AN INTRODUCTION TO TRANSFORMATIONAL SYNTAX

Roger Fowler

Roger Fowler

An introduction to **Transformational Syntax**

Routledge & Kegan Paul
London and Henley

An introduction to

Transformational Syntax

Preface

This is a textbook on transformational-generative syntax, a mode of grammatical description proposed by the American linguist Noam Chomsky in a little book called *Syntactic Structures* published in 1957. Since that time, TG, as it has come to be called, has undergone massive growth and change. In the first place, its general framework has become accepted by the majority of Western linguists as providing the most reliable and revealing version of linguistic analysis. This fact has to be acknowledged despite intractable opposition from a few representatives of older schools of linguistics or of more insular traditions; and despite many disagreements about details of the proposed analysis. Second, TG has benefited from very substantial and useful revisions over the years.

The net consequence of these developments is that, although there is an increasingly wide demand for information on TG, current writings in the field are forbiddingly specialized and somewhat disputatious, and the older books have become somewhat out of date. Available elementary textbooks fall into two categories: there are those which were published in the early and mid-1960s, and project a version of TG which is not entirely consonant with more recent statements of the approach; and there are newer books – an increasing number – which embody fragments of contemporary revisions in what is sometimes a puzzling way. I have attempted to provide a 'compromise' account. The primary intention is to describe a transformational model of syntax which is more up to date than the classic textbooks, based as they are on *Syntactic Structures*, can provide. In essence, this means incorporating the general changes announced in Chomsky's *Aspects of the Theory of Syntax* (1965) and foreshadowed in Katz and Postal's *Integrated Theory of Linguistic*

Descriptions (1964). At the same time, I have tried to avoid making *Syntactic Structures* 'unreadable' through uncompromising paraphrase of *Aspects*: I assume that any student who uses the present book as a way of gaining access to contemporary syntactic theory will be interested enough to read *Syntactic Structures*, which, though now superseded in many respects, remains the most succinct, powerful and attractive argument for a transformational approach to syntax.

Readers who are familiar with the history of TG will soon realize that the present book is not a faithful paraphrase of *Aspects*. I would claim that it builds on the basic framework of that account – although even that claim may be controversial. Prematurely, I feel, the whole position of *Aspects* is under attack from some quarters. What I have tried to do is tidy up such contradictions and omissions as appear in *Aspects* without, in my opinion, invalidating the overall position. In an attempt to reflect contemporary work I have gone beyond the letter of *Aspects* (hence my reading of it could be called inaccurate) in several respects: a more extensive use of feature analysis in syntax, and, in particular, a new treatment of *Det* and *Aux* which is not envisaged in *Aspects*. Beyond this up-dating of *Aspects*, I have tried to indicate directions of subsequent enquiry by other grammarians: for this reason, my treatment of pronouns, relative and appositive clauses, and conjoining has been worded in a tentative and open-ended way – these are current preoccupations in syntactic research and I want to suggest that further rethinking in these areas may bring important and radical revisions to the very basis of the grammar.

The grammar presented here, then, is by no stretch of the imagination 'final': it is a provisional grammar designed to help students read both classic and contemporary writings in TG. Certain obvious limitations of the present model of syntax make it clear that it is quite provisional: I would point especially to the difficulty of explaining adverbials and phrasal conjunction in this version of transformational syntax. A student who realizes just what these particular difficulties are will be well equipped to evaluate both older and newer solutions to such problems.

A word on how this book is to be read. It is a textbook and an instrument, to be used rather than consulted. The material is presented sequentially, with modifications *en route*, and so it should be read slowly, from beginning to end. It is not a reference book, and, to discourage its use as such, the index is minimal. I assume that the book will normally be used in a taught course, in conjunction with other reading materials. At the end of the book there is a short reading list: this reflects a range of important books and articles

which any student who has mastered this book might be expected to have read by the time he finishes his course. The teacher of a course on TG which uses this book will certainly want to assign particular readings chapter by chapter, and I would not wish to dictate what these should be, so have not made many specific recommendations. Likewise, the topics for discussion and exercises at the end of each chapter are, by and large, posed rather neutrally: most of them can be attempted using analytic terms which are not precisely those offered in this book. They can be supplemented by more detailed analytic tasks slanted towards different techniques of analysis.

The book contains no footnotes, because I think it will be easier to read without continual qualifications and references. However, I would not wish to conceal the fact that the materials are not on the whole original, but drawn eclectically from a wide range of authors within the general TG framework. The responsibility for their presentation, and certainly for the overall content of the book, is mine alone.

Of the many people who have helped me determine the content and organization of this book, I would particularly like to thank several generations of students of linguistics at the University of California (Berkeley) and the University of East Anglia, who have had this material tested on them in oral form and who, directly and indirectly, have been responsible for modifications too numerous to itemize but nevertheless quite invaluable.

Contents

one

What is a Grammar?

The word 'grammar' in present-day linguistics has at least two important meanings. On the one hand, we say that a speaker knows the grammar of his language. He usually does not know it consciously – unless he has some special training in linguistics, he cannot talk confidently about the nature of his grammar. A grammar in this first sense comprises the linguistic knowledge speakers possess which enables them to communicate in their language. 'Grammar' here is a psychological, mentalistic, concept. The second sense relates to the linguist, not to the speaker: the linguist is said to write a grammar of the language. This grammar is a formal, explicit, description of the language.

Now these two usages must be kept apart. One look at a printed grammar is enough to convince us that it is extremely unlikely that the speaker knows *his* grammar as an object of the shape the linguist provides when he writes *his* grammar. If we could magically 'tap' the speaker's hidden linguistic knowledge – by hypnosis, drugs or whatever other implausible technique – so that he could tell us directly what it is that he knows which we refer to as 'his grammar', he would not simply dictate Jespersen's *Modern English Grammar* or Chomsky's *Syntactic Structures* to us. The speaker does not store his linguistic knowledge in the format which the linguist adopts for explanatory purposes; nor, when he produces sentences, does he follow step-by-step the processes which the linguist spells out as he constructs derivations for sentences. This latter point is most important, and I will return to it: a linguist's grammar *generates* sentences; a speaker *produces* (and understands) sentences – the two processes are quite independent.

Although the two senses of 'grammar' must be dissociated, we can

learn a lot about how to write a grammar, and what to put in it, by speculating on the nature of the grammatical knowledge of speakers. We can profitably ask: what must a speaker-hearer know in order to communicate in his language? If we observe linguistic behaviour from a number of angles, we can begin to make observations which encourage us to predict certain necessary components of grammatical knowledge. First, native speakers know that, of the following three sentences, (1) is not a sentence of English, (2) is an ungrammatical sentence of English, (3) is a grammatical sentence of English:

(1) Quel est l'objet à la fois intégral et concret de la linguistique?
(2) Three tons are weighed by this truck.
(3) This truck weighs three tons.

To go into more detail, they know more about ungrammatical sentences; for example, that (4), (5), (6) and (7) are progressively more deviant:

(4) This circle is square.
(5) John alarmed an apple.
(6) John alarmed a the.
(7) Alarmed a the John.

More relevantly, perhaps, they know an enormous amount about grammatical sentences of English. For example, they know that (8) and (9) are similar in meaning, as are (10), (11) and (12) and, in a different way, (13) and (14):

(8) Her frankness astonished him.
(9) He was astonished by her frankness.
(10) The carpet was brown.
(11) The brown carpet . . .
(12) The carpet which was brown . . .
(13) He mounted his proud horse.
(14) He mounted his proud steed.

It goes without saying, of course, that speakers know which sentences are different, as well as which ones are alike. That is, they can tell sentences apart. This observation needs no illustration at this point, since the book as a whole is a discourse upon it.

Another area of linguistic knowledge concerns ambiguous sentences. Consider the following two examples:

(15) The chicken is ready to eat.
(16) I saw her in the street.

(15) can be associated with either 'X eats the chicken' or 'the chicken

eats X'. (16) means either 'I saw her when I was in the street' or 'I saw her when she was in the street'. A mature speaker of English knows enough about the structure of (15) and (16) to retrieve either (or, as alternatives, both) of the meanings for each of these sentences.

The linguist attempts to find a way of explaining these facts about speaker-hearers' linguistic capacities. He has to account for the structure of English sentences in a way which takes cognizance of speakers' intuitions of deviance, similarity, distinctness and ambiguity in their experience of English sentences. For instance, no analysis of (15) is adequate unless it assigns two alternative structural descriptions to that sentence, in recognition of the fact that speakers attach two different meanings to it. In this case, the grammarian will probably say that *the chicken* is the Object of the verb in one interpretation ('X eats the chicken'), the Subject of the sentence in the other ('The chicken eats X'). 'Subject' and 'Object' are descriptive concepts which the linguist proposes as a way of explaining certain structural facts about English. Notice that, while the motivation for these concepts comes from an enquiry into 'what the speaker knows' – here, the speaker's perception of ambiguity – they are no more than theoretical terms, aids to expressing a hypothesis about linguistic knowledge. It is not necessary to assume that English speakers' brains contain two compartments labelled 'Subject' and 'Object'.

A linguist writes a grammar in an attempt to expose the structure of the sentences of a language. His structural analysis is well-motivated to the extent that he bears in mind that this set of sentences relates to a shared linguistic competence in speakers of the language under description. The problem 'What do speakers know?' has an immense bearing on our more directly relevant question 'How shall I present the structure of the sentences by which speakers communicate?'

Briefly, a language *L* is a set of sentences. The linguist must account for all and only the grammatical sentences of *L*. ('*L*' is a standard abbreviation for 'any natural language'.) This obligation follows from my comments on sentences (1)–(3) above: the mature speaker-hearer can distinguish between grammatical sentences of *L*, ungrammatical sentences of *L*, and sentences which are not of *L*. If the set described did not have limits, the grammar produced would be utterly unprincipled: it would fail to divide off English from French sentences, and, since it would omit to separate off ungrammatical and grammatical sentences of *L*, it would be structurally anarchic. I will assume that we have procedures for discounting sentences which are not of *L* and sentences which are not grammatical sentences of *L*. (Actually, these procedures are not yet properly

established, but the problems are too complex to be discussed here.) If we can thus recognize grammatical sentences of *L*, we must go on to ask 'How many of them are there?' The answer to this question is known: the set *L* contains an infinite number of grammatical sentences. Almost every sentence we hear, or produce, is new to us. We are not normally conscious of the inventiveness of natural language; do not realize that few of our sentences are exact repetitions of already-used utterances. Of course, every society has a stock of routine utterances like 'Good morning', 'Dear Sir', 'Thank you', 'No Smoking', 'I love you', 'Any other business?' and so on. These utterances, which are frequently used, invariant, and tied to ritualized communication situations, are quite untypical of normal linguistic performance, which is diversified apparently without limit. One might object that this observation is either unprovable, or, if provable, irrelevant, since, because of human mortality, we cannot actually experience an infinite set of sentences. However, we need this assumption, because we must account for the creativity of language – we are interested in the newness of sentences, even if we cannot be concerned with their infiniteness. And there is, as it happens, a demonstration of the notion 'infinite set of sentences' which is not vulnerable to the embarrassing death of the grammarian before he finishes counting sentences. What we can show is that there is *no longest sentence* in a natural language, and therefore by implication that there are an infinite number of sentences. (This is not to say that there can be a sentence of infinite length, as has sometimes been claimed, quite erroneously.) For every sentence of the type (17), a longer sentence (18) is possible:

(17) John eats meat and vegetables.
(18) John eats meat, vegetables and fruit.

And for every sentence (18), a longer sentence can be constructed by adding one more item. I will give two more examples of constructions with this property; there are in fact several syntactic devices available for extending sentences indefinitely:

(19) John believed that Mary claimed that Peter maintained that Clive said that . . .
(20) This hot, sunny, lazy, predictable . . . climate suits me very well.

As the sentences of a language are infinite in number, the set which the linguist must describe cannot be coextensive with any finite corpus of sentences which, by observation and recording, he might collect. There is a second reason why the task of writing a grammar cannot be accomplished by merely cataloguing the structures found in an observed corpus of sentences. The fact is that the actual utter-

ances of speakers do not adequately reflect speakers' real competence in *L*. Actual speech, as any unprejudiced observation will confirm, is riddled with grammatical mistakes of all kinds: incomplete sentences, false concords, arbitrary changes of structure in mid-sentence, illicit conjoinings of constituents which ought not to be linked together – or at least not in the manner that they are – and so on. (I am not appealing to 'prescriptive' standards. By 'ungrammatical' here I don't mean structures which, in the manner of the eighteenth-century purifiers or the edicts of the French Academy, have been decreed to be unacceptable; but structures which native speakers, if they could be reliably consulted, would agree are ill-formed from the standpoint of their grammatical knowledge.) These errors stem from various kinds of psychological and situational 'interference': distraction, lapses of memory, shifts of attention, hesitation, etc. To describe such deviant sentences as these which occur in a corpus would be to describe linguistically irrelevant psychological factors as well as the linguistically relevant structural knowledge of speakers.

Thus a corpus of utterances is not the true subject-matter of linguistic description: it is only *data*—a set of observations from which, with caution, the linguist must draw his grammatical statements. In view of what has just been said, it is clear that the linguist's use of his primary data must involve two adaptations. (First, some 'idealization' is necessary so that the grammar does not take account of the deviant sentences which occur in the corpus. (Second) the linguist must devise rules which project from his finite, observed materials to an infinite set of sentences. That is to say, the grammar must have *predictive* power.

All this adds up to the fact that a grammar is not a simple reflection of linguistic usage. A few years ago, linguists used to be attacked, for instance in the editorials of educational journals, for abandoning all standards and saying that 'anything goes': in fact, linguists until quite recently believed that any sentence which was produced ought to be described by a grammar. But now a major reorientation has taken place – it has been realized that speakers' actual linguistic performance is not a very accurate indication of their underlying linguistic competence. Many features of linguistic performance, many aspects of texts and utterances, have to be discounted when writing a grammar. At this stage, I might mention just one other characteristic of discourse which a grammar does not seek to represent. It is well known that some words, and some constructions, occur more frequently than others: e.g. words like *the* and *and* are much more frequent than *discourse* or *dog*; complex sentences more frequent than simple sentences. Furthermore, the types of sentences which

occur in discourse correlate broadly with the circumstances in which discourse is used – there are typical styles for advertising, informal conversation, political rhetoric, scientific writing, etc. But as far as the grammar is concerned, no one sentence, or type of sentence, is more predictable than any other. Grammar does not take account of probabilities. If a sentence occurs in a text or discourse, the grammar will describe its structure; it will not explain why that sentence rather than some other was selected. The explanation of *why* sentences occur in discourse is the task of stylistics and sociolinguistics, not of grammar.

A grammar which meets the requirements outlined above is called a *generative* grammar. Such a grammar is predictive or projective in the sense that, given a finite body of data (including a collection of observed sentences), it offers a system of rules so framed as to account for an infinite set of potential sentences. In this way a grammar 'generates' or 'enumerates' or 'describes' or 'defines' the set of sentences which makes up the language. In an explicit and formal manner, the grammar assigns at least one structural description to each sentence in the language (allowing that many sentences are ambiguous and must therefore receive two or more structural descriptions). We can test individual sentences – 'Is this sentence generated by the grammar of English?' – by retracing a formal derivation: by working through a series of rule-applications by which the sentence is derived. (For the notion of 'derivation', see below, pp. 45–7.) A generative grammar allows each structural description to be associated unambiguously with one derivation. Remember that a derivation is *not* an account of how a speaker produces a sentence. As we will see when we have looked at some derivations, such a proposal would be completely nonsensical. Early critics of transformational-generative grammar believed, quite mistakenly, that 'generate' meant 'produce' – that such a grammar focused on the speaker's end of the communicative process. Actually, a generative grammar is quite neutral with respect to speaker or hearer: it makes no claims to explain how a sentence in actual linguistic performance is either produced or comprehended.

One further clarification of terminology is necessary. A generative grammar does not have to be a transformational grammar. 'Transformation' refers to a particular kind of rule, and a generative grammar may or may not utilize transformational rules. In practice, most modern generative grammars happen also to be transformational. But in principle a generative grammar without transformational rules could be written. We may note also that transformations are not restricted to syntax: there are transformational rules in phonology, also. Note that the present book is about transforma-

tional *syntax*, and this is not the same as transformational *grammar* because grammar includes more than syntax.

I have said that a generative grammar 'assigns structural descriptions to sentences' and that in this way the linguist accounts for their structure in a manner which is consistent with what he can deduce about speaker–hearers' linguistic knowledge. We must now ask 'What do structural descriptions [*SD*s] tell us about sentences?' Given any one of the infinite set of sentences of *L*, all fully competent speakers of that language will agree, within reasonable limits, on its meaning. Equally, discounting peculiarities of accent and personal voice quality, speakers agree on what it sounds like. To put it another way, speakers are able to correctly associate a *semantic interpretation* with a *phonetic representation* for each of an infinite set of sentences of *L*. It would seem reasonable to expect a structural description to reveal those qualities which speakers attribute to sentences as they achieve sound-meaning associations. Let us consider a simple sentence:

(21) The cat sat on the mat

is readily interpretable somewhat as follows: it concerns a cat (known to be a certain kind of animal), particularized as one cat (rather than as more than one) and as a specific cat rather than any old cat (*the*, not *a*); identifiable behaviour (sitting) is attributed to the cat; a location is specified; this location is identified as a particular kind of inanimate object; the position of the cat relative to this object is given ('on'); the whole semantic complex – cat-sitting-location-mat – is set in past time. All this is roughly what the sentence means, to any unprejudiced English speaker. He possesses conventions for constructing this meaning, and he is also able to give these conventions realization in sound or script. These conventions of meaning and sound are community property: for every sentence, all speakers in the community agree on the mechanisms by which meanings are built up and associated with sounds.

Generative linguists, like traditional grammarians in general, deal with these facts by setting up three interrelated levels of description: a *semantic* level, a *syntactic* level and a *phonological* level. Alternatively, we could say that a grammar has three 'components', calling the components by their traditional names. (Note that 'grammar' is often used as an equivalent to 'syntax'; but our usage of the term 'grammar' is, perhaps *un*traditionally, more inclusive.) The *semantic* component is responsible, first, for assigning meanings to lexical items: it must incorporate a dictionary. Like ordinary dictionaries, this one must attempt to distinguish each lexical item from all others, by stating exactly what senses mature speakers attribute to

each item in the language's vocabulary. It must also try to set out the *structure* of the lexicon: the semantic relations (synonymy, antonymy, superordination, etc.) which exist between lexical items. A thesaurus aims to show these relationships, but conventional dictionaries do not usually attempt to define such relationships systematically. Second, the semantic component of a grammar should account for the fact that the meanings of individual words are, in sentences, amalgamated so that more complex meanings are formed. Since these 'larger' meanings are built up under syntactic constraints, the semantic component has to be arranged so that it can make reference to appropriate syntactic properties of sentences. The general design of the *syntactic* component of a grammar will be indicated in some detail in the next chapter. To put it in rather impressionistic terms, syntax distributes lexical items – and non-lexical formatives – in patterns, patterns which are spread out 'left-to-right' in time or space. Syntax lays the basis for translating an abstract meaning-complex into a piece of sequential behaviour. It does so by generating a linear string of words arranged in a regular pattern. This string constitutes the input to the *phonological* component of the grammar. For every word, and every string of words, there is an agreed realization in sound, a phonetic shape. The phonological component specifies what phonetic contour is to be attached to each of the infinite number of strings of words that the semantic and syntactic parts of the grammar produce between them. It is a set of instructions for pronunciation. Since many languages use a written, as well as spoken, medium, there is also a *graphological* equivalent to the phonological section of the grammar.

This book is about syntax. It will therefore have little to say about the details of phonological and semantic structure. But we must remember that the three components interlock, that none of them functions independently of the others. I have already mentioned, for example, that semantics must make reference to syntax to guide the formation of sentence-meanings out of the sub-sentence elements provided by the dictionary. Likewise, the phonological component cannot work unless it has a very precise analysis of the syntactic structure of the sentences for which it has to design a phonetic representation. Syntax is very definitely not autonomous, and so during the course of this book I will do my best to clarify the points at which it makes contact with the other components.

Exercises and topics for discussion

1. Discuss the distinction between 'prescriptive' and 'descriptive' grammar.

2. In Ch. 1, some of the reasons why a grammar cannot be simply a description of a finite corpus were given. Work out the arguments against 'corpus-bound' grammar in more detail.
3. Investigate the notion of 'grammaticalness' in the writings of some modern transformational grammarians.
4. Make a critical review of Noam Chomsky's expositions of the distinction between 'linguistic competence' and 'linguistic performance'. You may wish to consider also Ferdinand de Saussure's distinction between *langue* and *parole*.

two

Deep and Surface Structure

As it happens, the most important relationship between syntax, semantics and phonology can be presented straight away. The linguistic levels of 'meaning' and 'sound' both have to be invoked to help us define the central distinction in syntax – and with it the idea of a transformational grammar itself. Consider the following sentences:

(22) He took off his hat.
(23) He took his hat off.

These sentences have the same meaning; but they are different arrangements of words. Since the difference between (22) and (23) is immediately apparent at first glance, 'on the surface', as it were, let us say that (22) and (23) exhibit different *surface structures* (or *superficial structures*). To continue the metaphor, we will explain the synonymy of (22) and (23) by saying that they have the same *deep structure* (or *underlying structure*). Deep structure relates to meaning; surface structure relates to order of elements, and hence to sound, for in effect the surface structure determines the sequence of sounds which occurs in a phonetic realization of a sentence. Surface structure is a dimension with physical associations, since it is the point at which a sentence impinges on space and time. Deep structure, however, is an abstraction, a complex of meanings which is 'unpronounceable' unless it is rendered as a surface structure. Before we attempt to say more about the theoretical status of deep and surface structures, let us look at some more examples.

(22) and (23) illustrate the situation in which one deep structure is realized as two different surface structures. Another type of example of this same relationship is (24), (25) and (26):