of necessary processing mechanisms.

### Target Structures

Case grammar theories as developed especially by Fillmore (1968) and Celce-Murcia (1972) have argued convincingly that sentence processing is centered around the verb. However the meaning which a speaker wishes to transmit to a hearer may in fact not be centered around verbs at all. As Minsky (1974) writes:

In more extended discourse...I think that verb-centered structures often become subordinate or even disappear. The "topic" or theme of a paragraph is as likely to be a scene as to be an action, as likely to be a characterization of a person as to be something he is doing. Thus in understanding discourse, the synthesis of verb-structure with its case-arguments may be a necessary but transient phase. As sentences

-----  
\*This work was supported by the Office of Naval Research under Contract N00014-67-A-0305-0026, at the Coordinated Science Laboratory, University of Illinois, Urbana.

are understood, the resulting substructures must be transferred to a growing "scene-frame" to build up the larger picture.

By "target structure" I will refer roughly to what Minsky has called here "the larger picture."

However it may not always be the case that large target structures are built up incrementally in a discourse. Sometimes a large target structure may be created by a single sentence, as in

- 1) The girl's room seemed to the priest like the inner circle of Dante's Inferno.

Here in a few words one can evoke a picture of a dark, foul-smelling, hot and oppressive place, and the target structure is something like a list of such characteristics assigned to "the girl's room." This target structure bears little resemblance to any "deep structure" I can imagine obtaining by merely parsing this sentence.

### The Meaning of Syntactic Structures

Even in much simpler sentences, the target structure may not bear much relation to the "deep structures" of the sentences. This is most clearly demonstrated in the case of pictured descriptions of the world, where the parts of speech of words are not an accurate guide to their meanings. For example, why should explode and burn be verbs, while noise and lightning are nouns? Whorf (1940) discusses such examples at length, and argues persuasively that our use of nouns and verbs to describe the world does not reflect the inherent structure of the world, but rather reflects our linguistic training.

Whorf goes even further, to argue that we tend to "see" lightning as a thing because it is a noun, but burning as an action because it is a verb. Thus in this view, the words we have to describe the world warp our perception of the world. If this is true, then it may be that our perception is largely symbolic, as Minsky (1974) suggests, and syntactic structures may be meaningful in the sense that frames would be initially composed of syntactic structure fragments. The fact is that English is rather inconsistent in the types of words and structures it uses to describe phenomena, and therefore inconsistent in the meanings of syntactic relations. This fact may be intimately related to the fact that it is difficult to choose logically adequate arcs and links for representing relations between nodes of semantic nets, as demonstrated convincingly in Woods (1975). In any case, it seems to me that one should be wary about assigning much meaning to a syntactic structure.

## How Much Meaning is Attached to a Word?

The amount of structure one must postulate to be attached to a word varies tremendously depending on the particular word and the context of its use. Here I will try to show certain simple cases in which quite large amounts of knowledge must be attached to words. Take for example the following sentences:

- (2) John is a physician.
- (3) John is a Renaissance man.

In the first sentence (2), it seems that parsing plus simple semantic mechanisms can extract much of the meaning, if one assumes that "x is a y" means "x is a member of the set of all y." A system embodying these mechanisms (like Raphael (1968)) could properly set up "member" and "superset" links between John and physician. I will argue later that this analysis of meaning is insufficient, but for now let us assume that it at least captures much of the intended meaning of this sentence.

However, such an analysis will clearly not suffice for sentence (3). Why not? Partially because membership in the set of "Renaissance men" is not neatly delineated as is membership in the set of physicians; it would seem exceedingly strange to set up a single node in a semantic net labeled "Renaissance men." Another reason is that the sentence seems to really be a statement of the speaker's opinion, and as such a hearer would need to have separate nodes for "Harriet's list of Renaissance men," "Lyndon's list of Renaissance men" and so on for each speaker who has so classified a person to the hearer. Of course this too seems ludicrous. The underlying meaning conveyed by the sentence is not purely a matter of opinion on the part of the speaker, but something like "John is knowledgeable and accomplished in a number of different fields," a matter (potentially) testable by the hearer. That this is the true import of the sentence can be demonstrated by the fact that

- (4) John is a Renaissance man, but he has no knowledge about AI.

provides a reasonable contrast of propositions, whereas

- (5) John is a Renaissance man, but Harriet doesn't think so.

does not seem to be a reasonable contrast.

Sentence (4) suggests something else of interest. After

- (3) John is a Renaissance man.

we expect something like

- (6) He plays oboe in the Boston symphony, won the Nobel Prize in physics, has written 30 best-selling novels, and won the 1972 U.S. Demolition Derby Championship.

Sentence (6) supports the original assertion and gives a particular instantiation of the general concept of "Renaissance man" appropriate for John. I believe that sentence (3) sets up something like the skeleton of a frame for "Renaissance man" as its target structure, leading one to expect further explanation or expansion to fill in the empty slots of the frame.

Let us now reconsider sentence (2). When one says that "John is a physician," does one merely set up in the hearer links between two nodes named John and physician? Consider the following sentences:

- (7) John is a physician, but he does not know much about medication.
- (8) John is a physician, but he does not mind taking orders.

I believe that these sentences illustrate the fact that in understanding "John is a physician" one is not merely setting up a link between two nodes but is in fact assigning to John a number of default attributes of physicians, i.e., the hearer takes the word physician to refer simultaneously to an entire complex of attributes, knowledge, actions and attitudes. Thus (7) provides a reasonable contrast, since one would expect a physician to know about medication. A listener might infer from (8) that the speaker presupposed that all physicians resent taking orders. But more than presupposition of the speaker is at work here; the sentence

- (9) John is a physician, but he resents taking orders.

seems a ludicrous contrast, almost like "He may be fat but he runs slowly." One thing we know about doctors is that they ordinarily do not take orders, so it is quite plausible that they might not like to take orders. The point here is that even such unlikely inferences as (8) seem to be included automatically in the sentence "John is a physician." That there is a limit to such inferences is illustrated by the fact that

- (10) John is a physician, but he has a low bowling average.
- and

- (11) John is a physician, but he has a high bowling average.

both seem equally to be non-sequiturs.

## Understanding Poetry

The types of language which are most difficult for us to understand may best illustrate the amounts of knowledge that must be attached to words and the methods for combining this knowledge or transferring complexes of knowledge from one domain to another (as in metaphor). Everyday language communicates familiar relations, events and attitudes, and does so in familiar ways. In the terms of this paper, the target structures of such communication already exist as frames which need only instantiations and exceptions to become

complete representations. As such, the full meanings of words rarely become evident. Let us consider some examples which require overhaul or forcing together of existing structures.

An accident is an inevitable occurrence due to the action of immutable natural laws.

-Ambrose Bierce, The Devil's Dictionary

What makes this funny is that it points out a conflict between essential facts in our frames which explain personal actions, and our frames which explain physical events. One cannot simultaneously hold the view that accidents are the result of random chance and the view that "God does not throw dice," yet we do. A system which is not aware of the humor in this unresolvable conflict cannot be said to understand the passage.

Inebriate of air am I,  
And debauchee of dew,  
Reeling, through endless summer days,  
From inns of molten blue.

-Emily Dickinson,  
from "I taste a liquor never brewed"

In order to understand this passage one must force together frames for drinking immoderately invoked by inebriate, debauchee, reeling and inns and for experiencing nature, invoked by air, dew, endless summer days, and molten blue (skies). (See Charniak (1972) for a discussion of the difficulty of invoking a frame when it is not mentioned explicitly by name.) In this case note the difficulty of saying something similar in a sentence: "I get drunk from experiencing nature" does not seem to produce anything like the same target structure for me, probably because I am not forced to interweave the two frames as I am in the poem.

A great deal of communication, not only in poetry, depends on the use of analogies akin to the example above. Ortony (1975) has postulated three main functions which metaphors (as well as similes and analogies) serve in language; first, they allow one to transfer large amounts of information from one domain to another without explicitly spelling out all the details; second, they allow one to communicate the otherwise inexplicable; and third, they often make distinctions more vivid by adding otherwise missing perceptual or emotional content. His examples further support the general thesis of large amounts of knowledge being potentially attached to each word. For instance.

(12) The thought slipped my mind like a squirrel behind a tree.

enables one to express the fact that thoughts are sometimes elusive, quick, deceptively easy to catch, camouflaged, etc. and to say so in a very compact fashion. Such an idea cannot be stated directly since

we have no literal language for the behavior of thought; to say that "the thought escaped me" is still to rely on a metaphor.

Finally, analogies are useful for creating or teaching new frames. While the account below is invented, I have had a number of such exchanges with my daughter.

Adult: A veterinarian is a doctor for animals.

Child: Does he give the animals lollipops? Does he use a stethoscope? etc.

Clearly what is happening in such an exchange is that the child is creating a new frame for veterinarian, and is testing to see just how much can be transferred from her doctor frame. Questions often center on things that are difficult to imagine transferring, e.g., an animal eating a lollipop.

#### Relation to Current AI

A number of pieces of research, the work on conceptual dependency diagrams of Schank et al (1973), Schank (1973), the work on belief systems of Colby (1973), Abelson (1973) and McDermott (1974) as well as the already mentioned work of Charniak (1972) in understanding children's stories and Minsky's (1974) work on frames, stand out as most relevant to solving the kinds of problems I have presented. In addition, Collins et al (1974) have investigated the problem of teaching geographical concepts by use of analogies. I will, however, discuss only Schank's work in more detail here.

In Schank's schemes, verbs can have quite complex structures attached to them, which are not necessarily closely related to the syntactic deep structure of the sentences in which they occur, but which are related to what I have called target structures. For example, Schank represents the sentence

(13) John grew the corn with fertilizer.

as a conceptual dependency diagram which I read as something like

(14) John took fertilizer out of a bag and put it on the corn's ground, and this caused the corn to become larger.

Clearly the structure attached to the transitive verb grow constitutes a frame of sorts, which could potentially be used to transfer meaning to a new domain. It is not too difficult to imagine a program which could accept a sentence like

(15) Raising cattle is like growing corn, except that instead of spreading fertilizer, one feeds the cattle grain.

and use its knowledge of grow to produce a new frame for raise (cattle). Schank's system can also draw inferences from sentences, e.g., "The corn will become ripe," "John probably wants to eat the corn," etc.

from sentence (13).

Some key problems still have not been addressed, however. First, the inference process does not know when to stop generating inferences, whereas I believe people have quite stable and definite ideas about how far inferences should be drawn from a sentence. Second, as illustrated in a number of earlier examples, nouns (e.g., physician) may have substantial amounts of structure attached to them. In Schank's work, only verbs are represented by substantial structures. Third, while Schank has done work on piecing together sentences which form a coherent and expected discourse (a process he calls knitting), he has not to my knowledge worked on attempting to fit together apparently unrelated follow-up sentences. Indeed, without much more detailed frames for words, I believe that there is no way to understand the more tenuous connections between ideas that we discover routinely every day.

#### Acknowledgement

I would like to thank Andrew Ortony for suggesting a number of the ideas in this paper to me.

#### REFERENCES

- Abelson, R.P., "The Structure of Belief Systems," in Schank and Colby (eds.), Computer Models of Thought and Language, W.H. Freeman, San Francisco, 1973.
- Celce-Murcia, M., "Paradigms for Sentence Recognition," UCLA Dept. of Linguistics, 1972.
- Charniak, E., "Toward a Model of Children's Story Comprehension," MIT AI Lab AI-TR-266, 1972.
- Colby, K.M., "Simulations of Belief Systems," in Schank and Colby (eds.), Computer Models of Thought and Language, W.H. Freeman, San Francisco, 1973.
- Collins, A., Warnock, E.H., Aiello, N. and Miller, M.L., "Reasoning from Incomplete Knowledge," unpublished draft, BBN, Cambridge, MA, June 1974.
- Fillmore, C., "The Case for Case" in Bach and Harms (eds.), Universals in Linguistic Theory, Holt, Rinehart and Winston, 1968.
- McDermott, D.V., "Assimilation of New Information by a Natural Language-Understanding System," MIT AI Lab AI-TR-291, 1974.
- Minsky, M.L., "A Framework for Representing Knowledge," MIT AI Lab Memo no. 306, 1974.
- Ortony, A., "Why Metaphors are Necessary and Not Just Nice," in Educational Theory, Vol. 25, No. 1, Winter 1975.
- Raphael, B., "SIR: a Program for Semantic Information Retrieval," in Minsky (ed.) Semantic Information Processing, MIT Press, Cambridge, MA, 1968.
- Schank, R.C., Goldman, N., Rieger, C.J., and Riesbeck, C., "MARGIE: Memory, Analysis, Response Generation, and Inference on English," Adv. Papers of the 3IJCAI, Stanford Univ., 1973.
- Schank, R.C., "Identificaton of Conceptualizations Underlying Natural Language," in Schank and Colby (ed.), Computer Models of Thought and Language, W.H. Freeman, San Francisco, 1973.
- Whorf, B.L., "Science and Linguistics (1940)" in Language, Thought and Reality, MIT Press, Cambridge, MA, 1964.
- Woods, W.A., "What's in a Link: Foundations for Semantic Networks," in Bobrow and Collins, Representation and Understanding, New York: Academic Press, 1975.