

ON ARGUING ABOUT SEMANTICS


Jerry L. Morgan
University of Chicago

Recently several proposals have been made concerning the form and nature of deep structure, and the relation between it and semantic representation. I would like to examine some aspects of one of the most radical of these: namely, McCawley's proposal that the relevant underlying structure for grammatical description is in fact semantic representation. It seems to me that McCawley is correct in claiming, with Lakoff, Goss, and others that there is no autonomous level of deep structure intervening between semantic representation and surface structure. However, it seems likely that for those who reject this view and accept the position outlined by Chomsky¹ in recent papers, many rules long considered to be syntactic transformations will come to have the status of "semantic interpretation rules", and deep structure will turn out to be considerably less deep than has been supposed in much of the literature. If so, then the two opposing theories may eventually evolve into notational variants of each other. In any case, McCawley's model is a useful and important one, with considerable theoretical import, regardless of whether one considers it a model of syntax or of semantics. It is as a framework and a model for research in semantics that I wish to discuss it in this paper, with particular attention to the way it treats the semantic representation of lexical

items.

McCawley argues (McCawley 1967, 1969) that semantic representation must be in the form of tree-like structures which are highly abstract and complex configurations of semantic terms and referential indices. These semantic trees are "operated upon" by rules in the same manner that transformations can be said to operate on syntactic structures—permuting, copying, and deleting elements of the tree. The resulting constituents are then replaced by lexical items. The constituents for which lexical items are inserted need not have been constituents at the deepest level of semantic representation. For example, McCawley has analyzed kill as the derived constituent CAUSE(BECOME (NOT(ALIVE))), created by a rule of "predicate-raising" as in the following derivation.

(1)



(2)



(3)





But the rule of predicate-raising is optional. Whether or not it is chosen at each of the points where it is applicable determines the ultimate derived structure into which lexical items are inserted. This accounts for paraphrase relationships among such sentences as (6) through (9).

- (6) John killed Harry.
 (7) John caused Harry to die.
 (8) John caused Harry to cease to be alive.
 (9) John caused Harry to come to be not alive.

Rules like predicate-raising which are pre-lexical (i.e. which apply before lexical items are inserted) can in principle have no exceptions, since true exceptions are lexical in nature. Therefore rules like this may produce derived constituents for which there is no lexical entry. For instance, predicate-raising would produce from what underlies "cause to cease to be obnoxious" the derived constituent (10), for which there is no entry in the English lexicon (e.g. *deobnoxify or some such entity).



Similarly, non-existent items like *aggrise would presumably be treated as cases of accidental holes in the

lexicon, rather than as absolute exceptions to certain rules, as in Lakoff (1965).

McCawley's semantic representation is in some ways only a much deeper "deep structure"—a representation of the logical structure underlying utterances. It has interesting consequences, both in its implications about the nature of language, and in indirectly raising some questions about the difficulties inherent in the study of semantics. Its implications about the nature of language are of two sorts—the ultimate nature of semantic representation, and more specific consequences of the details of his model. What, after all, is the nature of semantic representation? If the representation of meaning in terms of structures of discrete elements is correct, does this reflect conceptual or cognitive structure, or is semantic representation a purely linguistic object? What is the relationship between semantically primitive terms and the mechanisms of cognition and perception? If a term is semantically primitive, does that mean it is cognitively primitive? The answers to questions like these are a long way off.

The implications of specific details of McCawley's model are slightly more concrete. For example, take the claim that lexical items can only be inserted to replace constituents. What can be a constituent at the point of lexical insertion is limited by the constituent structure of the original semantic representation, by the rules which operate to create derived constituents, and by

constraints on the operation of rules, both of the type described by Ross (1967), and of a type I have suggested may be necessary (Morgan 1968). Thus the constraint (due to Ross, loc. cit.) which permits (11) while preventing (12):

- (11) John saw Mary laying a wreath at the grave of the unknown what?
 (12) *What did John see Mary laying a wreath at the grave of the unknown?

may be what makes impossible a verb sneep with a meaning such that (14) is a paraphrase of (13)).

- (13) John saw Mary laying a wreath at the grave of the unknown hippie.
 (14) *John sneeped Mary laying a wreath at the grave of the unknown.

Therefore, given the set of universal pre-lexical rules, the set of universal semantic primitives, and the set of universal constraints on the operation of rules, these together define the universal set of possible lexical items in their semantic aspect; that is, they rule out as impossible in any language an infinite class of a priori possible "meanings"² a lexical item could have. Furthermore, this model predicts that there may be systematic differences in lexical inventory among languages. Insofar as two languages differ in pre-lexical rules, semantic primitives, or constraints on rule operation, these differences should be reflected systematically in the lexicon, since such differences would cause the two languages to differ in their range of possible derived constituents.

The point which raises difficult questions about methods of research in semantics is the claim that otherwise unanalyzable items like kill have analyzable internal semantic structure. Presumably the claim is that they have internal structure which is somehow psychologically real. Any theory of semantics which does not claim some sort of psychological reality for its constructs is sterile, little more than a game of how-language-would-work-if-I-had-designed-it. Given, then, that the constructs of a theory must have some kind of psychological reality, what kind of psychological reality can the semantic representation of a lexical item have? (Caveat: I am about to set up what I claim are straw men).

First there is the possibility that using a word is a conscious intellectual process of merely putting a name on a concept that may or may not have its own structure. This possibility is easily dismissed. Studies like those of Fillmore (1968: 69) offer abundant evidence that knowledge of the meaning of a word is knowledge that we have very little conscious awareness of; the difference between accuse and criticize, for example. This knowledge is more akin to the sort of subtle, unconscious knowledge one has of rules of syntax. As an illustration of this type of knowledge, consider time machines.³

Suppose I climb in my time machine and travel back in time to five years ago. When the machine stops, I, still sitting in my time machine see me (an earlier version)

reading a book. Call the earlier version of me "old me" and the token of me sitting in the time machine "new me". In this situation, what are the rules for reflexivization? For example, take the sentence

(15) I slapped myself.

If new me, sitting in my time machine, slaps new me, I (new me) can report it by "I slapped myself." Similarly, if new me observes old me slapping old me, I (new me) can report it by "I slapped myself." If new me reaches out from the time machine and slaps his younger counterpart, I can report it by "I slapped myself." But if old me reaches into the time machine and slaps new me, I cannot report it by "I slapped myself." On the other hand, consider

(16) I was slapped by myself.

The applications of this sentence are precisely the complement of those of the active version. It can be used to describe the one case the active version can't be used for, but not to describe the other three; in fact it is somewhat ungrammatical with those readings. The small group of people I have asked agree with these judgements. Whether or not all speakers agree is no doubt relevant, but it does not alter the conclusion I want to draw. The important thing is that each of the speakers involved was able to make clear judgements on these sentences, in spite of the fact that the situations the sentences are used to describe are such that no speaker would encounter them or anything like them in the process of language acquisition.

In fact, it's probably safe to say that none of my informants was ever required to make this type of judgement before. The question is, then, how were they able to make relatively quick and clear-cut judgements? For that matter, how were they able to make any judgement at all? I have no answer to these questions. But one conclusion is unavoidable. Some aspects of linguistic knowledge are totally unconscious abilities of tremendous subtleness and complexity. And knowledge of the meaning of lexical items is in many cases this kind of knowledge.

Still, perhaps this knowledge of the meaning of a word is holistic, with no psychologically real structure at all. Take the word banana, for instance. What is the semantic representation of this word? What am I asserting when I say, "Harry is holding a banana" or "I like bananas"? Am I asserting the same thing I assert when I say "Harry is holding a physical object which is inanimate, non-human, fruit, long and thin in shape, shouldn't be kept in the refrigerator..." and so on? There doesn't seem to be a paraphrase which is both uniquely a paraphrase of banana, and finite in length. Banana seems to be the name of a concept with very fuzzy edges. As an illustration of this, consider the following sequence of events, and imagine that at each step I utter the sentence "This is a banana." First I take a banana, peel it, cut it in pieces, throw it on the floor, stomp on it, stir it with a stick, add purple

dye, spray it with enamel, and set it on fire. At what point does the sentence "This is a banana" become inappropriate? Answers to this question vary considerably from person to person, and many people are unable to make any clear-cut judgement at all. If all words have this kind of "meaning", then the investigation of them is a task for psychologists, not linguists.

But there are cases of lexical items whose meanings have more clear-cut edges. Kill, for instance cannot be used unless several conditions are satisfied. The sentence "X killed Y" is not appropriate to describe an event unless Y was alive previous to the event, Y ceased to be alive in the event, and X bears some causal relationship to the event--the same conditions which determine the applicability of "X caused Y to die" and other paraphrases. Then it is a priori possible that words like kill are merely abbreviations for strings of words, a sort of shorthand way of communicating. But this is not the case. We can show this with the aid of some observations made by Rudolph de Rijk. He proposes that "cease to know" is an accurate paraphrase of forget on the following grounds. First of all, "cease to know" is an intuitively quite plausible description of the meaning of forget--it is intuitively obvious that one can forget only those things that one can know, and to forget something, one must first know it. Further, "cease to know" and forget have the same presuppositions, namely that the speaker presupposes

that their complement is true, as is shown in the following sentences:

- (17) Harry has ceased to know that Shakespeare wrote Hamlet.
 (18) Harry has forgotten that Shakespeare wrote Hamlet.
 (19) *Harry has ceased to know that Marlowe wrote Hamlet, but Shakespeare really wrote it.
 (20) *Harry has forgotten that Marlowe wrote Hamlet, but Shakespeare really wrote it.
 (21) *I have ceased to know that tomorrow is my birthday.
 (22) *I have forgotten that tomorrow is my birthday.

Moreover, any complement that "cease to know" can have can also be used with forget:

- (23) Harry ceased to know { his teacher's name.
 how to swim.
 why birds sing.
 what time it was.
 that yesterday was his wife's birthday.
 Swahili.
- (24) Harry forgot { his teacher's name.
 how to swim.
 why birds sing.
 what time it was.
 that yesterday was his wife's birthday.
 Swahili.

And if a given complement is pragmatically strange for "cease to know", it is equally strange for forget:

- (25) My brother has ceased to know how to lay eggs.
 (26) My brother has forgotten how to lay eggs.

So far this strongly suggests that forget and "cease to know" have the same semantic representation--forget might just as well be an abbreviation for the three words "cease to know" (de Rijk did not, I should mention, propose this;

it is my own straw man). However, it won't work. For, as de Rijk points out, there are cases where forget is not equivalent to "cease to know". For instance, suppose I have some friends in Chicago. Then, if they all suddenly move to Australia, I can say, "I have ceased to know where to turn for help in Chicago." But I cannot say, with the same meaning, "I have forgotten where to turn for help in Chicago." The problem here is that "cease to know" is ambiguous. It can mean either that I have undergone some mental change in relation to what it is that I knew, or that the thing I knew has changed. But forget has only the former meaning.

Thus, if at some point I know who my brother's favorite movie star is, but later this information disappears from my mind, I can say with equal appropriateness (27) or (28).

(27) I have ceased to know who my brother's favorite movie star is.

(28) I have forgotten who my brother's favorite movie star is.

On the other hand, if my fickle brother changes his mind, but I don't know who his new favorite is, I can properly say (27), but not (28).

Also, words like patricide provide evidence against the "abbreviation" hypothesis. Patricide is a paraphrase of the phrase beginning kill and ending father in each of the following pairs of sentences.

(29a) Harry decided against patricide.

(29b) Harry decided against killing his father.

- (30a) Mary decided against patricide.
 (30b) Mary decided against killing her father.
 (31a) I decided against patricide.
 (31b) I decided against killing my father.
 (32a) All the boys decided against patricide.
 (32b) All the boys decided against killing their
 fathers.

Obviously, the meaning of patricide cannot be stated in terms of a string of words. Nor is "killing one's father" an accurate representation, as can be seen by substituting it for patricide in each sentence. It should be clear, then, that paraphrase relationships must be accounted for in terms much more abstract than equivalence relations between strings of words.

One more straw man. It is still possible that there is no psychologically real internal structure to the meaning of lexical items--rather, that knowledge of meaning is actually in the form of a characterization of the infinite set of situations under which it is appropriate to use the item in question, and paraphrase relationships are accounted for by coinciding or overlapping characterizations of this sort. There would be no internal structure to lexical items in a model like this; in fact, it's not clear that the notion would make any sense in such a model. But I think there is some evidence that at least some lexical items do in fact have psychologically real internal structure. This evidence resides in certain types of ambiguity. One of these types, to my knowledge first pointed out by Robert Binnick, is in sentences involving transitive verbs like to jail, as in (33) and (34).

- (33) The Sheriff of Nottingham jailed Robin Hood
in Reading.
- (34) The Sheriff of Nottingham jailed Robin Hood
for four years.

Each of these sentences has two readings. In (33), the first reading is the syntactically predicted one where the scope of "in Reading" is the remainder of the sentence--in other words, an event occurred in Reading in which the Sheriff jailed Robin Hood. But the sentence has another reading, which as far as I can see can be described only in terms of the adverbial phrase having as its semantic scope a portion of semantic structure internal to the verb jail; that is, the reading which is paraphrased by (33')

- (33') The Sheriff of Nottingham caused Robin Hood to be in jail in Reading.

in which the scope of "in Reading" is "be in jail". Similarly, (34) has the (somewhat strange) reading that the Sheriff jailed Robin Hood, and it took him four years to complete the act. But it also has a more natural reading, paraphrased by (34')

- (34') The Sheriff of Nottingham caused Robin Hood to be in jail for four years.

where the scope of "for four years" is "be in jail".

A similar phenomenon occurs with ambiguity in scope of words like again. For instance,

- (35) John knocked Harry down again.

is ambiguous, in that it can be used to mean either that John, having knocked Harry down before, did it again; or that Harry was down for some other reason (he fell, or he

was sleeping, for example) and when he got up, John knocked him down, so that he came to be down again. The ambiguity here is quite naturally accounted for by saying that the scope of again can be either the construction "knock down" or merely "down". Here again, the same sort of ambiguity can be found below the level of lexical item, as in (36).

(36) Bill arrived at 3:00 and left again at 4:00. This sentence can mean either that Bill had left once before and did so again at 4:00, or only that Bill, having been elsewhere before 3:00, came to be elsewhere again at 4:00, with the scope of again internal to the verb left.

Another phenomenon of sub-lexical scope occurs in connection with quantifiers like almost and nearly. A sentence like (37) is ambiguous in the scope of almost.

(37) I almost drank all my milk.

It can be either the entire verb phrase "drank all my milk" or the noun phrase "all my milk" as paraphrased in (38):

(38) I drank almost all my milk.

The solution which suggests itself is that the quantifier can appear in either position in semantic representation, with corresponding difference in meaning, and an optional rule moves it to a higher place in the tree, creating the ambiguity.

This also applies at the sub-lexical level. The sentence

(39) I almost killed John.

has several readings due to ambiguity in the scope of almost, at least one of which involves the scope of almost being internal to killed; i.e., the reading paraphrased by

(40) I caused John to become almost dead.

The full range of readings becomes more apparent in the pseudo-cleft construction:

- (41) What I almost did was kill John.
- (42) What I did was almost kill John.
- (43) What I did to John was almost kill him.

The solution that was suggested earlier can be quite naturally extended to account for the sub-lexical reading if one accepts the hypothesis that kill has internal semantic predicates.

Furthermore, there is evidence that, in relation to certain constraints, these predicates behave like structures, rather than like an unordered bundle of features. There are constraints on rules which move quantifiers--one is that they can't be moved over not. For example, the almost in (44) cannot be moved up as it can in the non-negative

(44) I didn't drink almost all my milk.

version; this sentence cannot be transformed into (45).

(45) I almost didn't drink all my milk.

Similarly, while (46) has a reading equivalent to (47)

- (46) John almost spent all his money.
- (47) John spent almost all his money.

the sentence (48) does not have the reading of (49)

- (48) John almost didn't spend all his money.
- (49) John didn't spend almost all his money.

since this would involve raising almost up over not.

In the same way, (37) has the reading of (50)

(50) I caused John to come to be almost not alive.

but it does not have the reading of (51):

(51) I caused John to come to be not almost alive.

The hypothesis that lexical items have internal semantic structure would account in a natural way for these phenomena. And as McCawley has observed, such structure is remarkably similar to the sort of structure that has traditionally been used to describe syntax. But if lexical items do have such structure, how do we go about finding out for a given item whether it has structure, and if so, what structure? Much of the work being done on the representation of lexical items is being done by brute force of intuition. Other fruitful work can be and is being done in refining paraphrases of lexical items, and in exploring the limits of applicability of the item, thereby more nearly approaching an understanding of what structures and elements must be involved in the semantic representation of the item, and whether the limits of its use are sharp or fuzzy.

Thus, it is not sufficient to say that decapitate means "cut off the head of" and consider the question closed. Since knowledge of meaning is in fact unconscious and very subtle, it will take some rather outlandish research methods and a lot of imagination to approach it. For instance, if John has two heads, and Harry cuts off one of them, can one properly say "Harry decapitated John"? For me, no. This suggests that a better paraphrase for decapitate is "cause to come to have no head". But then if John's behavior ultimately causes him to be beheaded, can

one say "John's behavior decapitated him"? If one dissolves John's head in acid, has one decapitated him? This type of probing, outlandish as it may seem, is an indispensable tool for exploring both the particulars of meaning, and where the speaker's ability to make clear judgements ends. And it will most likely produce some surprising results, just as the consideration of time machines in connection with reflexivization does.

But there are problems with a semantic theory like McCawley's that may make progress difficult. First of all, since semantic structure is in terms of abstract semantic predicates, we may find abstract elements which do not have any clear one-to-one correspondents in natural language. Know, for example, has associated with it the speaker's presupposition that its complement is true. When this is factored out, what is left? Is know think plus presupposition? It's hard to tell, since it's difficult to talk about it. When we appear to have reached the limit of analyzability of the meaning of an item, how will we know that we have reached semantic primitives, rather than merely having reached a point beyond which we have no way to talk about what is there?

Furthermore, there are some indications that the present system of representation may be inadequate. The questions of adequacy arise in connection with the representation of presupposition. It seems fairly clear that presuppositions should be represented in terms of trees;

their content seems to be structured the same way the content of assertions is. The question is how the relation between a sentence and its associated presuppositions should be represented. Many kinds of presuppositions "act like" previously uttered sentences. For example, the presupposition in (52)

(52) The chair you're sitting on is green.

is precisely what is conjoined to the left in (53).

(53) You're sitting on a chair₁ and the chair₁ you're sitting on is green.

The presupposition that determines the definite article is not, on the other hand, made explicit in the order of

(54) The chair₁ you're sitting on is green, and you're sitting on a chair₁.

which is in fact anomalous in some way.

Similarly, presuppositions affect stress exactly as do previous utterances and/or previous parts of the same sentence. Thus,

(55) How does it feel to be a beautiful girl?

is a felicitous question only when its utterer believes his addressee is a beautiful girl, whereas

(56) How does it feel to be a beautiful girl?

may be asked felicitously under just about any circumstances. The stresses here are not contrastive; the difference between the two sentences doesn't reflect a contrast with any opposed concept. On the contrary, the stress on feel in (55) appears to be placed there by the same rule which seems to move stress leftwards one constituent from

stressed elements which are second occurrences; this is what is responsible for the fact that pronouns, and anaphoric expressions in general, lack stress.

Again, consider the difference between (57) and (58):

(57) I'm going outside to see what's happening
in the world.

(58) I'm going outside to see what's happening
in the world.

These sentences have a common presupposition, namely, "I am here." The difference between them seems to lie in the relation between "here" and "in the world"--one has the presupposition that "here" and "in the world" refer to the same place, the other that they refer to different places. There is a generalization to be captured here. The effects of presupposition can be described by rules which are independently necessary to describe the effects of actual previous utterance, or even left-conjoined clauses of the same sentence. But to describe both with the same set of rules, it is necessary to represent presuppositions as trees to the left of the sentences they are associated with.

Moreover, it may be necessary to have some presuppositions tucked away inside trees, rather than having all presuppositions associated with the sentence as a whole. For instance, as I noted above, the verb know must somehow have specified in its lexical entry that there is associated with it the speaker's presupposition that its complement is true. The same is true of several other verbs: regret, annoy and so forth, as in (59) through (62).

- (59) I regret that I live in Chicago.
- (60) I don't regret that I live in Chicago.
- (61) It annoys me that Daley is so paranoid.
- (62) It doesn't annoy me that Daley is so paranoid.

And the construction usually called "counterfactual conditional" as in (63)

- (63) If John were here, he'd help.

involves the presupposition that the condition is not true--in this case, that John is not here. Now what happens with counterfactual conditionals containing verbs like know and regret, as in

- (64) If John were here, I'd know it.
- (65) If he had done that, he'd regret having done it.

If presuppositions are stated merely as a list associated with the entire sentence, these two sentences each have a pair of contradictory presuppositions.

There are also verbs with negative presupposition, like pretend, whose complement is presupposed by the speaker to be false. Thus (66) and (67) are anomalous if uttered by me.

- (66) I'm pretending to be me.
- (67) I'm pretending to be a human being.

But I can quite properly say

- (68) I'm pretending to be Greta Garbo pretending to be me.
- (69) I'm pretending to be Snoopy pretending to be a human being.

What is involved here seems to be related to the notion "world-creating verb" proposed by Lakoff (1968) to account for certain phenomena of reference and pronominalization.

It appears that presuppositions are defined locally, relative to the scope of world-creating verbs, rather than associated with a sentence as a whole. Semantic trees begin to look more and more bizarre.

The difficulties of the type of analysis involved in a theory with pre-lexical syntax are obvious, especially since there is lacking a well-justified criterion for what constitutes evidence for and against a given analysis. Researchers in this area are in the position of physicists trying to discover sub-atomic particles solely by inference from the behavior of other entities. Unfortunately, in semantics, one is never sure whether the "other entities" are trustworthy.

FOOTNOTES

¹See, for example, Chomsky (1968, 1969). Similar viewpoints are argued in various papers by Ray Jackendoff and Joseph Emonds.

²I do not mean to imply by this term that words have meaning in isolation. By the "meaning" of a lexical item I mean the specification of the elements and structures it contributes to the meaning of sentences in which it occurs.

³If memory serves, I owe the time machine question to James D. McCawley.

REFERENCES

- Binnick, Robert. 1968. "On the Nature of the 'Lexical Item'" Papers from the Fourth Regional Meeting, Chicago Linguistic Society, ed. by Bill J. Darden, Charles-James N. Bailey, and Alice Davison, 1-13. Chicago, Department of Linguistics, University of Chicago.
- Chomsky, Noam. 1968. "Remarks on Nominalization." Readings in English Transformational Grammar, ed. by R. Jacobs and P. Rosenbaum. Blaisdell (in press).
- _____. 1969. "Deep Structure, Surface Structure, and Semantic Interpretation." MIT.
- de Rijk, Rudolph. 1968. "A Note on Pre-lexical Predicate-raising; or How to Refute Generative Semantics in the Event That It Can Be Refuted." Dittoed, MIT.
- Fillmore, Charles J. 1968. "Types of Lexical Information." Working Papers in Linguistics No. 2, 65-103. Computer and Information Science Research Center, Ohio State University.
- _____. 1969. "The Lexicology of Guilt." Paper presented at the Southeastern Conference on Linguistics, March, 1969.
- Lakoff, George. 1965. On the Nature of Syntactic Irregularity. Indiana University dissertation.
- _____. 1968. "Counterparts, or the Problem of Reference in Transformational Grammar." Paper presented at the summer meeting of the Linguistic Society of America, July, 1968.
- McCawley, James D. 1967. "Meaning and the Description of Languages." Kotoba no uchū, vol. 2, nos. 9(10-18), 10(38-48), and 11(51-57).
- _____. 1968. "Lexical Insertion in a Transformational Grammar without Deep Structure." Papers from the Fourth Regional Meeting, GLS, 71-80.
- _____. 1969. "Semantic Representation." Paper presented to the Symposium on Cognitive Studies and Artificial Intelligence Research, March, 1969.
- Morgan, Jerry L. 1968. "On the Notion 'Possible Lexical Item.'" Paper presented at the winter meeting of the Linguistic Society of America, December 1968.
- Ross, John R. 1967. Constraints on Variables in Syntax. MIT dissertation.