

CHARLES E. M. DUNLOP

CONCEPTUAL DEPENDENCY AS THE LANGUAGE OF THOUGHT¹

The rule of the game, therefore, is not for the reader to say 'You can't do that', because what we describe can be and has been done, to varying degrees of success. Rather, you may say, 'That isn't quite right', or 'You've oversimplified a very deep philosophical problem' . . .

(Rieger 1975, p. 195)

AI must come to terms with the fact it is concerned with many issues that are also of interest to philosophers. I hope that the cooperation here will be of more use than was the head-butting that has gone on between AI people and linguists.

(Schank 1980, p. 178)

ABSTRACT. Roger Schank's research in AI takes seriously the ideas that understanding natural language involves mapping its expressions into an internal representation scheme and that these internal representations have a syntax appropriate for computational operations. It therefore falls within the computational approach to the study of mind. This paper discusses certain aspects of Schank's approach in order to assess its potential adequacy as a (partial) model of cognition. This version of the Language of Thought hypothesis encounters some of the same difficulties that arise for Fodor's account.

Two influential and much-discussed themes in the philosophy of mind/cognitive science literature are (1) that understanding natural language involves mapping its expressions onto some kind of internal representation scheme, and (2) that the resulting representations exhibit a syntax such that they are susceptible to computational operations. These ideas have been taken seriously in Roger Schank's artificial intelligence projects at Yale University, where a variety of computer programs have been based upon them. Schank's ultimate research goal is not only to build AI systems that exhibit human-like linguistic behavior, but also to provide a psychologically accurate computer model of various human mental processes involved in the understanding of language. His work, therefore, falls squarely within the 'computational' approach to the study of mind.

Synthese 82: 275–296, 1990.

© 1990 Kluwer Academic Publishers. Printed in the Netherlands.

Although computational theories of mental phenomena are committed to the apparatus of internal representations, they are often quite vague as to the *nature* of those representations. Schank's work has the merit of describing the hypothesized representation scheme in considerable detail. Since he regards it as providing a psychologically accurate necessary condition for the use of natural language, his theory may be viewed as one version of the Language of Thought hypothesis (which Schank calls 'Conceptual Dependency' theory, abbreviated as CD).

This paper discusses some aspects of Schank's representation scheme, with an eye toward assessing its adequacy as a (partial) model of cognition.² I shall begin by outlining various features of the theory, and at points later on I shall try to connect it with recent philosophical work by Jerry Fodor and others. Schank's own account, I shall argue, contains many of the same pitfalls that may be found in the philosophical discussions of the Language of Thought. Some of these pitfalls will be familiar to philosophers with an interest in the computational approach to mind; others, I believe, have received little if any treatment in the literature.

1. A PRELIMINARY ARGUMENT

Considerations of parsimony, Schank believes, argue for some sort of universal, internal representation scheme in the processing of sentences (Schank 1975, p. 28; Schank and Riesbeck 1981, pp. 14–16). To illustrate, suppose that the representation scheme was not universal, but rather language-specific. In that case, a speaker of English would map sentences onto English-specific representation structures, while a speaker of French would do the same vis-à-vis French-specific structures. Let E stand for some English sentence, and F for its French equivalent, while $R(E)$ stands for the internal representation of E , and $R(F)$ for its French counterpart. Then, a translation of E to F would require (i) mapping E onto $R(E)$; (ii) correlating $R(E)$ with $R(F)$; and (iii) mapping $R(F)$ onto F . Assuming that $R(E)$ is not identical to $R(F)$, an increasing number of representations along with language-specific correlation rules of type (ii) would be required as more languages entered the picture. Matters would be considerably simpler, however, if one universal representation scheme underlay a variety of languages.

In that case, translation of an English sentence into its French equivalent would require only mapping *E* onto *R*, and then *R* onto *F*.

This appeal of an interlingual representation scheme, incidentally, was recently noted in a *Time Magazine* account of Japanese language-translating machines. The Fujitsu Company has developed devices that provide at least rough translations between English, Japanese, French, and German (additional languages are planned as well). As a researcher from Fujitsu put it: 'If we did not use interlingua, then each pair of languages would require the development of a specific set of grammatical rules and a bilingual dictionary. Interlingua acts as the hub of a wheel'.³

Of course, the utility of an interlingua for machine translation does not prove that monolingual speakers employ a language of thought, although it does explain why Schank, whose projects include translation, might find it appealing. Further arguments will be considered later on. Note also that the argument just canvassed on behalf of an interlingual representation scheme is silent as to the *nature* of that scheme. But Schank goes on to say a good deal about that topic. His account is distinctive in its commitment to the primacy of meaning over syntax; Conceptual Dependency is essentially a meaning representation scheme. Thus, Schank's language-processing systems do not attempt to build separate syntactic representations of natural language sentences; in fact, syntactic considerations come into play only when required in order to help resolve ambiguities, find linguistic units that have been predicted by semantic features, etc.

2. CONCEPTUAL DEPENDENCY THEORY: VOCABULARY

Schank's work on natural language understanding has focused primarily on the representation of actions – broadly construed so as to include natural forces as 'agents'. Thus, 'John hit Mary' and 'Hurricane Gilbert hit Mexico' fall under the purview of actions, so conceived. With an eye toward obvious objections to this view, Schank emphasizes that his aim is to capture how ordinary speakers conceptualize the world, irrespective of whether such an account will withstand close ontological scrutiny (Schank 1973a, p. 206; Schank 1975, p. 41; cf. Rieger 1975, p. 187). It remains to be seen, of course, whether ordinary speakers do in fact view the world in the way that Schank's account dictates. If

they do, and their view harbors some incoherence, then a computer model of their view will naturally exhibit the same incoherence.

The core of Schank's internal representation scheme, Conceptual Dependency, involves an ACT, which is an action performed on some OBJECT. The actor is known as a PICTURE PRODUCER (PP). ACTs are directed toward a LOCATION, which indicates their DIRECTION, and may result in an OBJECT's being in a particular STATE, or a RECIPIENT's coming to possess an OBJECT. From these ingredients, a conceptualization may be formed, i.e., a representation structure which indicates what was done by whom, to what, etc.

These ingredients of CD theory sound suspiciously close to categories of natural language, but Schank takes pains to insist that the conceptual level is *extralinguistic*:

We have required that the meaning representation that we use be language-free [W]e began to believe that language and thought were separable structures (Schank 1975, p. 7).

What does this mean? In its most benign sense, the point would appear to be that internal representations are not identical to natural language sentences. But while this may be true under a narrow interpretation, Conceptual Dependency vocabulary bears a striking relationship to familiar terms in natural language. Consider, for instance, the eleven⁴ primitive ACTs in CD theory (Schank 1975, pp. 40-44; Schank 1981, pp. 17-25):

Physical ACTS

PROPEL	Apply a force to
MOVE	Move a body part
INGEST	Take something to the inside of an animate object
EXPEL	Take something from inside an animate object, and force it out
GRASP	To physically grasp an object

ACTS That Cause State Changes

PTRANS	To change the location of something
ATRANS	To change an abstract relationship


ACTS Used Primarily as Instruments of other ACTS

- SPEAK To produce a sound
- ATTEND To direct a sense organ toward a particular stimulus

Mental ACTS

- MTRANS To transfer information
- MBUILD To create or combine thoughts

Schank's assumption is that the meaning of action-sentences can be captured by way of these eleven primitives. Consider, for example, the sentence 'John ate a frog'. Its CD analysis looks like this (Schank 1975, p. 24):



JOHN is the PP; INGEST is the ACT, and the symbol ' \Leftrightarrow ' marks a mutual dependency relation. The OBJECT of the ACT is designated by 'frog'; 'D' indicates the ACT's DIRECTION; and the INSTRUMENT of John's ACT is yet another ACT in which JOHN MOVED his hand.

In what sense is this representation scheme 'extralinguistic'? To be sure MTRANS and MBUILD are not (so far as I know) terms of any natural language, and MOVE is defined more restrictively than its English language counterpart, although the meaning of INGEST offers no surprise. JOHN, according to Schank, is not the English name 'John', but rather a pointer to all the information that we have about John (apparently, a reduction of objects to bundles of properties is ultimately envisaged⁵). But if the categories just outlined actually represent the categories in which we think, there is little evidence that it is 'language-free'; indeed, it seems perverse to maintain that category


terms that can be *defined* are extralinguistic. Moreover, Schank offers no argument for such a requirement. The success of his computational models does not derive from that assumption, and other considerations (philosophical, psychological, linguistic) may point in the opposite direction. Jerry Fodor once suggested that 'the language of thought may be very like a natural language' (Fodor 1975, p. 156).

The MTRANS primitive, denoting transfer of information, deserves special notice. Information conveyed in an MTRANS is always a complete conceptualization (corresponding to a proposition), with MTRANS taking the recipient case. Information transfer here is thought of in terms of a proximate source and a receiver. The eye, for example, may be regarded as the proximate information source. What is the 'receiver' here? Schank answers that information first goes to a Conceptual Processor, where all conscious thoughts occur. This CP, then, is viewed as the recipient. From the Conceptual Processor, information may go to an Intermediate Memory (as when we remember a telephone number just long enough to dial it), or to Long Term Memory (in cases where we need the telephone number on a later occasion). These sorts of mental categories have been argued for by various cognitive psychologists, and some interesting experiments have been cited in support of them. But I want to focus here for a moment on an odd restriction that Schank imposes on Long Term Memory, or LTM. He writes that it

contains all the information that is known by a person. We postulate that only true facts are stored in LTM and that false things are derived from them. (Schank 1975, p. 44)

Schank's 'postulate' (that only truths are stored in Long Term Memory), besides being unnecessary, has little to recommend it. After all, whether or not something is true often cannot be perceptually discriminated by an observer, and there is no reason to suppose that false beliefs arising in such circumstances cannot get into Long Term memory.⁶ To clarify, suppose that I read in an encyclopedia that Abraham Lincoln was the sixteenth President of the United States, and suppose also that someone else (owing to a misprint) reads in a different encyclopedia that Lincoln was the fifteenth President. It is scarcely plausible to suppose that the information *I* got goes into my LTM, while the other person's information does not get represented in LTM; our relationships to our respective sources of information were virtually identical. Well, perhaps Schank means that what goes into each of our LTMs is

not the encyclopedia's information *simpliciter*, but rather the information *that the encyclopedia reported the information*. (This might provide the basis for Schank's claim that falsehoods are *derived from* our stock of stored truths.) There may, of course, be situations where this occurs, but no plausible psychological theory can insist that all information stored in LTM is of this kind; we frequently fail to remember where we got a particular piece of information. In such cases it is the *information*, not the *source + report*, that we normally retain in long term memory. There is evidence that Schank agrees, for he represents the sentence 'John read about Nixon in the Encyclopedia' as follow (Schank 1975, p. 61):



Schank adds that 'What John MTRANSed here were assorted unknown facts about Nixon'; he does not say that John MTRANSed the fact that the encyclopedia was reporting facts about Nixon. Since the putative facts may be incorrect, but may nonetheless be accepted by John, the 'postulate' that Long Term Memory contains only truths is untenable. In terms of our example, we should say also that at least some of what goes into John's Long Term Memory is not knowledge.

Why, then, does Schank endorse the curious doctrine that Long Term Memory contains only truths? The primary reason, I suggest, is that he takes himself to be proposing a knowledge representation scheme. In Schank's words, "'Know' is represented as 'being in the LTM of'" (Schank 1975, p. 62). Moreover, I suggest, Schank realizes that knowledge requires truth. Thus, in order to accommodate these two ideas – (1) that a knowledge representation scheme is being proposed, and (2) that knowledge requires truth – it is stipulated that the system's Long Term Memory only contains truths. From here it is not

too difficult to conclude (as Schank in fact does conclude) that knowledge and belief are 'virtually identical'.

Although this treatment of knowledge representation is quite common in the artificial intelligence literature, it fares rather badly on two grounds. First, by requiring that all stored propositions be true, it ignores the question of how *mis*representation is possible, although this is a question that models of human cognition can scarcely afford to dismiss. And second, it fails to recognize that knowledge representation, properly conceived, involves considerably more than just the well-organized storage of true propositions. To determine just what else is needed comes as no easy task, as post-Gettier epistemology has made clear. But knowledge representation systems that ignore this question are, I submit, parading under a false banner.

Earlier, it was mentioned that CD theory provides a computational advantage for translation from one language to another. The primitive ACTs of CD theory are also interesting insofar as they allow for a common representation of sentences that have significant overlapping conceptual content. This too has a computational advantage, insofar as allowable inferences associated with ACT primitives only need to be stored once in the system. For example, (1) 'John gave Mary a book', (2) 'Mary received a book from John', (3) 'John bought Mary a book' and (4) 'John stole a book for Mary' would all be diagrammed using the ATRANS primitive (for abstract transfer of possession). This ATRANS primitive serves as the common connecting point for probable inferences concerning transfer of possession (Who has the object now? Who had it before? Did John want Mary to have a book?), thereby obviating the need to store the same inference rules redundantly. Of course, the four sentences in this example have important differences of meaning also. Such differences would tend to be brought out by different Instrumental case diagrams (indicating the means by which the ACT was accomplished).

Despite the inferential (computational) advantage accruing to ACT primitives, there are times when this approach does not work very well. For example, one of the inferences from EXPEL is that the EXPelled object was previously INGESTed. But the sentence 'John spat at Mary' is represented in terms of John's EXPelling saliva (Schank 1975, p. 58), although in most instances saliva is manufactured, not INGESTed. It is hardly surprising that a mere eleven ACT primitives should fail to capture the core meaning of most natural-language action verbs. The

problem becomes particularly acute, however, for action terms which operate against a background of social institutions and conventions. As Schank acknowledges, 'if we have 'John kissed Mary', our mapping of kiss into 'MOVE lips towards' will not simplify the problem one bit' (Schank 1975, p. 81). Similar points could be made about 'The policeman gave me a parking ticket' and 'My friend has been married and divorced three times'. In all such cases, ACT primitives appear to be of little service; what is needed in these particular cases is knowledge of appropriate social institutions.

Schank does not indicate how he thinks we come to possess our representation language. Suppose, as some writers have, that it is innate. In that case, one would expect its ACT categories to be universal. This need not imply that all the categories of CD are employed in every possible culture: in discussing the notion of transfer of ownership, Schank remarks that 'it is possible to conceive of a culture and therefore a language that would have a different set of those abstract relations or none at all (and thus no ATRANS)' (1975, p. 55). In such an instance, the universal ACT primitive ATRANS would simply not get actualized. But another speculation suggests that CD might not turn out to be universal after all: 'If in fact, there exists a culture where life is viewed as a continuum rather than a series of distinct actor-action events, Conceptual Dependency would not do as a conceptual model of such a culture' (Schank 1973a, p. 206). In that case, Schank's recourse would presumably be either to hold (1) the (very unparsimonious) view that each of us has multiple innate representation languages, perhaps only one of which actually gets employed by a given individual, or (2) that the representation language is not innate. On the latter alternative, one's representation language would presumably be acquired in a linguistic environment, which raises the suspicion that it may not really be a necessary condition for the understanding of natural language (cf. Fodor 1975, Chap. 2, for further arguments). For if the representation language itself could be acquired (without an underlying language), why could a natural language not be acquired this way also?

3. CONCEPTUAL SYNTAX

Despite the subordination of natural-language syntax to meaning in Schank's sentence analyzer, there is considerable emphasis given to the

formal structure of Conceptual Dependency diagrams. Here are some examples (Schank 1975, pp. 38-9):

$$(1) \quad PP \iff ACT$$

$$(2) \quad ACT \xleftarrow{O} PP$$

$$(3) \quad ACT \xleftarrow{D} \begin{cases} \rightarrow PP_1 \\ \rightarrow PP_2 \end{cases}$$

$$(4) \quad ACT \xleftarrow{R} \begin{cases} \rightarrow PP_1 \\ \rightarrow PP_2 \end{cases}$$

$$(5) \quad ACT \xleftarrow{I} \begin{array}{c} \wedge \\ \parallel \\ \vee \end{array}$$

Rule (1) says that Picture Producers (usually, animate agents) can perform actions; Rule (2) means that ACTs can have objects; Rule (3) means that ACTs can have directions (a variable in place of PP_1 or PP_2 indicates an unknown position). Rule (4) indicates that ACTs can have recipients. And Rule (5) says that ACTs have instruments that are themselves complete (completeness is shown by the double lines between two arrowheads). Rules (2) through (5) collectively represent conceptual cases – modifiers of ACTs – and a specific number of them (either two or three) is required by every ACT.

In some instances, the CD syntax rules *permit* a particular structure; in other cases they *require* it. One such requirement involves the Instrumental case. Returning to the CD representation of 'John ate a frog', notice that while the Instrumental case depiction is a reasonable inference, it does not represent information explicitly given in the sample sentence (John could have done this even though he possessed no hands). In fact, although Schank will insist that ACTs always have instruments, he frequently omits their depiction in instances where the

instrument was not expressly specified, and cannot be known with virtual certainty.⁷ If the instrument is not filled in by the *analysis* of a sentence into CD, it can be obtained by an inference mechanism that relies on default judgments.

Whether or not instrumental ACTS are explicitly shown, however, the requirement that every ACT must have an Instrumental case encounters a familiar logical difficulty.⁸ Since the instrument of an ACT is itself a complete (and distinct) conceptualization, the occurrence of one ACT will actually require an infinity of ACTs. For example, suppose that John eats ice cream, and that the instrumental ACT is his moving a spoon to his mouth. Now, the conceptual representation of 'John moves a spoon to his mouth' will itself require an instrument, e.g., 'John activates the muscles in his arm', and so on. The problem is that if the performance of any given one ACT presupposes an infinity of ACTs, no ACTs will be possible at all. Schank, unfortunately, views the Instrumental case requirement only as providing a notational inconvenience:

Since an analysis of this kind is not particularly useful and is quite bothersome to write, we do not do so. Rather, whenever we represent a conceptualization we only diagram the main conceptualization and such instrumental conceptualizations as might be necessary to illustrate whatever part we are making (Schank 1975, p. 33; cf. Schank 1973a, p. 201).

He continues:

[T]he ACT in a conceptualization is really the name of a set of *sequential* actions that it subsumes (and are considered to be part of it). These instrumental conceptualizations *are not causally related* since they are *not actually separable* from each other. In actuality, they express one event and are thus considered to be part of one conceptualization. The rule is then, that one conceptualization (which may have many conceptualizations as part of it) is considered to be representative of one event (*Ibid.*, p. 34, italics added).

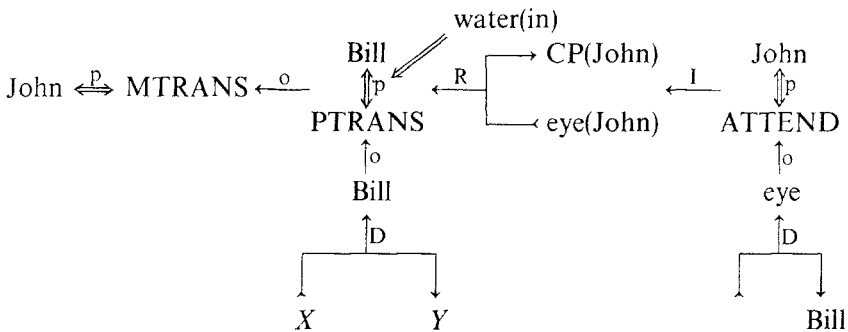
It is certainly plausible to maintain that the name of an event somehow encompasses its constituents, as the phrase 'the third game of the 1988 World Series' might subsume a variety of events making up that baseball game. It is true also that in ordinary circumstances we do not carry the analysis of an act into its constituents very far. But these points do not obviate the fact that the Instrumental case requirement is logically committed to an infinity of constituents for any ACT. Each of those constituents must in turn be an ACT performed by an agent. Moreover, it is not at all clear what Schank means in claiming that the subsumed instrumental conceptualizations are not 'actually separable', since he

also maintains that they are sequential. If they are sequential, they certainly are separable, and they should therefore at least be candidates for causal interaction. In short, the universal requirement of Instrumental cases is highly problematic.⁹

4. PERCEPTION AND IMAGE-REPRESENTATION

As noted earlier, MTRANS involves the movement of a conceptualization into an agent's CP, IM, or LTM, either from the CP of some other agent, or perhaps through memory or sense perception. So, if Mary informed Bill that his car had been wrecked, a conceptualization is 'transferred' from Mary's CP to Bill's, although the term 'transfer' is something of a misnomer, since the information does not vacate Mary's CP in going to Bill's. On the other hand, if Bill momentarily noticed some event that he was later able to recall, there would presumably have been some 'movement' of information into LTM once the noticed event was no longer in Bill's consciousness.

Schank's treatment of perception raises some very interesting and difficult issues. Consider his representation of the sentence 'John saw Bill swimming' (Schank 1975, p. 61):¹⁰



Given Schank's account of various ACT terms, this diagram appears to say that John transferred from his eye to his Conceptual Processor the conceptualization that Bill was swimming.¹¹ It should be remembered that this is intended, not as an account of perception *simpliciter*, but rather as a description of how a layman represents the perceptual process. Of course, few laymen could be expected to use categories such as 'conceptual processor' and 'conceptualization' (see the end of

Section 5 below for further discussion), but perhaps the point here is just that John became aware that Bill was swimming, on the basis of seeing him. No doubt there is something right about this account, insofar as John's conceptualization was derived from sense perception.

At the same time, as one's account of perception becomes more sophisticated, some limitations of conceptual primitives emerge. It is important to bear in mind that, in Schank's account, the object of an MTRANS is always a conceptualization. But, according to the above diagram, the *conceptualization*, before going to the CP, was somehow *in the eye!* There is, I think, no need to be overly literal about what this means: it might involve a retinal image, or a series of images,¹² or perhaps even a percept or something of the sort. Nonetheless, even on this liberal reading, the MTRANS account of perception involves a major problem. The problem is that a conceptual representation must be *constructed from* the available sensory information;¹³ it does not occur 'ready-made' at the sensory level. And what ends up being cognitively available to an agent is considerably *less than* what was given in perception. In the current example, visual information received by John includes such items as whether Bill was nude or wearing swimming trunks, whether he was in a lake or a pool, whether anyone else was near him, whether the sun was bright, etc. – none of which appears in the CD graph displayed above. I do not mean that this omission represents a mistake; no plausible CD diagram would include a description of everything 'contained in' the sensory data, for we are very seldom aware of everything we perceive. There is, in other words, a considerable gap between sensory reception and conceptual representation. Although the process by which the one gives rise to the other is not well understood, the phenomenon itself seems uncontroversial. What the MTRANS account fails to detail, or even to acknowledge, is the process by which a specific piece of information is extracted from the greater wealth of information available at the sensory level.

This point is important enough to warrant consideration from a slightly different point of view. A sensory representation, be it an image, image series, or percept, has the informational richness of a picture.¹⁴ An interpretation of it leads to an informationally impoverished cognitive counterpart. And this interpretation requirement shows why a sensory content (imagistic or otherwise) is not by itself sufficient to yield the content of a Schankian conceptualization. If one asks, 'What proposition does this image (or sensory content) convey?', there

is no single answer, and *a fortiori* no single *correct* answer. The reason is that propositions have truth values, whereas icons do not. As Jerry Fodor has written:

To a first approximation, the kind of thing that can get a truth value is an assignment of some property to some object. A representational system must therefore provide appropriate vehicles for expressing such assignments. Under what conditions, then, is a representation adequate to express the assignment of a property to an object? Well, one condition which surely must be satisfied is that the representation specify *which* property is being assigned and which object it is being assigned to. The trouble with trying to truth-value icons is that they provide no way of doing the former (Fodor 1975, p. 181).

Fodor's idea is that a given image can represent any number of propositions. The *same* image (sensory content), for example, could yield the proposition that Mary is fat, that Mary has a perm, that Mary is wearing clothes, and so on. We may, of course, assign an interpretation – perhaps consisting of a conjunction of propositions – to an sensory representation, but the point is that the assignment involves extrasensory factors.¹⁵

It appears once again that the attempt to account for perception in terms of MTRANSing a conceptualization from a sensory component to LTM has serious limitations. In part, the problem arises from attempting to make conceptual primitives bear too much weight. But the underlying issue is really just a version of a major issue in the study of perception: How does sensory information mediate between the external world and a perceiver's mental model thereof? No doubt it was not Schank's intention to address this problem in any depth, and, as noted earlier, it may be that his account does approximate the way in which 'ordinary' people think about perception. No argument is provided, incidentally, to show that this is the case. In fact, it strikes me as unlikely that the layman has any view at all about whether propositional structures are contained, say, in a retinal image, or constructed from that image. But if the layman opts for the first account, as is suggested by Schank's analysis, the present discussion indicates that the 'common sense' account cannot be extended very naturally into a more accurate and complete version.

5. IS CD NECESSARY?

Does the use of CD graphs, including ACT primitives, make for a reasonable viewing of the way in which human beings represent events? Let us begin by canvassing some arguments on its behalf:

- (1) Sentences that we understand are generally not stored in the form in which we encounter them. 'If two sentences with different words mean the same thing, then it is not easy to recall which particular words were used after a certain time' (Schank 1975, p. 17).
- (2) Although sentences in a natural language may be ambiguous, there must be some way of representing unambiguously each possible meaning, since 'the original meaning that the speaker chose to impart was unambiguous' (Ibid., p. 15).
- (3) Human hearers are able to supply information that was not explicitly given in a natural language sentence. The inferential mechanisms that humans can apply to natural language sentences would not be economically accounted for in terms of natural language storage (Schank 1981, p. 16).

To take these points in order: it is true that understanders engage in automatic paraphrasing, and without special effort they cannot repeat verbatim much of what they have heard. But it does not *follow* that the paraphrase must be represented in anything like CD graphs. While CD representation certainly provides a possible account, it is not obviously preferable to some alternative. Might it not be the case that my representation of the meaning of natural language sentences is stored in the form of natural language expressions?

Whatever the ruling on this issue, however, the appeal to CD representation as a *required* vehicle for natural language understanding must face a logical dilemma. For, if understanding natural language requires us to map natural language expressions onto CD, then in what does our understanding of CD consist? Clearly, it cannot be maintained that we in turn map CD expressions onto some other type of representation, for an infinite regress would then be under way. So it must be possible for us to understand CD expressions without mapping them in turn onto a meaning representation language. But if we can understand CD representation without doing any sort of mapping, then why should it not be possible to understand natural language without doing any mapping either?

Fodor has considered this argument, and has offered a reply to it. He draws an analogy between a human's mapping of natural language onto an inner representation scheme, and a compiler's producing machine code from a higher-level language. Although a compiler (or

interpreter) is required to render the higher-level language usable, there is obviously no need for a second compiler to make the object code usable. Fodor writes:

What avoids an infinite regression of compilers is the fact that the machine is *built* to use the machine language. Roughly, the machine language differs from the input/output language in that its formulae correspond directly to computationally relevant physical states and operations of the machine

[T]here are two ways in which it can come about that a device (including, presumably, a person) understands a predicate. In one case, the device has and employs a representation of the extension of the predicate, where the representation is itself given in some language that the device understands. In the second case, the device is so constructed that its use of the predicate (e.g., in computations) comport [*sic*] with the conditions that such a representation would specify. I want to say that the first is true of predicates in the natural languages people learn and the second of predicates in the internal language in which they think (Fodor 1975, p. 66).

This analogy between machines and humans breaks down in one crucial respect, however. For as Fodor himself has pointed out elsewhere (Fodor 1981, ch. 8), the meaning of anything couched in machine's 'representation language' is entirely parasitic upon the higher-level source code (whose meaning may in turn be dependent upon a programmer's intentions). A given bit pattern, in other words, might represent an ASCII alphabetic character, a positive integer in base ten, or a two's-complement negative number. There is nothing intrinsically representative about the machine-level 'language'. Yet a 'language of thought' such as Schank's (or, for that matter, Fodor's) must possess its representational power *independently* of any higher-level language, since it is supposed to be a precondition of learning or processing a higher-level (natural) language. Machine language, however, possesses no (classical denotative) semantics of its own; in a word, what the computer 'understands' has no meaning. What kind of understanding is this supposed to be? Needed, of course, is a convincing account of how Conceptual Dependency (or the Language of Thought) gets its reference to the world. This is the problem that John Haugeland dubbed 'the mystery of original meaning' (Haugeland 1985, pp. 119ff), and it remains a major stumbling block for theories of mind that appeal to a Language of Thought.¹⁶


Turning at last to the second argument for CD representation (concerning the need for unambiguous representation), although the various meanings of ambiguous sentences should be susceptible to unambiguous

representation, it does not *follow* that the representation must be interlingual, no matter whether the ambiguity is syntactic or semantic in origin. The multiple interpretations of syntactically ambiguous sentences (e.g., 'Flying planes can be dangerous') as well as semantically ambiguous sentences (e.g., 'I went to the bank') can be specified unambiguously in English. In fact, we quite commonly do so.

The third point (concerning inference) has already been touched upon (see Section 2 above). There I argued that, although inferential parsimony was alleged to derive from a universal, language-independent representation scheme, neither inferential adequacy nor language-independence can in fact properly be claimed for it. Here I want to add that CD theory fails to provide a plausible account of inferences involving intensional contexts.

Consider the sentence 'John believes that the Evening Star is red'. Even though the Evening Star is identical to the Morning Star, the substitution of the phrase 'Morning Star' into our example sentence may not be truth preserving. As is well known, extensionally equivalent terms are not intersubstitutable in belief-contexts, or indeed in *any* contexts describing a psychological state. Notice also, however, that even with *intensional* equivalence added, the terms may not be intersubstitutable in those contexts either; the sentence 'John believes that horses are four-legged and eat grass' does not entail 'John believes that horses are graminivorous quadrupeds', despite the relevant synonymies. This is not merely a point about the relationship between sentences, for there is a related psychological point here as well. Insofar as the ascription of beliefs to an agent provides a vehicle for explaining that agent's actions, we need to know, not just what the agent believes, but (so to speak) how he or she believes it. The mental representations of the two sentences about horses are presumably different¹⁷ and can therefore be expected to produce somewhat different causal consequences. Thus, it should not be supposed that both descriptions of John's belief are equally apt.

The connection of these reflections with ACT primitives is as follows. Even if ACT primitives captured the meaning of natural language verbs, it does not follow that they can enter into the propositional content of anyone's psychological attitudes. Take, for example, Schank's CD representation of the sentence 'John remembered that he forgot to bring his sandwich to school' (Schank 1975, p. 60):



According to this diagram, the *content* of John's memory is approximately this: he remembered that he failed to MTRANS into his CP the fact of PTRANSing his sandwich to school. What has happened here is Schank has incorporated a theory about *how* people represent things into an account of *what they represent*. And in this case, the account is highly implausible. Memories, like beliefs, figure into explanations of action; hence, if John really remembered that he failed to MTRANS something, he might be expected, say, to issue a report couched in those terms. But, unless John happens to be one of Schank's disciples intent upon making an idiosyncratic point, he will do no such thing. I conclude that there is a large class of sentences—those describing the contents of psychological attitudes for which the use of ACT primitives in CD representations will not work. They simply do not accurately depict the contents of anyone's psychological states. Or, perhaps more accurately, they fail to allow for a distinction between *de dicto* and *de re* construals of those states.

6. CONCLUSION

There is no doubt that computer programs built upon Conceptual Dependency theory have produced (in carefully delimited microworlds¹⁸) some interesting and striking results. Viewed from a pragmatic perspective, this approach is salutary; it may, for example, lead to increasingly user-friendly machines with natural language interfaces. But the goal of many researchers in artificial intelligence, Schank included, is to produce a psychologically valid computer model of human mental processes. I have set forth a number of serious difficulties which Conceptual Dependency theory – the bedrock of Schank's approach – must face. These difficulties, I believe, collectively constitute a strong reason for doubting that Conceptual Dependency theory provides (a component of) a plausible theory of mental representation.

I shall conclude by reiterating three problems that strike me as most important. (1) The reduction of human actions to a small group of primitives has little promise of succeeding, as does the parallel project of reducing objects to properties; (2) Primitive ACTS cannot plausibly serve to capture the contents of intentional states; this provides a second reason for thinking that Conceptual Dependency diagrams cannot get by with ACT primitives everywhere replacing natural language action verbs; and (3) The kind of project envisaged by Schank ultimately requires a 'naturalized' semantics for the 'language of thought', whether that 'language' appears in a human being or a machine. Very little work has been done to date in this area, and in this respect, Schank's program is no worse off than many other computational theories of mind. I do not know whether such a project of 'naturalized' semantics can possibly succeed. If points (1) and (2) are correct, however, the 'language of thought' requiring a naturalized semantics will not look very much like the language of Conceptual Dependency Theory.

NOTES

¹ An ancestor of this paper was written while I was on sabbatical leave from the University of Michigan, Flint, during which time I held a fellowship in the Computer Science Department of Wright State University. Revisions were made while I held a Visiting Lectureship at the University of Waikato. I am grateful to these three institutions for their support, and to James H. Fetzer, David Hemmendinger, and Edwin Hung for helpful comments on earlier versions.

² CD theory has played a less explicit role in Schank's most recent work (Schank 1977, 1980, 1982, and 1986), where the aim has been to develop accounts of 'higher-level' knowledge structures. Nonetheless, much of CD theory is embedded in the later work as well; consequently, its difficulties tend also to be absorbed into the more recent accounts.

³ *Time Magazine* (International Edition), July 24, 1989, p. 64.

⁴ The number of primitives Conceptual Dependency theory has varied; at one time fourteen were employed. This is understandable, given Schank's experimental approach to the issue of representation.

⁵ Christopher Riesbeck writes:

What is lacking is a well-defined internal structure for PPs [Picture Producers]. Presumably a PP is a bundle of features, but how many features there are, how many it takes for an object to qualify as a certain kind of PP, how features relate to each other, how features which are discrete relate to the perception of a world that is not, all these questions are unanswered (Riesbeck, in Schank 1975, p. 114).

Although Riesbeck registers no skepticism about the *possibility* of providing 'a well-defined internal structure for PPs', the history of unsuccessful attempts by phenomenologists

to 'reduce' physical objects to collections of sensory data suggests that a workable "feature" account of PPs is not likely to be forthcoming.

⁶ This kind of reflection led Fodor to argue for 'methodological solipsism', one point being that psychological explanations involve opaque construals of mental states (Fodor 1981, Chapter 9). If you and I are presented with different examples of a metallic substance, gold in color, we may both acquire the belief that 'My sample is genuine gold'. And *pace* Schank, this belief may go into both of our LTMs even though I was viewing iron pyrite, and you were viewing the genuine article. In this case my LTM contains a false belief, while yours contains a true one.

⁷ For example, the representation of 'John punched Mary' would show John's moving his fist as the instrumental ACT, while the CD representation of 'John went to New York' would not specify the instrumental ACT (he might have taken a plane or a bus or a train, etc.).

⁸ The universal requirement of an Instrumental case was explicitly absent in Schank's early discussions. See Schank 1972, especially pp. 572 and 574.

⁹ It should be remembered that Conceptual Dependency theory is intended, in part, as a model of human psychology. But no psychological argument is provided for claiming that we must represent every action as having an instrument. Here, defenders of 'basic actions' might extend their ontological claims into the psychological realm, the idea being that basic actions constitute a necessary ingredient in our representation of actions.

¹⁰ Notice that no distinction is drawn here between (1) 'John saw Bill swimming', and (2) 'John saw that Bill was swimming'. In fact, Schank's version of CD theory does not appear to distinguish these two expressions. Since the object clause of (1) is the object of an MTRANS, it can go into LTM, and therefore serve an object of knowledge. But this must mean that the object clause of (1) is being regarded as propositional.

¹¹ Swimming, of course, means more than moving oneself about in the water, but Schank employs the current shorthand just for the purpose of illustration.

¹² See, for example, Dretske (1981, ch. 6, p. 145). The entire Chapter 6 of his book is an excellent and fascinating study of perception from the standpoint of information theory.

¹³ In the example under discussion, much of the sensory information is visual, but important clues may also include proprioceptive data. Information concerning turning of the head, for example, is valuable for determining whether it is the perceived object or the perceiver that is moving.

¹⁴ This analogy is suggested by Dretske, *op. cit.*, pp. 137-143, and does not imply that sensory representations need be imagistic.

¹⁵ It is worth noting that other sorts of representation schemes are subject to the same point. The inscription 'my banjo', for example, could be taken as a referring phrase, or as a series of alphabetic characters. Context and convention, not the 'representation itself' are the crucial determiners.

¹⁶ For interesting, related discussions of this issue, see Fetzer (1988a, 1989) and Searle (1980).

Fodor has recently tackled the problem of a naturalized semantics (Fodor 1987, Chapter 4), espousing a version of the causal theory of content. The causal connection between an object (e.g., horse) and an appropriate internal token ('horse') may be mediated by intentional links. This would appear to undercut the claim that a naturalized semantics

has been achieved. Fodor deals with this issue by claiming that 'For purposes of semantic naturalization, *it's the existence of a reliable mind/world correlation that counts, not the mechanisms by which that correlation is effected*' (p. 122, Fodor's italics). Who says so? It looks to me as if this is pure stipulation, and that semantic naturalization is achieved by leaving something essential—a description of the causal mechanisms—out of the story. Once the underlying intentional mechanisms are described, however, the claim to semantic naturalization for internal tokens loses its plausibility.

¹⁷ Even if they were the same, however, it would not follow that the respective beliefs were the same. As Putnam has repeatedly urged—most recently in Putnam 1988—agents' beliefs are a function, not only of mental representation, but of the contexts in which agents are embedded.

¹⁸ Dreyfus 1979 has built an extensive case for his contention that the results obtained in such microworlds do not admit of generalization, and do not constitute an appropriate model for human mental processes.

REFERENCES

- Dennett, Daniel C. and Douglas R. Hofstadter (eds.): 1981, *The Mind's I*, Basic Books, New York.
- Dretske, Fred I.: 1981, *Knowledge and the Flow of Information*, The MIT Press, Cambridge.
- Dreyfus, Hubert L.: 1979, *What Computers Can't Do: The Limits of Artificial Intelligence* (revised edition), Harper Colophon Books, New York.
- Fetzer, James H.: 1988a, 'Signs and Minds: An Introduction to the Theory of Semiotic Systems', in *Aspects of Artificial Intelligence*, Kluwer Academic Publishers, Dordrecht, pp. 133–161.
- Fetzer, James H. (ed.): 1988b, *Aspects of Artificial Intelligence*, Kluwer Academic Publishers, Dordrecht.
- Fetzer, James H.: 1989, 'Language and Mentality: Computational, Representational, and Dispositional Conceptions', *Behaviorism* 17, 21–39.
- Fodor, Jerry A.: 1975, *The Language of Thought*, Thomas Y. Crowell Company, Inc, New York.
- Fodor, Jerry A.: 1981, *Representations*, The MIT Press, Cambridge.
- Fodor, Jerry A.: 1987, *Psychosemantics*, The MIT Press, Cambridge.
- Haugeland, John: 1985, *Artificial Intelligence: The Very Idea*, The MIT Press, Cambridge.
- Putnam, Hilary: 1987, *Representation and Reality*, The MIT Press, Cambridge.
- Rieger, Charles J. III: 1975, 'Conceptual Memory and Inference', in Roger Schank, *Conceptual Information Processing*, Elsevier, Amsterdam.
- Schank, Roger: 1972, 'Conceptual Dependency: A Theory of Natural Language Understanding', *Cognitive Psychology* 3, 552–631.
- Schank, Roger: 1973a, 'Identification of Conceptualizations Underlying Natural Languages', in Schank and Colby, *Computer Models of Thought and Language*, W. H. Freeman and Company, San Francisco, pp. 187–247.
- Schank, Roger C. and Kenneth Mark Colby: 1973b, *Computer Models of Thought and Language*, W. H. Freeman and Company, San Francisco.

- Schank, Roger: 1975, *Conceptual Information Processing*, Elsevier, Amsterdam.
- Schank, Roger and Robert Abelson: 1977, *Scripts, Plans, Goals and Understanding*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.
- Schank, Roger: 1980, 'Language and Memory', reprinted in Barbara J. Grosz, Karen Spark Jones, and Bonnie Lynn Webber: 1986, *Readings in Natural Language Processing*, Morgan Kaufman Publishers, Inc., Los Altos, California, pp. 171-91.
- Schank, Roger C. and Christopher Riesbeck: 1981, *Inside Computer Understanding*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.
- Schank, Roger C.: 1982, *Dynamic Memory*, Cambridge University Press, Cambridge.
- Schank, Roger C.: 1986, *Explanation Patterns*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.
- Searle, John R.: 1980, 'Minds, Brains, and Programs', reprinted as Chapter 22 of Dennett and Hofstadter, 1981.

Dept of Philosophy
University of Michigan, Flint
Flint, Michigan 48502
U.S.A.