

PRIMITIVES AND WORDS

Yorick Wilks  
 Istituto per Gli Studi  
 Semantici e Cognitivi  
 Castagnola, Switzerland

We may usefully distinguish between internal and external questions when discussing the use of primitives for representing natural language content and doing related semantic computations. Here I shall give a few examples of internal questions; go on to explain why I shall turn immediately to external questions; and finally discuss two of the latter: the justification of primitives in general, and the distinction, if any, between primitives and words.

What I mean by "internal questions" about primitives are detailed considerations about what semantic primitives to choose, or how to insert them into larger structures in particular cases so as to represent some complex concept or conceptual relation, etc. These are questions that can only arise when the general notion of semantic primitive has already been accepted. It is not possible to discuss such internal questions while one is at the same time answering external questions, such as the justification of semantic primitives in general.

There are fairly straightforward internal criteria for the selection and maintenance of a primitive vocabulary: the vocabulary should not be obviously redundant, with two primitives covering the same, or nearly the same, range of meaning. If one could show of anyone's suggested set of primitives that this was so, it would follow that he did not have a good set. Secondly, a primitive vocabulary should not be obviously oriented towards a particular subject area, if it is at the same time claimed to be a general set. So, for example, if, in a proffered set of primitive actions, we saw a majority of primitives concerned with human bodily actions such as moving, expelling, ingesting, etc., we might well wonder how such a system would cope with the expression of general actions such as "divide", "separate", "specify", "undertake", "delay", etc.

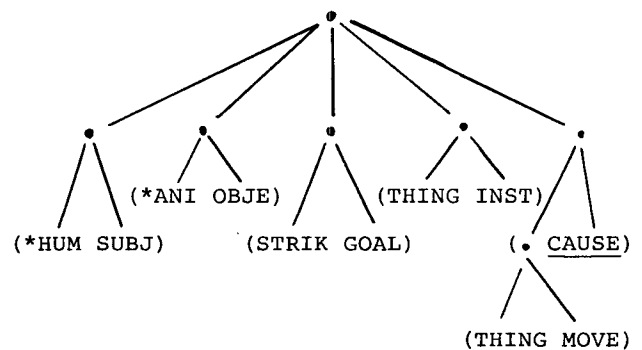
Now there would be a drawback in centering a contribution, for a meeting as general as this one, on only internal questions. To my knowledge there are only two effective systems (in the AI approach to natural language understanding\*) that make use of primitives, namely Schank's (1973a) and my own (Wilks 1972), and therefore any such concentration would be something of a

\*It will be clear here that I refer only to AI systems. There were, of course, working primitive systems in the earlier days of Computational Linguistics, of which the best known was perhaps Gardin's (1965), and in Generative Linguistics there is the well known work of Fodor, Katz, Postal, Lakoof, McCawley, etc.

private fight. External questions, however, are of wider interest and it is my contention that, if we follow certain external questions through, we can see that other workers in natural language understanding are also using primitives, though they may not be aware of the fact.

One issue bridges the gap between internal and external questions in an interesting way, and will serve me here as a new point of departure. A point of difference between Schank's view and my own has always been over the appearance in semantic representations of what appeared to be simply the surface words of the language. So, for example, in a representation of "John shot Mary" by Schank (1973b) there will normally appear the English word "gun". Whereas, in my representations, all one sees are structurings of the 80 primitives, and "gun" is not one of them. I have suggested that Schank's diagrams are therefore of mixed type, as between primitives and words. However, I now believe this criticism of mine to have been badly put, and to rest upon a confusion, namely that there is some clear distinction between primitives and non-primitives (i.e. surface vocabulary). Historically, what happened in Schank's system is clear: he wanted a surface-free semantic representation without words in it, and has gradually achieved this, most recently by the elimination of many noun-words in favour of Fodor and Katz type primitive lists, as developed by Weber. I am proposing now in my own system to reverse the process, as it were, and starting from an inventory of only primitives (about 54 in 1967, now grown to 80, see Wilks (1972)) to begin to insert nonprimitives, i.e. words, into the formula structures that define word senses, but with one important proviso that I shall now explain.

Let us consider the English action "fire at", as a single unit, and its expression in terms of more primitive concepts. I have expressed it until now (1975a) as the formula, or tree structure, below:



The structure need not detain us except to note that the right-most primitive CAUSE is the head, or principal primitive, of the action, and the rightmost THING that is caused-to-move is, of course, the bullet, while the leftmost THING, that is the INSTRUMENT, is the gun.

What I now propose to do is to insert English words into the formula: in particular the words "gun" and "bullet", as follows (in list form):

```
((HUM SUBJ) ((ANI OBJE) ((STRIK GOAL)
((gun INST) ((bullet MOVE) CAUSE))))
```

but to do this if and only if there are also formulas elsewhere in the system defining the meanings of "gun" and "bullet". Thus, in the decomposition of a formula, two kinds of entities may be encountered, primitives and "words", and if an atom is found that is not a primitive then it must be a word and so a formula must exist elsewhere for it. Now, there is a danger of circularity in all this: in that, looking for the formula defining "gun" as that atom appearing in the formula for "fire at", we might find a formula that told us (in primitives) that a gun was an object used by human beings for firing at something. True, but since the head (rightmost element) of the formula for "gun" will be THING, this re-entrant method of formula construction cannot give less information than formulas consisting only of primitives, as mine have until now.

Now, does this method of re-entrant formulas -- re-entrant in the sense of mentioning formulas inside each other (as that for "fire at" now mentions that for "gun" and "bullet") -- force me to withdraw the "mixed-type" criticism I made in the past of Schank's descriptions. Well, yes and no. No, in that Schank's use of, say, the word "gun" in the conceptual dependency representation for "shoot" was not re-entrant in the sense I have defined it. There was, if I understand him, no formula elsewhere for "gun" in his example: there was the English lexeme and no more. Hence, if a program parsing his system searched in a text for "gun" as the instrument of "shoot", and found "bow and arrow", it would be helpless because it did not find exactly what it was looking for -- it would have no descriptive formula for the sense of "gun" to help it know when it had found roughly the same sort of thing. Just as in "John shot her with a colt", it would have no information with which to separate the horse and gun senses of "colt": it would either find the sought English lexeme or, as in these cases, not.

But then again, on the other hand, the criticism is changed because the claim that there should not in principle be mixed type (word and primitive) semantic descriptions is implicitly withdrawn by the above proposal for re-entrant formulas. Now there need be no serious "theoretical" considerations involved in such a proposal. It can be seen as simply a notational convenience: in that "gun" in a formula for "shoot" is now just a shorthand form for the formula for "gun" existing elsewhere in the system. This makes the formulas easier to read for a human user by avoiding the insertion of too much repetitive material in terms of primitives into the body of the formula itself.

Yet, there is a more fundamental point buried in all this, very much an external question, and one which led me to mis-state my criticism of mixed type descriptions in the past. Let me now go into this a little: though it is not, as I shall show, intended to lead to conclusions that should give any comfort, either to those who have used mixed type descriptions in the past or to those who wish to avoid primitive descriptions altogether.

My basic error was to defend the use of primitives by implicitly assuming that they were in some way essentially different from natural language words, and therefore, that structures of primitives were in some way different from structures of natural language words, which is to say, assuming that a semantic description in terms of primitives was something other than a reduced micro-language, with all the normal weaknesses and vaguenesses of a natural language. I now see more clearly that that is not and cannot be so. Or, to put the message of this position paper in a slogan form, there is no escape from a natural language and certainly not via the primitive route.

Let me take first the justification aspect of this. Those of us who have used primitives for some time have all I think made use of the innocent notational convention of writing primitives in upper case letters thus: MAN, CAUSE, TRANS, etc. And I would not wish to change this; it is useful for indicating whether we are, at any given moment, describing a structure in the primitive or the surface language. The trouble is that such a usage does inevitably carry the suggestion that the upper case entities are something other than (over and above, or "deeper" than, but anyway different from) the ordinary words they look like. And this, of course, is error, pure and simple. There have certainly been claims from some quarters that these entities are quite other: that CAUSE has almost a halo around it, and represents, or refers to, some other kind of entity in the brain or mind directly.

At its crudest, this is just the old referentialist fallacy moved up a level, as it were. The low level fallacy, still alive and well, is that the meanings of words are physical objects. "Chair", the story goes, signifies by referring to things like the one I am sitting on, and so therefore do "mind", "action", "friendship" and "cunning", though perhaps in a slightly more roundabout way. I do not want to discuss this view here, but only to point out that the view of primitives under discussion holds that primitives, like words, have their meaning/significance in the same sort of way, but in their cases by referring to certain ill-defined mental entities. I have argued in detail elsewhere (Wilks 1974) that this could not conceivably be so, or be known even if it were so.

But this comic-strip philosophy is not the heart of the matter: what this view of primitives does, in real terms, is to lead

to research that attempts to justify particular sets of primitives directly in some way, and these days that usually means psychologically\*. But what I am arguing is that, if formulas, templates, conceptual dependency structures, etc. are simply usages of a language of primitives, then no direct justification of the vocabulary makes sense, and certainly not any justifications of a correct vocabulary, any more than it would make sense to try to establish the correct vocabulary of English or any other natural language. It follows from this that there can be a variety of primitive languages for semantic descriptions, no one necessarily better or worse than any other, any more than my vocabulary is better or worse than yours if I know 100 English words you don't, and you know 101 that I don't. In the case of each primitive vocabulary, the only ultimate test will be the success or failure of linguistic computations that make use of it.

Now there are limits to this sweeping parallel between primitive and natural languages. Clearly, and as I mentioned earlier, a primitive vocabulary should not have synonymous primitives, whereas in a natural language it is often a point of aesthetic pride to have as large a range of semi-synonyms as possible: "cemetery" and "graveyard" in English, for example. Again, there is at least one interesting piece of indirect justification of a primitive vocabulary, namely the project at Systems Development Corporation that put the whole of Webster's Third International Dictionary onto tape and counted the rank frequency list of words used in the definition of other words. That was, ignoring very frequent words like "a" and "the", down to the 80th rank order, pretty close to my own list of primitives, and naturally I was pleased. That is what I, like anyone else, would have hoped for; since primitives are in fact used to define the senses of other words, we should be happy to be close to the list used unconsciously by the makers of a large and efficient dictionary.

Nevertheless, and in spite of these two caveats, it seems to me that a primitive vocabulary is nothing other than a small natural language, and therefore not open to methods of justification unavailable to any other natural language, and that it is a mistake to pretend otherwise. To sum up, if words of a natural language like English are not justified nor gain significance by their direct reference to things, but only by their function within the overall language, then we may expect precisely the same to be

-----  
\*I cannot see that the notion of the "psychological justification of primitives" makes sense, though I would be happy to be shown. I do not of course refer here to work like (Johnson-Laird 1974) showing that humans seem not to store surface language. Such results are quite consistent with a "primitive hypothesis", but do not support it, since they are also consistent with the hypothesis that human semantic representation is not linguistic at all, but consists of, say, binary numbers!

true of a language of primitives, and no amount of writing in upper case letters is going to change that fact.

Now, if I take this to heart, that there is no difference at bottom between primitives and other words, then I find that another criticism I have made implicitly or explicitly in the past must also be completely recast. This concerns the initial representation of sentences in PLANNER-type formalisms, where it seems not unnatural to represent "John is at the station" as (AT JOHN STATION).

I have been tempted to criticise such representations, to myself at least, on the grounds that they were simply the English words of the sentence (or something very like them) rearranged in some plausible way, and that therefore nothing had been shown or structured. But, as I said, if I take to heart the point that primitive vocabularies are not essentially different from those comprising more obviously surface words, then I cannot maintain that criticism in the same way. Why shouldn't "John", "at" and "station" be somebody's primitive vocabulary?

Now the structural criticism of this method of coding up sentences is unaffected. That is to say, little is revealed by such a method of expression, unless done in a very systematic manner\*, and it normally relies over much on our intuitive appreciation of the structure of the original sentence, in such a way as to be not really the structuring of an example but a mere reprojection of the example itself. To see this, one only has to think of what it would be like to express the first sentence of this paragraph by such a method, the one beginning "Now the structural criticism...etc.".

However, the point at issue here is not this structural one but that of the status of the items in the description: as I raised the question, but did not answer it, why should those English words not be declared to be part of the primitive vocabulary? For, if as I have just argued no serious distinction of type can be maintained between words and other primitives, what could be wrong with that?

Well, it is easy to see what is wrong, given that one accepts one other principle: namely, that one's system, whatever it is, should be extensible in a non-trivial manner. What we are now discussing is the fallacy of the map, to adapt a philosophical cliché. We have a system claiming to represent the structure of natural language but which in fact represents it in the way a map would if its scale was one-mile-to-one-mile. There would be something wrong\*\* with such a map, that much is clear, and similarly there is something wrong with a system which is only extensible on the same scale as what it represents: it adds to its primitives (ordinary words in

-----  
\*The only attempt I know to do this sort of thing systematically is (Sandewall 1972).

this case) at the same rate as it adds to the sentences covered. I am not trying to smuggle back any of the distinction between primitives and non-primitive words that I have abjured, but am only pointing out that, if one chooses surface words as one's primitives, there is nothing theoretically wrong, but just the practical (and in view insuperable) difficulties of (1) inability to state significant semantic generalizations, and (2) the inability to extend one's coverage of the language in anything other than a mile-to-mile manner. It is for this reason that it still seems possible, to me, to give up believing in any difference in principle between primitives and other words but still advocate strongly the use of a sensible selection of words as a reduced, or primitive, sublanguage for semantic expression.

One final point, which is not argument but merely the drawing of a bead on a more distant target. It is the case that the use of a PLANNER, or Predicate Calculus, type of representation in "semi-English" that I have just discussed is intimately associated with the view that the study of reasoning is dissociable, or decouplable, from the study of the semantic structure of natural language, and can be pursued in isolation. In (Wilks 1975b) I have raised a number of doubts about that view, but here I want to add another which follows directly from the argument of this paper.

If it is true, as I have argued here, that there is no escape from a natural language into some other realm, and that a language of representation is just another natural language whether of primitives or of "semi-English", then it follows that there is no special extra-terrestrial sphere for the examination of reasoning\*, but only translations into another natural language. Hence there is no reasoning about natural language separate from natural language, and all we can do is to choose the language in which we prefer to model the reasoning and over which we prefer to compute. Thus, to compress things somewhat, we have the choice between computing about reasoning in a primitive-like language, or one reducible to it by the "re-entrant" method I described, or in one like PLANNER semi-English in which little is made explicit, and which would require another system to make its internal relationships explicit for any but the most trivial examples.

-----  
\*\*One way out I have not discussed here is for someone to argue that, for the representation of language, the mile-to-a-mile scale is not a fallacy, because no significant linguistic generalizations are possible.  
-----

\*I am not denying here that some non-linguistic forms may explicate our reasoning about, say, position in space or numerical relationships; nor am I denying that there is reasoning in the sphere of vision, which is also possessed by dogs, and cannot therefore be linguistic. Granted all that, I am arguing that the sort of reasoning required to understand the argument of this paper, or of a standard newspaper editorial, must (a) be linguistic in nature, but (b) will not be explicated in an interesting and non-circular manner by a system using only semi-English.

It may turn out that it is more sensible to say that language understanding depends on reasoning, rather than vice-versa. Everyone in A.I. seems to believe it without question, and I have done no more here than raise a few small doubts that it might, after all, turn out to be the other way round.

#### REFERENCES

- Gardin, J., SYNTOL, in Artandi (ed) Rutgers Series on Systems for the Intellectual Organization of Information, New Jersey, 1965.
- Johnson-Laird, P., "Memory for Words", Nature, 1974.
- Sandewall, E., "PCF-2, A first-order calculus for expressing conceptual information", Dept. Computer Science, Uppsala Univ., 1972.
- Schank, R., "The Fourteen Primitive Actions and their Inferences", Stanford Univ., AI Lab. Memo No. 183, 1973a.
- Schank, R., "Identification of Conceptualizations underlying Natural Language", in Schank and Colby (eds), Computer Models of Thought and Language, San Francisco, 1973b.
- Wilks, Y., Grammar, Meaning and the Machine Analysis of Language, London, 1972.
- Wilks, Y., "One Small Head", Foundations of Language, 1974.
- Wilks, Y., "An intelligent analyser and understander for English", Comm. ACM, 1975a.
- Wilks, Y., "Natural Language Systems within the AI Paradigm", Stanford Univ. AI Lab, Memo No. 337, 1975b.