Richard Hudson

GRAMMAR WITHOUT TRANSFORMATIONS

It is now nearly twenty years since Noam Chomsky's *Syntactic Structures* appeared, and during these twenty years many things have changed in linguistics—not least, the interest that the rest of the world now takes in what we linguists do. The reason for this is clearly because Chomsky claimed to have discovered a window into the human mind, via the study of the structure of languages. The argument is a simple one: a linguist can write a grammar for a language, with some degree of confidence that his grammar is the "right" one for that language, but since the language only exists in the heads of its speakers until linguists or other grammar-writers come along, the grammar that he has written must correspond in some direct sense, and in detail, to the knowledge that exists in the speaker's mind. Therefore, in order to study that particular part of our minds, all we need to do is to write grammars, making sure that they are right, of course, and then study the grammars instead. We can also compare the grammars for radically different languages and see to what extent they are similar or different, and then draw conclusions about general properties of all human beings' minds—and we can even speculate as to how our minds have developed that way, whether by genetic programming or by other means. It

https://doi.org/10.1177/039219217602409605 Published online by Cambridge University Press

is no wonder that linguists are not the only people who think linguistics is important.

However, there is a snag: linguists do not find it at all easy to agree on which grammar *is* the "right" one for any language, though there is a fair measure of agreement as to how grammar-evaluation should be done. Indeed, one of the few things all linguists would probably agree on is that *none* of us has yet written the right grammar for a language, so all we have is a number of partial approximations to right grammars for an ever growing number of different languages; what we disagree about is who is in front at this rather early stage in the race.

This is unfortunate for the nonlinguist hoping to use linguists' grammars as a window into the human mind, of course, since any grammar that we can offer him may turn out to be wrong just at the point that interests him. Perhaps this is rather too pessimistc a view—there are in fact some things on which we seem to agree, not only regarding the phenomena needing to be covered by a grammar, but even the contents of a general theory of language. On the other hand, Noam Chomsky has claimed that the right grammar for a language will turn out to be some kind of "transformational generative" grammar (the kind he himself launched), and it is worth pointing out that this claim is *not* one of the things on which all linguists would agree.

The purpose of this article is to introduce a radical alternative to transformational-generative grammar, from which one might perhaps draw different conclusions about the human mind. In the next section I shall discuss some of the shortcomings of the Chomskyan theory, before explaining how the alternative works, but it may be helpful first to explain briefly how linguists argue about their theories and grammars, to explain the nature of the controversy. The main point is that we are concerned with facts—linguistics, in other words, is an empirical discipline, and linguists' theories have to fit these facts. On the other hand, the connection between the hardest facts and the linguist's general theories about language is an indirect one, mediated by a chain of more abstract facts and less general theories, which is why linguists can disagree so easily over theories while more or less agreeing on the basic data.

The chain connecting the basic data of linguistics, which are

things people say, to the linguist's general theories of language, can be thought of as containing five links, including the two end-points: (1) utterances, (2) judgments on utterances, (3) structural analyses of sentences, (4) rules for generating structural analyses, (5) metatheories about rules. Each pair of links is held together by the fact that each later link has to fit the preceding one, as theories fitting data, thus metatheories can be seen as hypotheses about rules, and rules as hypotheses about structural analyses, and so on. However, it is not possible to work at any point in the chain beyond the first two links (utterances and judgements on utterances) without having to consider *all* the links at the same time, for fear of going wrong—and the main message of the next section is that transformational-generative linguists seem to have gone badly wrong because their structural analyses (link 3) are based on false assumptions about the nature of structural analyses and the rules for generating them, which is a matter of metatheory (link 5).

In order to argue this case, I shall refer mainly to the link between judgments and structural analyses, showing that the analyses (link 3) generated by the rules of a transformational-generative grammar (link 4) do not fit the interpretations (or judgments) which hearers put on utterances (link 2), and that this is because of assumptions made in the metatheory (link 5) for which there is no independent motivation. In the following section, I suggest that if we make a different set of metatheoretical assumptions, we can write a different kind of rule, which generates a different type of structural analysis, which fits much better with the judgments and interpretations that native speakers give to utterances.

It may be helpful to take an easy example of an argument of this type before we get into the meat of the article. Let us suppose that a linguist makes the following metatheoretical assumption (link 5): the form of a sentence can be related to its meaning by means of just two kinds of rule, or statement namely: "word-meaning" rules, which give the meanings for individual words in the sentence, and "structure-meaning" rules, which give meanings for generalized structures defined in terms of word-classes such as "noun". (The first type of "rule" is generally represented simply as a dictionary entry, pairing some form with

some meaning, which is why I used the word "statement" as an alternative to "rule" above). This assumption allows us to move quite satisfactorily down the chain for the majority of utterances: it allows rules which generate structural analyses that fit the interpretations by native speakers of many utterances. Take a simple example like "Cows moo". We find that native speakers interpret an utterance of this sentence in a way that can best be represented by an analysis at three levels in which there is a phonological analysis (for the pronunciation), a syntactic analysis (for the structure), and a semantic analysis (for the meaning); moreover, the meaning can easily be divided into the meanings of the individual words (*cows* and *moo*) plus the meanings of the syntactic structure (for matters such as time reference, whether it is a statement or a question, whether the noun is actor or "patient" or what, and so on).

If all structural analyses were like this one, they could be generated by means of rules which comply with the metatheoretical assumption concerned in this example: that "word-meaning" statements and "structure-meaning" rules exhaust the whole of meaning. However, things are clearly not as simple as this, and we can show that the metatheory in this case is wrong, by pointing to idioms, such as *pull someone's leg* and *kick the bucket* (which each have, alongside a literal meaning, an idiomatic one: namely "tease" and "die" respectively). Thus if we now add "He's pulling your leg", to our stock of utterances, we find that native speakers allow an interpretation which needs a structural analysis that cannot be generated satisfactorily by means of a simple combination of word-meaning and structure-meaning rules. A third type of rule is needed to take account of groups of words. To remedy this situation, we have to go back to the metatheory and revise the assumptions—i.e., the claims of the original metatheory have, in fact, been refuted.

WEAKNESSES OF THE TRANSFORMATIONAL-GENERATIVE APPROACH

The basis of the transformational-generative approach is a metatheoretical assumption about the nature of structural analyses: that structural analyses should be of an extremely simple type,

called phrase-markers. This type of analysis is generally represented by means of the familiar tree-diagrams:

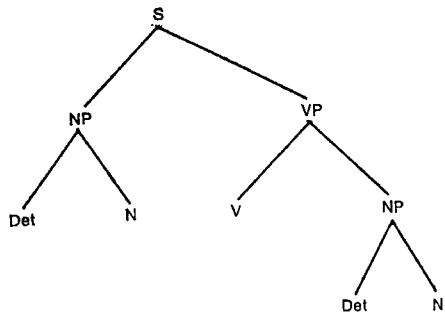

This would be a suitable analysis, say, for "The dog buried a bone". It is extremely simple in that it is equivalent to a labelled, bracketed string, as Chomsky says: in other words, we can take a sentence like "The dog buried a bone" and divide it into successively smaller parts by means of brackets, labelling the bracketed parts (as "S" for sentence, "NP" for noun-phrase, "VP" for verb-phrase, "Det" for determiner, "N" for noun, "V" for verb, and so on):

$$(_{S}(_{NP}(_{Det}\textit{The})\,(_{N}\textit{dog}))\,(_{VP}(_{V}\textit{buried})\,(_{NP}(_{Det}\textit{a})\,(_{N}\textit{bone}))))$$

The reason for insisting on the simplicity of this type of structure is that the alternative described in the next section generates structures which are quite a bit more complex, so it is important to see where the relative simplicity of phrase-markers lies. First, the relations between parts and wholes are very simple: all the parts of a given whole must be next to each other (no discontinuity), and a given part must not belong to more than one whole (no "double-motherhood")—both of these restrictions follow from the assumption that the phrase-marker should be equivalent to a bracketed string, since there is no way of using brackets to show relations such as those represented below:

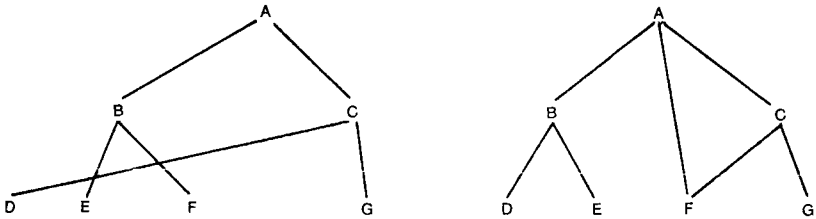

The second respect in which phrase-markers are simple is that the labelling is extremely sparse, consisting of just one class-name per node—so, for example, it is possible to label some part of a sentence "noun-phrase", but not to label it "noun-phrase, definite". It has never been clear why this initial assumption was made—one suspects it was due to the historical accident that similar assumptions were made by structuralist linguists at the time when Chomsky's views were developing. At any rate, it clearly does not follow from the assumption that structures should be equivalent to labelled bracketed strings, since the labelling can be as rich or as sparse as desired in either case.

This set of assumptions about structural representations has always been treated as axiomatic by Chomsky, though he has qualified the assumption about sparse labelling to the extent of allowing much richer labelling on the bottom nodes of a tree, and even suggests that it may be needed on higher nodes too, though without explaining how this is possible in the current transformational-generative framework. If we accept that structural representation must be as simple as this, then two major conclusions follow: first, that the rules for generating such structures should be, or at least can be, "phrase-structure rules"—rules which relate wholes ("mothers") directly to their immediate parts ("daughters"), such as the rule " noun-phrases consist of a determiner followed by a noun" (formalised as NP → Det + N). Allowing discontinuities, double-motherhood or rich labelling of higher nodes would make such rules unworkable.

The second consequence of accepting these assumptions is that a single structural representation per sentence is not enough, since the utterances we actually encounter in normal speech include

some which do in fact involve discontinuity or double-motherhood. For instance, "What do you think she said?" involves two complications: *what* belongs both to the main clause, whose verb is *think* (as question-word) and also to the subordinate clause whose verb is *said* (as object), so it has two "mothers" (the main clause and the subordinate clause), and the subordinate clause is discontinuous (*what... she said*). Clearly sentences like this raise problems if we assume that structural representations mustn't show discontinuity or double-motherhood. One reaction, possibly the most obvious one, is to reject the initial assumption: since discontinuitiy and double-motherhood are clearly properties of the interpretations that speakers give to some utterances, we must allow the structural representations to show such relations, and modify our rules and metatheory accordingly. Chomsky's reaction, however, was to preserve the assumption about simple structural representations but to argue that complex relations (such as discontinuity, double-motherhood and others that have not been mentioned here) can be shown by attribuiting a whole *series* of phrase-markers to each sentence, rather than just a single one. Taking the example, "What do you think that she said?", we can show that *what* belongs to the subordinate clause, as object, by generating one structure (the "deep structure") in which *what* is next to *said*, as its object, as well as another structure (the "surface structure") in which *what* is at the beginning of the main clause (see Fig. 1). This allows discontinuity and double-motherhood to be shown by changing the relations among the elements between deep and surface structure—so Fig. 1 shows that *what* is separated in surface structure from the rest of the embedded clause of which it is a part in deep structure, and moreover that it has different "mothers" in the two structures.

Having decided that we need more than one structure per sentence, the question arises as to what kind of rules are needed in order to generate a series of partly similar structures; and the answer is, reasonably enough, that we need phrase-structure rules to generate one of the structures (the deep structure), and rules of a new type, called *transformational rules*, to convert this structure into the surface structure. (N.B. Transformational rules differ from phrase-structure rules in their formal properties, in that they can relate whole trees to one another, rather than just individ-

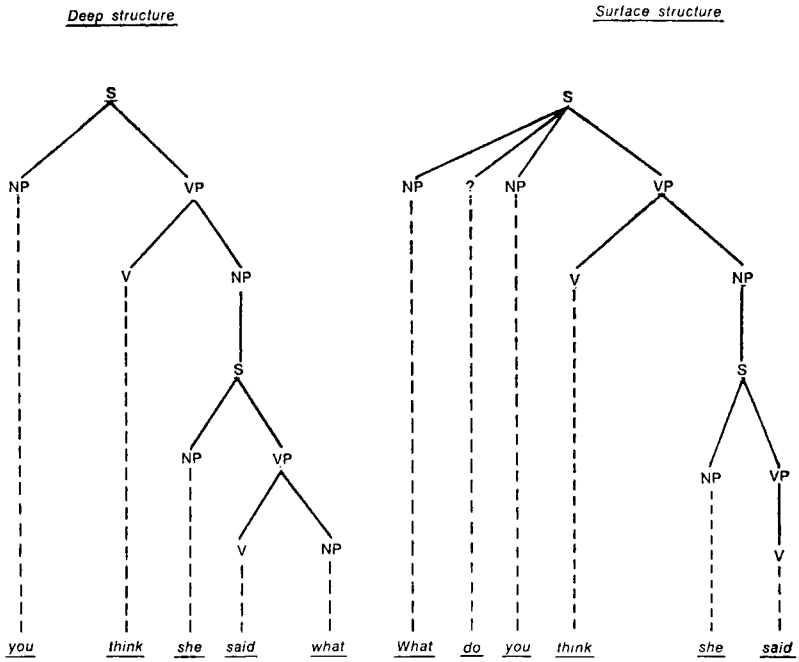

Fig. 1

ual nodes within a single tree). Transformational rules are able to do a variety of things to a structure: they can change the order of elements, they can change the part/whole relations, they can delete elements, they can add elements of certain types, and they can change the classification of elements (within limits). If a particular metatheory allows rules of such versatility to be included in the grammar, it is quite reasonable to make as much use as possible of their versatility in order to handle more complex constructions, and after twenty years of development this theory generally leads to grammars with a large number of transformational rules, at least a few of which apply to *every* sentence—even apparently simple ones like "Cows moo". Moreover, since the transformations apply one at a time, a more or less large number of intermediate structures is generated between the deep and the surface structures. For obvious reasons,

https://doi.org/10.1177/039219217602409605 Published online by Cambridge University Press

the total series of structures, including the deep, surface, and all the intermediate ones, is called a *derivation*.

We can now summarize the Chomskyan position. His metatheory claims that the rules of a grammar include phrase structure rules, which generate one structure per sentence (called its deep structure) and transformational rules, which convert this deep structure in a series of steps into a different, less abstract one called the surface structure). Thus the metatheory (link 5 in our chain) requires the rules (link 4) to generate for each sentence a whole series of phrase-markers (link 3) which, it is claimed, represents the interpretation imposed by native speakers on utterances of the sentences in question (links 2 and 1). It should be noticed that this constitutes a factual claim about the human mind, namely that speakers are incapable of imposing on an utterance any structure more complicated than that represented by a phrase-marker, but that they are capable of constructing a whole derivation, which may consist of tens or even hundreds of individual simple structures for one particular utterance, with (in some cases) quite abstract and complex relations among the structures. This claim seems somewhat implausible, for a number of reasons, not least of which is the fact that ten years of work by psychologists testing Chomsky's claims has found little evidence to support this position.

One reason for doubting the claim is the sheer *quantity* of structure that a speaker or hearer is supposed to build in order to understand the relations among elements in a normally complex utterance: for instance, a fairly standard transformational-generative grammar might generate no less than thirteen different syntactic structures for an ordinary sentence like "Haven't you got one that's easier to hold than this one?", to say nothing of the phonological and semantic structures that are needed. Even if there was time to build up all these structures in ones' mind at an utterance being formulated, it seems unlikely that such a cumbersome system would have been adopted as a result of whatever pressures have led to human language being as it is. It is true, of course, that we seem to be talking about performance (how language is used) rather than competence (what people have to know in order to speak their language), and that Chomsky's grammars are claimed to be about competence, not perfor-

mance. However, this is not really the point, since the competence in question includes what speakers know intuitively about utterances, and Chomsky's claim is quite clearly that the only way in which we can represent relations among elements in an utterance is by building up a complete derivation for that utterance. In other words, for "Haven't you got one that's easier to hold than this one?" the speaker and hearer have in some sense to internalize thirteen or so separate structures—otherwise they will not have understood, for example, that the word *that* represents the object of *hold*, that this one is the subject of *than this one is (easy to hold)*, and so on. If this is not in fact the way in which we represent these relations to ourselves, then the structures generated by a transformational-generative grammar have no psychological reality at all, and lose their interest for anyone who wants to know about the human mind.

There is also the problem of the *contents* of the structures included in a derivation, especially of the more abstract structures. One standard analysis for "Cows moo" gives it a deep structure containing an element "present" (for the present-tense affix) just to the left of the verb: "cow-plural-present-moo" (with appropriate bracketing and labelling of course). But it seems unlikely that this element "present", or any neural equivalent, is represented by real-life speakers to the *left* of the verb—if indeed it is represented, as an affix, anywhere at all. Not even the most determined transformationalist would claim that his intuitions support this analysis—or, probably any other deep-structure analysis, yet deep structures are meant to have psychological reality. This is not to deny the psychological reality of the *phenomena* which deep structures (and semantic structures, where these are separate from deep structures) are intended to reflect—such as the fact that present tense is always realized in whichever verb is the first in a clause. What is in question is the way in which these phenomena are reflected in a transformational-generative grammar—the complaint being that the structures generated by such grammars have a great deal to do with the initial metatheorical assumptions about the nature of syntactic structures, and very little with the way in which speakers actually represent the structural patterns of utterances to themselves.

Very many technical problems in writing transformational-gen-

erative grammars arise from the metatheoretical assumptions, involving difficulties in capturing certain types of generalization which need to be captured, and stating certain types of restriction which need to be stated. These problems are likely to impress professional linguists more than non-linguists, and it is in any case difficult to compress linguistic arguments in an academically respectable way since each conclusion tends to lie at the end of a long chain of arguments. Every linguist is aware of at least some of these problems, and several leading transformational linguists have recently become so acutely conscious of them as to suggest that transformational-generative grammar is based on fundamentally wrong premises (e.g., George Lakoff, Paul Postal and Paul Schachter, among others). The following limited examples will give an idea of the kind of problem that arises, but further discussion of these and many other problems can be found in my *Arguments for a Non-transformational Grammar* (University of Chicago Press, 1976).

As examples let us consider English constructions in which there is a verb ending in *-ing* acting as a " gerund " (that is, a verb behaving in some sense as a noun), as in "I resent having to work in the evening". In such sentences there are good reasons for identifying the ing-form verb plus the elements belonging to it (in this case, *having to work in the evening*) as a sentence, which allows the same range of alternative structures as other sentences, including various kinds of objects, passives, negatives and so on, except that a subject is optional (compare "I resent her having to work in the evening") and the first verb has to be an ing-form (*having, taking* and so on). On the other hand, there are also good reasons for classifying such constructions as noun-phrases, like *my brother* or *the people who live here*, since they can occur in precisely the same range of environments, and there is even the similarity that an initial nounphrase can be marked as a possessive, with *'s* (compare "John's having to work in the evening" with "John's hat". The normal transformational treatment of such constructions treats them as being sentences *within* noun-phrases (which have no other parts), but for a number of reasons this analysis fails—not least because it is quite unmotivated, except for the need to treat these things as both sentences and noun-phrases Instead of this type analysis what is

needed is a simple one in which *having to work in the evening* and other comparable items are classified as being equally sentences and noun-phrases—but this analysis is quite impossible in the transformational-generative framework, because of the initial assumption that only one label must be assigned to each node.

Another problem is that there is no way, within the transformational-generative framework, of preventing the verb in such constructions from being analyzed either as "present" or as "past" in deep structure—the standard analysis allows either "past" or "present" to appear in the deep structure of these constructions, indeed cannot prevent it from doing so, given the assumption that deep structures are generated by phrase-structures rules. It is easy to see that this is undesirable: for semantic reasons, the tense needed after present tense *resent* would have to be "present", not "past", whereas after *remember* it would have to be "past", not "present". Indeed, the best solution would be to prevent *any* selection of tense from being made in such structures—to treat them, in other words, as "tenseless"—but this cannot be done either, given the types of rule found in a transformational-generative grammar.

To summarize these criticisms, it seems that transformational-generative grammar cannot be the right metatheory, since on the one hand the structural analyses it generates are psychologically implausible, and on the other hand there are fundamental technical problems in writing effective transformational-generative grammars. So an alternative metatheory has been developed which no doubt has problems of its own, but does not share these particular ones.

A NON-TRANSFORMATIONAL ALTERNATIVE

I should emphasize that the theory to be outlined is not the only alternative which has been put foward: in Europe, there is dependency theory (going back to the work of the French linguist Lucien Tesnière in the 1930s), in Britain there is systemic grammar (developed by Michael Halliday during the 1960s), and in the United States there are theories of Sydney Lamb (stratificational grammar) and Kenneth Pike (tagmemics). All these theories make initial assumptions about the nature of

structural analyses that differ from those of Chomsky. However, like all theories these have both weak and strong points, and I believe that the one expounded below combines the good points of several (notably systematic grammar and dependency theory) and avoids most or all of their weakness. For lack of a better alternative name, I have called this theory "daughter-dependency theory", which can be conveniently abbreviated to DDT. What follows will, I hope, give a general idea of how the theory works; but the reader is referred for a proper account to *Arguments for a Non-transformational Grammar*.

In contrast with Chomsky's assumptions about the nature of structural analyses, I have assumed that structures can be as complex as is necessary in order to show all the grammatical relations that need to be shown, in a single structure, whether they involve discontinuities, double-motherhood, rich classification of larger items, or whatever. In fact it turns out that, at least for English the amout of complexity needed is not really very great, since there are severe restrictions on the circumstances under which discontinuity and double-motherhood can occur, and the classification needed for clauses, phrases and words is also quite limited. Bearing in mind the types of relation that need to be shown, we can develop the metatheory to allow rules of just the right type for generating structures of the kind we need. Hardly any of Chomsky's rule-types appear in the list—there are no phrase-structure rules and no transformational rules, and the only types of rule found in both theories are those which Chomsky calls "subcategorization rules", which play a minor role in his grammars but a central role in DDT grammars.

All these differences arise from the fact that the metatheory in DDT has been repeatedly revised in order to make sure that the rules generate a structure for each sentence which incorporates *all* the relations among its constituents that need to be shown so that none are left over to be incorporated into some additional, more abstract structure, as in transformational-generative grammar. Thus, a single structure has to show both relatively concrete relations, notably sequence, and also relatively abstract ones, such as class-membership and dependencies between elements. An example is shown in Fig. 2, which is a

(simplified) DDT structure for the sentence that is represented transformationally in Fig. 1.

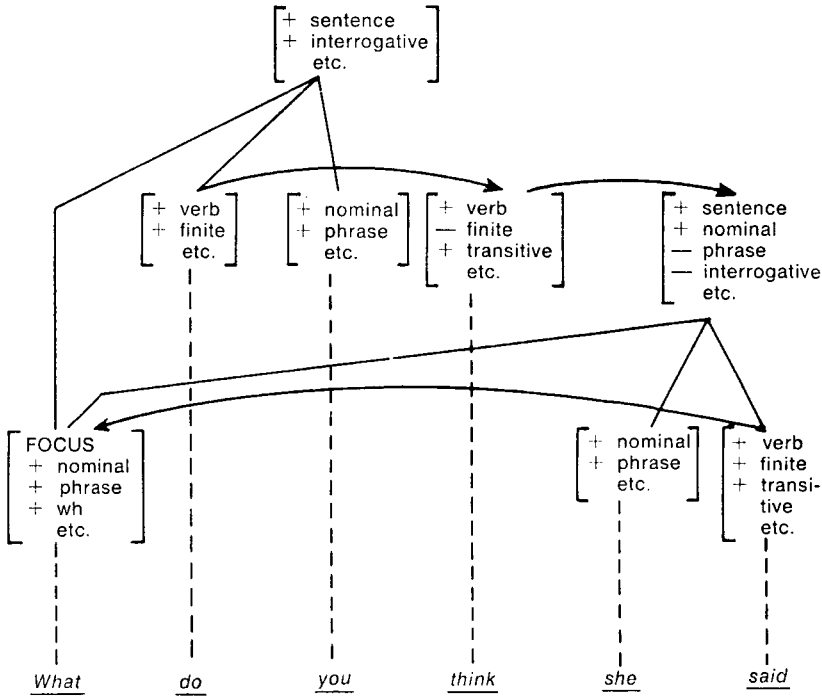

Fig. 2

The lines represent dependency relations—for example, the line linking *think* to *she said* (more precisely, linking the bundles of features that represent those items) shows that the latter depends on the former, as an object depending on a transitive verb; and the one linking *do* to the bundle of features at the top shows a part/whole dependency: i.e. the part (*do*) depends on the whole, in the sense that *do* can occur only as a "daughter" of a sentence. The diagram shows the discontinuity of the embedded clause and the double-motherhood of *what* directly, by means of the dependency lines connecting *what* to the rest of the

https://doi.org/10.1177/039219217602409605 Published online by Cambridge University Press

sentence; and it therefore makes it quite unnecessary to generate any other structure for this sentence. All that needs to be shown is already shown in this single structure, without recourse to deep structures and transformation. Thus DDT structures seem to be much more plausible candidates for psychological reality than transformational derivations, since it is easy to imagine speakers building up structures of this degree of complexity as they interpret "What do you think that she said?" in their minds. All that the speaker has to do is to identify the words, classify them, identify the larger units (phrases and clauses), classify them, and work out what depends on what—which is no more difficult than finding one's way round a department store, one would think.

Having reduced the number of *syntactic* structures to one, the question arises whether the other aspects of a sentence's linguistic patterning—its phonological and semantic aspect —could not perhaps also be built into this same single structure, so that the speaker's task in processing a sentence is reduced to reconstructing a single rather complex structure, rather than a different one for each of the three levels as in all other models with which I am familiar. I believe that this is a real possibility, but at present it is just a conjecture.

If this is the DDT position on the nature of structural analyses (link 3), what about the rules needed for generating such structures (link 4)? These rules are quite different from those found in a transformational-generative grammar, since there are no phrase-structure or transformational rules. Instead, there are the following types of rules, each internally quite simple: (a) rules for defining the total range of possible grammatical classes, in terms of which not only words but also clauses and phrases can be classified; (b) a lexicon, which lists all the words duly classified in terms of the classes defined in (a); (c) rules for defining possible dependency relations between parts and wholes (e.g. a finite verb can only occur as part of a sentence) or between parts (e.g. an object noun-phrase can only occur with a transitive verb); (d) rules for sequence, stating the order in which elements must occur, if in the same structure; (e) rules for adding various extra bits of information, such as the occasional functional label (e.g. FOCUS over *what* in Fig. 2).

The main difference between these rules and those of a transformational-generative grammar is that the latter's phrase-structure rules combine in a single rule-type two functions which are separated in a DDT grammar: namely defining dependency relations between part and whole, and defining sequence relations between the parts. For instance, the phrase-structure rule "NP → Det + N" says on the one hand that determiners (Det) and nouns (N) can occur only as daughters of noun-phrases (NP), and on the other hand that they must occur in the order determiner-followed-by-noun. By separating these two types of information from each other, it is possible to increase greatly the flexibility of the grammar without having to introduce rules with the power of transformational rules—so that the rules only introduce an element (by dependency) if it is overtly present in the "surface structure", and only put one element in front of another (by sequence rule) if they are going to stay that way. Consequently, there are no rules for deleting abstract underlying elements, and none for changing the order of elements.

This leads to the final link in the chain, that of the metatheory. Assuming that the rule types we have just listed are adequate not only for English (the main language on which I have worked) but also for other languages, we can build up a metatheory in which there are quite precise details about the types of rules that people internalize when learning their own language. I do not know whether the rule-types will turn out to be *specific* to language, but I suspect that many of them at least represent the sort of mental operations that are needed in other activities of life, such as classification and grouping items into larger units. Further work will no doubt show flaws in the present concept of DDT—indeed, this has already happened—but I believe that in general outline this theory gives a more reliable window into the human mind than transformational-generative grammar.