

UC Berkeley

UC Berkeley Previously Published Works

Title

Some Remarks on AI and Linguistics

Permalink

<https://escholarship.org/uc/item/359059z2>

Journal

Cognitive Science, 2(3)

Author

Lakoff, George

Publication Date

1978-07-01

Peer reviewed

another character from achieving his goal, then that action should be interpreted as the former character acting to prevent his own goal from failing. This knowledge would be used to interpret John's action as a plan to help him win the race.

The opposite of goal competition is *goal concord*. In goal concord, one character helps another perform an action because it suits his own needs. For example consider

- (30) John and Bill were surrounded by Indians and were trying to fight their way out. When John saw an Indian getting ready to fire an arrow at Bill, he yelled at Bill to duck.

To understand why John tried to warn Bill in Story (30), the reader would have had to understand that if Bill were hurt, it would make it harder for John to fulfill his own goal of trying to escape. That is, the reader had to realize that John and Bill had a goal in common in order to interpret John's behavior. If Bill were about to shoot an Indian, on the other hand, it is unlikely John would warn the Indian. John would not take such an action because he is in a goal competition situation with respect to the Indian, unlike the goal concord situation he is in with respect to Bill.

Goal conflict, goal competition, and goal concord are three instances of goal relations that occur frequently in natural language texts. To understand stories in which these goal relationships occur, a reader must have knowledge about the situations in which these relationships occur, and must know how to identify these situations when they arise. A detailed analysis of these situations and the processes needed to understand them appear in Wilensky (1978).

REFERENCES

- Charniak, E. On the use of framed knowledge in language comprehension. In press, 1978.
 Cullingford, R. E. Script application: Computer understanding of newspaper stories. Yale University Research Report 116, 1978.
 Minsky, M. A framework for representing knowledge. AI Memo No. 306, MIT, 1974.
 Schank, R. C., & Abelson, R. P. *Scripts, plans, goals and understanding*. Hillsdale, N.J.: Lawrence Erlbaum Assoc., 1977.
 Wilensky, R. Understanding goal-based stories. Ph.D. thesis, Yale University, 1978.

Some Remarks on AI and Linguistics

GEORGE LAKOFF

University of California, Berkeley

The April 4, 1976, issue of *Cognition* ran an attack by Drescher and Hornstein (D-H) on AI-based linguistic research. There were replies in the May 2, 1977, issue of *Cognition* by Schank and Wilensky (S-W) and by Winograd (W) and a rebuttal by D-H.

The attack seems to have been interpreted by S-W and W, as well as by other AI researchers, as an all but official statement of the position of Chomsky's interpretivist school (CIS), based mostly in the Northeast. In fact, the tone of the replies would better fit replies to Chomsky himself than to two graduate students at The University of Massachusetts and Harvard. The exchange left out the opinions of linguists, both generative and nongenerative, outside the interpretivist school—namely, most linguists. It was for this reason that the editors of *Cognitive Science* asked me to comment briefly on the exchange. I spent many years as a generative linguist and am now sympathetic to much of the AI work on language.

One thing that the exchange never mentioned but that seemed to be lurking in the background was money—in the form of research funding. With government funding sources running low and with a decision by the Sloan Foundation to pour millions of dollars into Cognitive Science, the competition for research funding has been keen. A number of people I have spoken to, both in the AI and linguistics communities, viewed the timing and nature of the D-H attack as being related to funding issues. Was it an accident that such an attack on AI should come when it did, after AI work had been ignored by transformational linguists for so many years? And why did the attack take the form it did—claiming that AI research was "unscientific"? Did it have anything to do with the fact that granting institutions with the most money primarily supported "scientific" research? I do not pretend to know the answers to these questions. But whatever the answers are, it is important in understanding the nature of the exchange to know that such questions were on people's minds. The charge of being "unscientific" is not mere name-

calling: it is a charge with financial consequences. It has to do with what research gets funded and therefore with what research gets carried out.

As for what was said in the exchange, I tend to agree with S-W and W. But I am much more concerned with what was not said. D-H assumed that they had universal grammar and language acquisition on their side. S-W and W never really questioned that; they simply observed that they were working in another paradigm. But as it turns out, actual research in universal grammar and language acquisition by no means supports the CIS paradigm.

Universal Grammar. One might get the impression from the exchange in *Cognition* that most of the work on universal grammar is being done by CIS linguistics. In fact, most empirical research on universals of grammar is being done outside of the interpretivist school. The early foundations were laid by Greenberg, a structuralist. Most contemporary research is being done not by transformational grammarians, but by relational or functional grammarians of various theoretical persuasions. (There is no room to survey, or explain, all that is being done. The following list merely names some of the principal figures and theories in this area.) There are functional relational grammarians, like Keenan and Comrie, who view grammar as being relationally based and who see grammatical relations not as being basic but as being predictable from semantic and pragmatic relations. There are relational grammarians, such as Postal, Perlmutter, and Johnson, who take grammatical relations as basic and not derivable from semantics or pragmatics. More recently, Postal and Johnson have developed a theory of arc-pair grammar, which takes grammatical structure as being in network form (similar to various AI representations) and which abandons transformational derivations. Then there are functional grammarians of various sorts—Givón, Silverstein, Van Valin, and Foley, to name only a few—who base universal grammar on semantic and pragmatic relations directly and deny the relevance of both grammatical relations and transformations. Most of the best current work on universals is being done in one or another of these traditions. This is not an accident. It has only been through a rejection of various central doctrines of transformational grammar that most of the universals so far discovered or hypothesized in these traditions could have been adequately formulated.

Language Acquisition. The principal research in language acquisition over the past several years—especially the work of Bates, E. V. Clark, Slobin, Ervin-Tripp, MacWhinney, Nelson, Newport, E. O. Kennan, and the Rome group (Antinucci, Parisi, Volterra, and their associates)—indicates that much of syntactic structure as acquired by children is a consequence of pragmatic and discourse functions, stages of sensori-motor and cognitive development, the development of processing capacities, social development, and various

aspects of meaning. This work indicates that grammars as children acquire them are not the kind of objects that the interpretive school assumes they are. It indicates that there is no independent or autonomous "language faculty" that is acquired independently of other faculties and that the problem of language acquisition is far more complex and interesting than "grammar construction" in the old generative sense. Even if one takes language acquisition to be the central problem of linguistic theory, the recent work in that field indicates that the Chomskian view of the problem is not only shaky, but most likely wrong in the most fundamental way. The AI approach fares somewhat better, though to a limited extent. It does allow for the possibility of including pragmatic and discourse functions, processing capacities of some sort, and aspects of meaning. But there is as yet no serious developmental work within the AI framework.

Perhaps the most glaring area about which nothing was said in the exchange was that of linguistic phenomena which indicate that either processing models or AI-like representations are needed to account for empirical linguistic facts. As it happens, examples abound in the literature. I will mention only a few.

Indirect Speech Acts. These are cases in which the meaning conveyed by a sentence in context is more than or different than the literal meaning. "Would it be possible for you to take out the garbage?" is usually understood as a request, not merely a question about possibility. What a sentence means in context is a function of its literal meaning (if it has one), knowledge of the world, and general principles of social interaction and conversation. These are generally taken to be in the area of "performance" as CIS linguistics use the term. Yet there are a wide variety of examples in which the well-formedness of sentences depends upon the conveyed meaning as the sentence is used. Compare:

- (1) Can you pass the salt?
- (2) Can you reach the salt?
- (3) Could you please pass the salt?
- (4) *Could you please reach the salt?

Sentence (4) is ill-formed because of a peculiar English constraint on *could*-questions and sentence-internal *please*: the content of the verb phrase must directly express the content of the request. This is not true of *can*-questions that convey requests. There are an enormous number of constraints on English morphemes and syntactic constructions that have to do with what is meant in some social situation as opposed to what is actually present in the surface or even logical structure of the sentence. The literature on the subject is large and a good place to start would be *Speech Acts* (Morgan & Cole, 1975).

Amalgams. English has a variety of sentence types in which the illocutionary force of the sentence is shifted in mid-stream and one or more sentences or sentence fragments appear in the place of noun phrases.

- (5) John invited you! *He never guess how many people to God knows what kind of a party.*
- (6) John is going to I think it's Chicago on didn't he say it was Saturday.
- (7) The seventy-sixers are going to win because who can handle Dr. J?

The occurrence of such constructions is tightly constrained:

- (8) *John invited Max guessed how many people to Sam doesn't know that it's a party.
- (9) *John is going to Sam is sorry that the place he lives is Chicago on nobody said that anything happened last Saturday.
- (10) *The Seventy-sixers are going to win because might Barry guard Dr. J?

The constraints on such sentences are partly syntactic partly semantic, but mostly have to do with pragmatics—that is, the conversational interaction between the participants *as the sentence is being said* and the conveyed, not literal, meaning that an extended version of the inserted fragment would have said at that point in the processing of the sentence (for a discussion, see Lakoff, 1974).

Interjections. Expressions like *ah, uh*, and *oh* were traditionally treated expressions of emotion—independent of the structure of the sentences in which they occurred. However, James, in a series of remarkable studies, showed that such expressions have as much syntactic, semantic, and pragmatic complexity as any morphemes or syntactic constructions in the language. Moreover, they are constraints that have to be learned about English and do not come naturally to nonnative speakers. James even shows that they obey Ross' constraints on movement transformations, which D-H describe as universals of grammar. Here are some examples of the kinds of facts James considered:

- (11) John threw the ball—oh—up.
- (12) *John threw his dinner—oh-up.
- (13) With a hammer—oh—you can build a/step ladder.
- (14) *With a hammer—oh—Bill hit Fred.

For an exhaustive discussion see the James (1972, 1973) and especially her dissertation (James, 1973). Part of her conclusion is that some of the constraints on such expressions involve the way in which issues of memory and emotion arise while sentences are being processed.

Correction and Editing Devices. Among the grammatical constructions of English ruled out by the CIS use of "performance" are those which are used to correct or "edit" sentences while they are being said. Correction is usually thought of as being only a matter of stopping and saying part of your sentence over differently—with the effect of erasing a tape and saying what you really meant. DuBois, however, has shown this to be an entirely erroneous view. There are a variety of correction and editing constructions in English, for example, *that is, well, I mean*, and *or rather*, and they all work differently—they have different pragmatic functions and different syntactic constraints. Evidence for syntactic constraints comes from sentences like:

- (15) I looked up her dress—I mean, her address.
- (16) *I looked up her dress—I mean, up her address.

Unpublished research by Monica Macaulay on "mid-sentence editing" has turned up further regularities, among them cases showing that these constructions obey the coordinate structure constraint.

- (17) I bought 5 bottles of cream soda—well, six bottles.
- (18) *I bought a dozen bagels, 5 bottles of cream soda, and two pounds of lox—well, six bottles.

The uses of the editing constructions are also remarkable. For example, *Well* is used to indicate, in the act of speaking, that the assumptions and standards behind the conversation are being shifted and that Grice's conversational maxims are being maintained relative to some unspoken new set of assumptions, whereas otherwise they might appear to be violated.

- (19) I've known thousands of people like that—well, hundreds.
- (20) She's a great actress—well, she can carry a tune.
- (21) Did you murder your grandmother?
Well, I did put just a little arsenic in her soup.

DuBois also noted that such "correction devices" are used in fluent prose.

- (22) I would like to comment on the racial, or rather racist, policies of the South African government.

The *or rather* is not a correction, but rather a rhetorical device used grammatically in a fluent well-formed sentence. Its effect has to do with how the sentence is understood up to the point of its introduction—that is, its meaning is a function of processing. For the initial study in this area, see DuBois (1974).

Speech Formulas. Bolinger (1976), in his paper "Language and Memory," and Becker (1975) observed that there are an enormous number of speech formulas in speech and writing of all kinds, and that productive

instructions are rarer than might otherwise be thought. Speech formulas vary from such things as *hello* and *good morning*, to expressions that fit regular patterns of the language, but are used very differently, for example, *we have a nice day, your place or mine?, what have we here?, I'll buy that, what's opening?*, etc. There are thousands of such expressions in English. They vary along a continuum from nearly complete frozenness to near productivity. In general, they are tied to stereotyped situations—the sort that are characterized by what Schank and Abelson call “scripts,” as opposed to plans.” In fact, speech formulas have a number of similarities to scripts. Both are tied to specific kinds of contexts. Both have the same form as productive utterances—sentences, on the one hand, and plans, on the other. Both are in various forms of otherwise productive structures—perhaps with some minor idiosyncrasies not found in the productive structures. For example, Schank and Abelson observe that the order in which the check comes in a restaurant script is a matter of ritual and is independent of any aspect of the plan or the nature of the script. Similarly, the same is true of the order of *solemnly* in the formula *I solemnly swear...* Both scripts and formulas seem to develop historically out of their productive counterparts. Schank and Abelson argue that scripts preempt plans. Similarly, speech formula uses preemptive productive uses. For example, *I'll buy that* will almost always be used in its usual sense and not in the act of buying.

To me one of the most interesting similarities between scripts and speech formulas is their use in situations where there is a drain on mental energy. With are part of our “automatic” behavior, and, as such, seem to require less mental energy.” For example, both seem to be characteristic of senility. One of the most striking things about the aged as they become senile is an increase in the use of speech formulas rather than productive expressions. Similarly, senility seems to involve an increased reliance on ritualized scripts rather than productively constructed plans. The link between speech formulas and available mental energy is evident in activities like baseball, where players on the field chatter primarily in formulas (*Chuck it in there! Turn his hat around!*), while the really creative use of language in baseball comes mostly from bench jockeys, who have fewer demands on their immediate attention. In Kempler recently linked the automatic character of speech formulas to a particular well-known occurrence in certain types of aphasia. To my knowledge it has not been studied whether such aphasics can function better in ritualized script situations than in situations where they have to make up productive plans. It would be interesting to find out.

Generative grammar has nothing to say about speech formulas, since it is currently the study of productive constructions. Yet such formulas dominate natural language. All approaches to language at present are not any better. But the similarities between speech formulas and scripts indicate that there is a generalization to be stated. And it crosses the boundary of what CIS

linguistics call the “language faculty.” This suggests that speech formulas can only be understood by giving up the idea of an autonomous language faculty and taking a broader perspective, as AI theorists do.

Semantics. Most approaches to semantics in the tradition of generative grammar, whether generative semantics or interpretive semantics, and in the logical tradition as well (including Montague grammar), try to base a theory of meaning on a theory of truth. In all of these traditions sentences are taken as wholes. The idea that the processing of the sentence itself might have something to do with its meaning or its truth conditions is not considered. If you take sentences containing correction or editing constructions such as those mentioned above, it becomes clear that the meaning of the sentence depends in an important way on how it is processed. But some subtler cases have been discovered where satisfaction conditions depend on mid-sentence processing. There is Fillmore's celebrated example “If you want to save your life, press the little red button in front of you right... NOW!” In tense logics, sentences are evaluated for satisfaction conditions as if the whole sentence occurred at a single time. But in Fillmore's example, the time at which the instruction is to be carried out depends on the internal processing of the sentence, namely, when the word *NOW* is uttered. Other examples have been constructed by Morgan.

(23) John Smith is Harry's murderer.

(24) Harry's murderer is John Smith.

Sentences (23) and (24) are usually taken to have the same truth conditions. But what about (25) and (26)?

(25) Harry's murderer, whose name I am about to reveal to you, is John Smith.

(26) John Smith, whose name I am about to reveal to you, is Harry's murderer.

In Sentence (25) the truth of the nonrestrictive relative clause depends upon when in the whole sentence the clause is uttered relative to when in the sentence the name is uttered. Sentence (26) is simply semantically aberrant.

The above are a few of the facts that have turned up in the past decade that indicate that processing considerations are important in linguistics. Such cases have led me to try to construct theories of grammar in which linguistic rules play a direct role in processing. I started in 1975 (Lakoff & Thompson, 1975a, b) with a theory of cognitive grammar. But as I became aware of various holistic aspects of syntax and semantics, I turned instead to the development of a theory of linguistic gestalts (see Lakoff, 1977). This work is part of a broader theory of the effect upon language of various aspects of human experience—growth and development, sensori-motor abilities,

perception, social interaction, culture, psychological development, personality structure, etc. It is my belief that much if not all of what is universal in linguistic structure can only be explained in such terms. That is, for me a theory of universal grammar explains nothing—it only presents a puzzle: Why should linguistic universals be as they are?

It has now been over twenty years since Chomsky tried to characterize what for him seemed sensible goals for a "scientific" study of language. These goals were based on assumptions which seemed reasonable then to many people (including me), but which with the hindsight of twenty years of linguistics research now seem gratuitous, or implausible, or downright wrongheaded. Among them is the assumption that one can speak of "the language faculty," an autonomous entity, independent of sensor-motor and cognitive development, perception, memory, attention, social interaction, personality, and other aspects of experience. Another is the assumption that language acquisition consists of constructing grammars on the basis of purely linguistic data. Another is the assumption that linguistic structure is independent of linguistic function. Another is the assumption that phrase structure rules and transformations are the right kinds of devices for characterizing linguistic rules. Twenty years ago these assumptions were not seriously brought into question. The battles then were being fought against structuralism and behaviorism. But a lot of work has been done over those two decades. And to maintain those assumptions in the face of what has been learned in that time takes a lot ofchutzpah. And it takes a lot more to claim that research that does not make such assumptions is "unscientific."

My view of AI research is that it is one among many current healthy trends. I do not think the result of AI work on language is all that impressive as yet. I think it shows promise for the distant future but, more important, it provides a formal framework for investigating interesting problems that could not be investigated within generative grammar.

The discerning reader will have noticed that the kinds of facts mentioned in this paper are largely different from those dealt with by both CIS and AI researchers. That, I think, is a function of the current linguistic scene. There are getting to be almost as many approaches to linguistics as there are linguists. The different approaches are generally concerned with different ranges of facts—depending on the kinds of facts that the approach is best fitted to handle, or the kinds of facts that appeal to the linguist doing the research. I think this is a healthy development, except that it tends to be difficult, especially for researchers outside of linguistics proper, to keep up with empirical discoveries in a wide range of linguistic areas. I hope ways can be found to bridge this information gap.

REFERENCES

- Becker, J. Theoretical Issues in Natural Language Processing Schank, R. and Nash-Webber, B. (Eds.) 1975.
- Bolinger, Language and memory. *Language Sciences*, 1976.
- Dresher, B. E., & Hornstein, N. On some supposed contributions of artificial intelligence to the scientific study of language. *Cognition*, 1976, 321-398.
- Dresher, B. E., & Hornstein, N. Reply to Winograd. *Cognition*, 1977, 377-392.
- DuBois, J. Syntax in mid-sentence. *Berkeley Studies in Syntax and Semantics*, 1974.
- James, D. *Proceedings of the Chicago Linguistics Society*, 1972, 8.
- James, D. *Proceedings of the Chicago Linguistics Society*, 1973, 9.
- James, D. The syntax and semantics of some interjections in English. University of Michigan Papers in Linguistics, 1973.
- Lakoff, G. Syntactic amalgams. *Proceedings of the Chicago Linguistics Society*, 1974.
- Lakoff, G. Linguistic gestalts. *Proceedings of the Chicago Linguistics Society*, 1977.
- Lakoff, G., & Thompson, H. Introducing cognitive grammar. *Proceedings of the Berkeley Linguistics Society*, 1975, 1.
- Lakoff, G., & Thompson, H. Dative questions in cognitive grammar. *Functionalism*, Chicago Linguistics Society.
- Morgan, J., & Cole, P. *Speech acts*. New York: Academic Press, 1975.
- Schank, R., & Abelson, R. *Scripts, plans, goals, and understanding*. Hillsdale, N.J.: Lawrence Erlbaum Assoc., 1977.
- Schank, R., & Wilensky, R. Response to Dresher and Hornstein. *Cognition*, 1977, 133-146.
- Winograd, T. On some contested suppositions of generative linguistics about the scientific study of language. *Cognition*, 1977, 151-180.