# 12 Generative Grammar

THOMAS WASOW

## 1 Introduction

### *1.1 "Grammar"*

To most people, the word "grammar" suggests a normative enterprise, dictating what constitutes correct language use. For example, many educated English speakers would identify at least five supposed grammatical "errors" in the following sentence:

(1) Hopefully, we will be able to easily figure out who to talk to.

Yet native speakers of American English also certainly recognize that (1) would be an entirely acceptable and natural sounding sentence in ordinary discourse. Indeed, the supposedly "correct" alternative (2) would be an awkward and affected way of expressing the thought.

(2) I hope that we shall be able easily to figure out to whom to talk.

Modern linguistics has little use for this prescriptive conception of grammar. Linguists are more interested in the knowledge of English that allows native speakers to judge (1) as fully acceptable and (2) as somewhat less natural. The prescriptions of traditional grammar are attempts to impose the speech patterns of one region, class, ethnicity, or generation on speakers belonging to other groups. They may be of interest to sociologists, historians, and political scientists, but they tell us very little about the nature of language.

Language is a natural phenomenon, constituting an essential component of every human society. Linguistics is concerned with studying languages and language in general, much as biology studies living species and life in general. From this scientific perspective, the norms of prescriptive grammar are to linguistics as the American Kennel Club's breed standards are to biology: arbitrary evaluative standards of no relevance to objective description.

Linguists use the term "grammar," then, to refer to structural properties of language that have evolved naturally and that native speakers of the language have mastered without explicit instruction. These are largely properties of languages that are not even mentioned in traditional grammars, though some are addressed in foreign language instruction. They include facts about word order, for example, that *we*, *will*, and *be* in (1) must appear in that order, or else the sentence becomes unacceptable. They also include facts about the proper forms of words in particular contexts, for example, that replacing *figure* in (1) with *figured*, *figures*, or *figuring* makes the sentence unacceptable. Put in more technical jargon, "grammar" is taken by linguists to encompass syntax and morphosyntax. The term may also be construed more broadly to include principles relating linguistic forms to the meanings they express (semantics) and / or the sound patterns of languages (phonology).

## *1.2 "Generative"*

The term "generative" is associated with the tradition of grammatical research initiated and inspired by the work of Noam Chomsky. This term is sometimes construed very narrowly to refer only to work directly derivative from Chomsky's. Here it will be used more broadly to refer to work generally within the Chomskyan tradition, irrespective of whether its formalism and terminology come directly from Chomsky.

Among Chomsky's most important insights is the observation (noted independently over a century earlier by the great German linguist Wilhelm von Humboldt) that there are infinitely many well-formed sentences in any natural language. This follows immediately from the fact that any limit one might impose on the length of sentences would be arbitrary: any supposedly longest English sentence S would be two words shorter than "I said S," which is surely well-formed if S is. On the other hand, a grammar, conceived of as a description of a language, should be finite.

How can we give a finite description of something infinite? Inspired by earlier work in mathematical logic and the foundations of computer science, Chomsky answered this question by proposing that we think of grammars as devices that put pieces of sentences together according to precise rules, thereby "generating" well-formed sentences. If some of the grammar rules can apply to their own outputs (in technical jargon, if some rules are "recursive"), then it is possible for finite grammars to generate infinite languages.

To illustrate this, consider the following very simple (nonlinguistic) example. The ordinary Arabic numeral system used to represent numbers has infinitely many well-formed expressions (one for each number) constructed out of ten symbols, namely, the digits "0" through "9." We can write a simple grammar for the numerals denoting positive integers with the following rules:

- Each of the digits 1, 2, 3, 4, 5, 6, 7, 8, and 9 is a numeral.
- If N is any numeral, then N0 is a numeral.
- If N is any numeral, then NN is a numeral.

One of many possible formalizations of this would be the following:

| | | |
|---|---|---|
| $N \to 1$ | $N \to 5$ | $N \to 9$ |
| $N \to 2$ | $N \to 6$ | $N \to N0$ |
| $N \to 3$ | $N \to 7$ | $N \to NN$ |
| $N \to 4$ | $N \to 8$ | |

Here N is the category of well-formed numerals, and the arrow can be interpreted to mean "may consist of." This little grammar generates the infinite "language" of numerals denoting positive integers, because it contains rules that are recursive (namely, the last two).

## 2 Tenets of Generative Grammar

Although the term "generative" orginally characterized a conception of grammars as such recursive rule systems, the term is now used somewhat more generally. In particular, what distinguishes work in generative grammar is the goal of describing languages systematically, as opposed to the more anecdotal approach of traditional grammars. While it is impossible to give a precise definition of generative grammar, there are several tenets shared by the vast majority of generative grammarians. These are summarized in the following subsections.

### 2.1 *Grammars should be descriptive, not prescriptive*

As discussed above, this proposition is generally accepted by modern linguists. Although it is not unique to generative grammarians, it is common to them.

### 2.2 *Grammars should characterize competence, not performance*

Despite its anti-prescriptivism, generative grammar is not an attempt to describe all or only the actual utterances of native speakers. This is implicit in the claim that languages are infinite: it would have been safe to assume that no sentence over one million words long will ever be uttered. But this upper bound exists because of limits on human memory and patience, not because of any linguistically interesting facts. Moreover, because of speech errors of various kinds, people frequently produce utterances that are not well-formed sentences, even by the judgments of the speakers. To distinguish between the idealized infinite languages that generative grammarians seek to describe and the far messier output of actual speakers, Chomsky introduced the terminology "competence" vs. "performance."

One common property of generative grammar in all its varieties is the focus on characterizing linguistic competence. Many generative grammarians would also like to develop models of linguistic performance, but most believe that a competence theory will be a necessary component of such a model. Put slightly differently, it is widely accepted that explaining how a language is actually used will require understanding speakers' knowledge of that language.

## 2.3 *Grammars should be fully explicit*

Traditional grammars presuppose some knowledge of the language under description and tend to focus on aspects of the language that are variable or have changed. Generative grammars are supposed to be precise rule systems that characterize the whole language, without relying on any prior knowledge of the language on the part of the reader. Many generative grammarians identify explicitness with formalization. Hence, the generative literature abounds with formalisms (though it is not always made clear how the formalisms are to be interpreted). Early work in generative grammar approached this goal of explicitness and formalization far more consistently than most recent work.

## 2.4 *Linguistic analyses should be maximally general*

If two grammars cover the same range of data, but one requires two distinct rules where the second has only one, generative grammarians take this as evidence for the superiority of the second grammar.

A famous example of this mode of reasoning is due to Postal (1964). He noted that what are called "tag questions" in English require a kind of matching between the tag and the initial portions of the main clause, as illustrated in (3). Following standard practice, asterisks are used to mark unacceptable strings.

(3) a. I have won, $\left\{\begin{array}{l}\text{haven't}\\ \text{*won't}\\ \text{*aren't}\end{array}\right\} \left\{\begin{array}{l}\text{I}\\ \text{*you}\\ \text{*we}\\ \text{*they}\end{array}\right\}$?

b. You will win, $\left\{\begin{array}{l}\text{won't}\\ \text{*haven't}\\ \text{*aren't}\end{array}\right\} \left\{\begin{array}{c}\text{you}\\ \text{*we}\\ \text{*they}\end{array}\right\}$?

Postal also observed that imperative sentences take only a restricted range of tags, though there is nothing overtly present in the initial portions of imperative sentences that the tags match.

(4) Close the door, $\left\{\begin{array}{l}\text{won't}\\ \text{*haven't}\\ \text{*aren't}\end{array}\right\} \left\{\begin{array}{l}\text{you}\\ \text{*I}\\ \text{*we}\\ \text{*they}\end{array}\right\}$!

If we analyze imperative sentences as having an initial *you will* at some level of analysis, he reasoned, we could use a simple rule to generate tag questions on both declarative and imperative sentences. Such an analysis is said to "capture a generalization" – in this case, the generalization that tags on imperatives and declaratives are fundamentally alike. The desire to capture generalizations plays a very important role in the argumentation of generative grammar.

## 2.5 The theory of grammar should make universal claims

To the extent possible, facts about individual languages should be derived from general principles that apply to all languages. Information stipulated in the grammars of particular languages should be kept to a minimum. This is motivated in part simply by standard scientific methodological considerations: more general hypotheses are both more parsimonious and more interesting than less general ones. But it is also motivated in part by psychological concerns – specifically, by Chomsky's "argument from the poverty of the stimulus," which will be discussed in the next subsection.

The focus on the development of a general theory of grammar – "universal grammar" (UG), as Chomsky dubbed it – is perhaps the most distinctive characteristic of the generative tradition. Although other linguistic traditions involve extensive cross-linguistic comparisons resulting in important hypotheses about universal properties of language (see Croft, Chapter 14, this volume, for a sample of such work), generative grammar approaches these issues in a distinctive way. Specifically, the universals of generative grammar tend to be formulated as rather abstract principles of grammatical organization that are not directly observable in the linguistic data. Rather, their discovery and testing typically involve a complex combination of empirical observations, methodological assumptions, and inferential processes. This is in sharp contrast with more observationally transparent universals like those of Greenberg (1963), and much subsequent work on language typology. Some examples of linguistic universals in the generative style will be provided in section 4 below.

## 2.6 Grammars should be psychologically relevant

Generative grammarians characteristically (but not universally – see, for example, Katz and Postal 1991) take their theories to be relevant to psychological questions. Chomsky has been particularly outspoken on this issue, asserting that "a particular generative grammar" is "a theory concerned with the state of the mind / brain of the person who knows a particular language" (Chomsky 1986: 3).

More specifically, Chomsky has argued that a rich theory of universal grammar is necessary to account for the possibility of language acquisition. The most striking fact about human languages, he claims, is the gulf between knowledge and experience, observing that the following question, formulated by Bertrand Russell, is particularly applicable in the domain of language:

> How comes it that human beings, whose contacts with the world are brief and personal and limited, are nevertheless able to know as much as they do?

The fact that every normal human masters a language with little apparent effort or explicit instruction suggests that humans are genetically endowed with a "mental organ" specifically adapted to acquire languages of a particular kind. This is known as the "argument from the poverty of the stimulus."

While Chomsky has emphasized the issue of learnability, others have argued that work in generative grammar is relevant to psychology in other ways. For example, Bresnan (1978) argued that a generative grammar should be an integral component of a theory of language use – that is of the mental processes involved in speaking and comprehension.

## 3 Common Formal Elements

Since Chomsky's seminar work in the 1950s, many different theories of grammar have been articulated that fit the general characterization in the preceding sections. Almost all can be viewed as extensions of what is known as "context-free (phrase structure) grammar" (CFG).

### *3.1 Context-free grammar*

CFG begins with the relatively uncontroversial assumption that words can be classified into categories, based on their morphological properties (that is, what changes in form they undergo through suffixation and the like), their distributional patterns (that is, what other words appear in their vicinity in sentences), and their meanings. The traditional categories of noun, verb, etc. (inherited from the grammatical studies of ancient Greece) are still quite generally employed, supplemented by a number of other categories, some of them idiosyncratic to particular theories.

A second generally accepted premise of CFG is that the words in sentences are grouped into phrases, which themselves are grouped together into larger phrases, and so on. It is common to represent the phrase structure of a sentence by means of a "tree diagram" like (5):

(5)

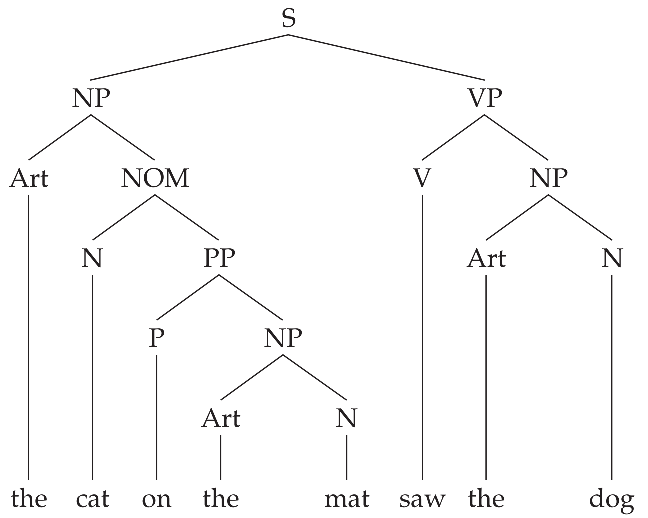


Phrases are identified by their distributional patterns and usually function as semantic units as well. Like words, phrases are generally classified into categories; the most widely used phrasal category labels – e.g., noun phrase (NP), verb phrase (VP), prepositional phrase (PP) – derive from the categories of words that appear in canonical instances of those phrases. These words are called the "lexical heads" (or sometimes just the "heads") of the phrases.

A CFG has two parts:

- A lexicon, consisting of a list of words, with their associated grammatical categories (referred to as "lexical categories").
- A set of rules of the form A → Φ where A is a phrasal category, and "Φ" stands for any string of lexical and / or phrasal categories. The arrow is to be interpreted as meaning, roughly, "may consist of." These rules are called "phrase structure rules."

The left-hand side of each rule specifies a phrase type (including the sentence as a type of phrase), and the right-hand side gives a possible pattern for that type of phrase. Because phrasal categories can appear on the right-hand sides of rules, it is possible to have phrases embedded within other phrases. In fact, some types of phrases (such as NPs and PPs) can be embedded in other phrases of the same type, giving CFGs the recursive character needed to generate infinite languages.

A CFG normally has one or more phrasal categories that are designated as "initial symbols." These are the types of phrases that can stand alone as sentences in the language. Most simple CFGs have just one initial symbol, namely "S." Any string of words that can be derived from one of the initial symbols by means of a sequence of applications of the rules of the grammar is generated by the grammar. The language a grammar generates is simply the collection of all of the sentences it generates.

In the 1950s and early 1960s, Chomsky, Postal, and others argued that simple CFGs lacked the descriptive power to account for all of the syntactic regularities of natural languages. Although some of those arguments have since been called into question, the conclusion remains generally accepted (see Savitch, et al. 1987 for a collection of relevant articles).

## *3.2 Transformational grammar*

Chomsky's earliest work suggests that the shortcomings of CFG could be remedied by associating with each sentence of a natural language, not just one tree but a sequence of trees. The initial tree in each sequence would be generated by a CFG (sometimes called the "base") and subsequent trees would be derived through a series of transformations – that is, rules that modified the trees in precisely specified ways.

This can be illustrated with the phenomena of tag questions and imperatives described above (see Baker's contribution to this volume for further illustrations). Space limitations require substantial simplifications: only non-negative sentences with pronouns as subjects and auxiliary verbs will be considered here. A simple transformational grammar for these phenomena might include the base grammar in (6) and the transformations in (7) and (8). Parentheses are used to indicate that an element is optional – for example, the fourth rule in (6) says a VP may consist of a verb, with or without a following NP. In (7) and (8), "$\Rightarrow$" means "may be transformed into."

(6) A lexicon for English, plus:
$S \rightarrow NP\ AUX\ VP$
$NP \rightarrow Pronoun$
$NP \rightarrow (Art)\ N$
$VP \rightarrow V\ (NP)$

(7) Tag formation transformation:

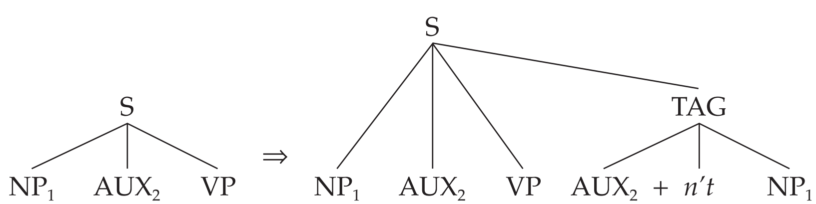


(8) Imperative transformation:

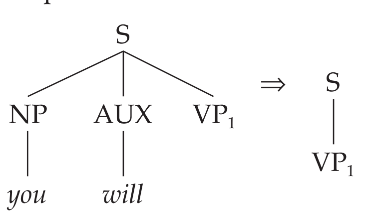

(7) takes as input trees for simple declarative sentences, and produces as outputs trees for the corresponding sentences with tags. It does this by copying the auxiliary verb, inserting a contracted negative, and copying the subject. (8) derives imperative sentences from declaratives starting with *you will* simply by deleting these two words (and the nodes right above them).

On this analysis, (4) is derived from the same base tree as *You will close the door*, by application of the two transformations, in the order given.

Early generative work was known as "transformational grammar," because the addition of transformations to CFG was seen as the crucial innovation. Throughout the history of generative grammar, transformational theories have had many advocates – always including Chomsky. Since the late 1970s, however, non-transformational alternatives have also been extensively developed.

## 3.3 Other enhancements to CFG

Several enhancements to simple CFG have been adopted in transformational and non-transformational generative theories alike. One of the earliest was the addition of a semantic component. It is evident that the acceptability of a sentence is influenced by what the intended meaning is, and it is often difficult to draw a sharp line between syntactic and semantic analyses. Consider, for example, the facts in (9).

(9) a. I excused $\left\{\begin{array}{l}\text{myself}\\ \text{*me}\end{array}\right\}$.

b. He excused himself. [*He* and *himself* must refer to the same person]

c. He excused him. [*He* and *him* must refer to different people].

The facts in (9a) are manifestly about the distribution of the words *myself* and *me*. The contrast between (9b) and (9c) is evidently a semantic one. Yet there is clearly a single generalization covering both contrasts, namely, that in the configuration $\text{NP}_1$-V-$\text{NP}_2$, $\text{NP}_2$ can be a reflexive pronoun (that is, a form ending in *-self* or *-selves*) just in case it refers to the same individual as $\text{NP}_1$. This generalization will be developed in more detail below. For now, it can serve as an illustration of the role of semantics in grammar.

Another enhancement of CFG that has been generally adopted in generative grammar is the use of non-atomic category labels for words and phrases. For example, in the mini-grammar presented in (6), AUX and V are distinct categories, with no more in common than, say, N and V. But what this grammar calls AUX has traditionally been treated as a species of verb. This makes sense when one considers the word *have*. In a sentence like *We have won*, (6) would treat it as an AUX (consider *We have won, haven't we?*); in a sentence like *We will have fun*, (6) must treat *have* as an instance of V (consider *We will have fun, won't we?*).

There are many more arguments for allowing words and phrases to be treated as the same in some respects but different in others. This is accomplished formally by replacing atomic category labels with more complex information structures. In particular, almost all varieties of generative grammar employ "feature structures" as category labels. Feature structures consist of pairings of features with values. A feature is simply a name for something used in classifying words or phrases; features are associated with multiple values, corresponding to properties of the words or phrases in question. For example, nouns can be subclassified into proper and common nouns, and into singular and plural nouns. Representing this with features would involve positing two features, say COMMON and NUMBER, each of which has two values (in English, at least). Then the two features could be used in representing the categories of some representative words as follows:

(10) $$\textit{child}\begin{bmatrix}\text{COMMON} & + \\ \text{NUMBER} & \text{sing}\end{bmatrix} \qquad \textit{London}\begin{bmatrix}\text{COMMON} & - \\ \text{NUMBER} & \text{sing}\end{bmatrix}$$

$$\textit{children}\begin{bmatrix}\text{COMMON} & + \\ \text{NUMBER} & \text{pl}\end{bmatrix} \qquad \textit{Alps}\begin{bmatrix}\text{COMMON} & - \\ \text{NUMBER} & \text{pl}\end{bmatrix}$$

All of the feature structures in (10) might also have something like [POS noun] (where "POS" is for "part of speech").

Treating categories as bundles of features makes it possible to represent large numbers of grammatical categories quite compactly, since every different combination of features and values is a different category. This allows grammarians to make fine distinctions, while still permitting reference to large classes of expressions. Some form of decomposition of categories into features has consequently been adopted in almost every variety of generative grammar. So long as there are only a finite number of features, each of which has only a finite number of possible values, this decomposition does not fundamentally alter the descriptive power of CFG. It does, however, make it possible to capture generalizations across categories of words and phrases, as well as characterizing categories at more or less fine-grained levels.

Some theories have taken this process one step further, however, allowing the values of features to be feature structures themselves. This constitutes a more fundamental enhancement of CFGs, allowing a great deal of information to be encoded into the representations of grammatical categories. As will become evident below, this increased descriptive power makes possible interesting alternatives to certain widely accepted transformational analyses.

One of the advantages of decomposing categories into features is that it permits efficient reference to classes of categories. For example, one can refer to all singular nouns with the feature specification

$$\begin{bmatrix}\text{POS} & \text{noun} \\ \text{NUMBER} & \text{sing}\end{bmatrix},$$

leaving other properties unspecified, including gender, case, and whether it is proper or common. This sort of "underspecification" is widely exploited in generative grammar.

One particularly influential case of underspecification is the suggestion by Chomsky (1970) that the phrase structure rules of languages could be reduced to a few very general schemas, with highly underspecified categories on both sides of the rules. This idea has been developed in many different ways, but has entered into most theories of generative grammar. In its simplest version, it holds that all phrases should be viewed as projections of lexical heads and that phrases uniformly have three levels: the lexical head, an intermediate level, and the full phrase. These are often designated as X, X′, and X″ (where X can stand for any combination of features). Then the principal phrase structure rules could be schematized as follows (where the superscripted asterisk is an abbreviation for zero or more occurrences of the immediately preceding symbol):

(11) $X'' \rightarrow (Y'')\ X' \qquad X' \rightarrow X\ Y''^{*}$

These rule schemas embody the claim that all phrases have the same basic structure, consisting of a lexical head, possibly followed by some other phrases (known as "complements") and possibly preceded by a single phrase (known as the "specifier"). Variants of this idea go under the label "X-bar theory." Although there are many different versions of X-bar theory, schematizing the phrase structure rules through underspecification of the categories is common to many generative theories.

The rule schemas in (11), as stated, do not appear to be good candidates for universal grammar, because they stipulate particular orders of elements. But there are languages (such as Japanese) in which lexical heads consistently come at the ends of phrases, and others (such as Irish) in which lexical heads come at the beginnings of phrases. It has been proposed (e.g., by Gazdar and Pullum (1981)) that the information about hierarchical structure and the information about left-to-right ordering of elements should be decoupled. That way, the schemas in (11) could be regarded as universal, up to the ordering of elements on the right-hand sides. This is another idea that has emerged in a number of different generative theories.

## 4 Some Phenomena Studied by Generative Grammarians

The literature of generative grammar is full of detailed examinations of myriad syntactic phenomena in a wide variety of languages. Most analyses depend on assumptions that are controversial. Nevertheless, the field has made numerous genuine discoveries. Although different schools of thought employ disparate formalisms and terminology, we know far more about the structure of language than we did in the 1950s, thanks to research in generative grammar. This section provides an overview of two areas in which generative grammarians have made clear progress.

## *4.1 Binding principles*

The examples in (9) above illustrate that English has two different types of pronouns, namely reflexives (*-self* / *-selves* forms) and non-reflexives. While *myself* and *me* both refer to the speaker (as does *I*), the environments in which they can be used differ. In particular, consider the following contrasts:

(12) a. *I support me.
b. I support myself.
c. They support me.
d. *They support myself.

(13) a. I don't expect them to support me.
b. *I don't expect them to support myself.
c. *They don't expect me to support me.
d. They don't expect me to support myself.

The following two generalizations (known as the "binding principles") roughly summarize the distributional difference between the two types of pronouns:

- A. A reflexive pronoun must have a local antecedent.
- B. A non-reflexive pronoun may not have a local antecedent.

For present purposes, "antecedent" can be taken to mean a preceding NP with the same reference. The term "local" is meant to convey the observation that the antecedent of a reflexive should not be too far away. However, giving a precise definition of "local" for these principles is not a trivial problem, as evidenced by examples like the following:

(14) a. The house has a fence around $\left\{\begin{matrix}\text{it} \\ \text{*itself}\end{matrix}\right\}$.

b. We wound the rope around $\left\{\begin{matrix}\text{it} \\ \text{itself}\end{matrix}\right\}$. [*it ≠ the rope*]

c. I wrapped the blanket around $\left\{\begin{matrix}\text{me} \\ \text{myself}\end{matrix}\right\}$.

These examples show that locality cannot be measured simply in terms of number of words or phrases intervening between the pronoun and its antecedent, for the three examples all have the same number of words and phrases.

There is a rich literature working out the details of the basic ideas in principles A and B above. These details need not concern us here. What is of interest is that English is by no means unique in having these two different kinds of pronouns. Indeed, a great many languages have parallel sets of pronouns that differ in just this way: one kind requires local antecedents and the other prohibits local antecedents. Just what counts as "local" (an issue we will return to) exhibits some cross-language variation, but the similarity is more striking than the difference. There is no a priori reason to expect languages to have more than one kind of pronoun, yet something like the principles above hold in language after language.

Notice, incidentally, that the binding principles interact in an interesting way with the analysis of imperatives suggested in section 3.2. Assuming that the principles are applied prior to the deletion of *you*, the principles correctly predict the following:

(15) a. Protect $\left\{\begin{matrix}\text{*myself}\\ \text{yourself}\\ \text{*himself}\end{matrix}\right\}$!

b. Protect $\left\{\begin{matrix}\text{me}\\ \text{*you}\\ \text{him}\end{matrix}\right\}$!

This provides further evidence that imperatives should be treated as having second-person subjects at some level of analysis.

## 4.2 *Filler-gap dependencies*

Context-free grammars provide a formal mechanism for expressing relationships between elements (words or phrases) that are close to one another in a sentence. But many languages have constructions involving dependencies between elements that may be far apart. An example of this in English is what are known as "*wh*-questions" – that is, questions requiring more than a yes-or-no answer, and hence containing one of the "*wh*-words" (*who*, *what*, *where*, etc.).

To illustrate this, consider the examples in (16).

(16) a. Pat relies $\left\{\begin{matrix}\text{on}\\ \text{upon}\\ \text{*of}\\ \text{*to}\end{matrix}\right\}$ a student.

b. $\left\{\begin{matrix}\text{On}\\ \text{Upon}\\ \text{*Of}\\ \text{*To}\end{matrix}\right\}$ which student does Pat rely?

c. $\left\{\begin{matrix}\text{On}\\ \text{Upon}\\ \text{*Of}\\ \text{*To}\end{matrix}\right\}$ which student does Kim say we think Pat relies?

(16a) shows that the verb *rely* requires a prepositional phrase complement beginning with *on* or *upon*; (16b) shows that, in a *wh*-question, although this phrase comes at the beginning of the sentence, rather than after the verb, the same restriction on the choice of prepositions is maintained; (16c) illustrates that this dependency between the verb and preposition holds even when lots of other material is inserted between them. In fact, there is no limit to the amount of additional text that can intervene.

Similarly, the dependency between verb form and the number (singular or plural) of its subject is preserved, even when the subject is a *wh*-phrase that is far away in the string of words.

(17) a. The $\begin{Bmatrix}\text{teacher}\\ \text{*teachers}\end{Bmatrix}$ dislikes one student.

b. The $\begin{Bmatrix}\text{*teacher}\\ \text{teachers}\end{Bmatrix}$ dislike one student.

c. Which $\begin{Bmatrix}\text{*teacher}\\ \text{teachers}\end{Bmatrix}$ would the parents all claim dislike one student?

d. Which $\begin{Bmatrix}\text{teacher}\\ \text{*teachers}\end{Bmatrix}$ would the parents all claim dislikes one student?

More generally, *wh*-phrases in such questions behave in some ways as though they were in a different position from where they actually occur. Dependencies like preposition selection or verb agreement, which are normally local, can hold between *wh*-phrases and elements far away in the sentence. This can be further demonstrated with the binding principles:

(18) Which dog do you think we saw scratch $\begin{Bmatrix}\text{*yourself}\\ \text{you}\\ \text{itself}\\ \text{it } [\textit{it} \neq \textit{which dog}]\end{Bmatrix}$?

On the surface, *which dog* does not look like the required local antecedent for *itself*, because of the intervening material *do you think we saw*. Moreover, *you* cannot serve as the antecedent for a reflexive object of *scratch*, even though it is closer to the object position. The binding pattern here is just what principles A and B would predict if *which dog* were in the subject position of *scratch*.

A very natural way to account for such relationships in a transformational grammar is to posit a rule that moves *wh*-phrases to the front of the sentence. Then the *wh*-phrases in (16)–(18) can be generated initially in a position close to the relevant verb or reflexive, and the dependencies can be licensed locally, prior to movement.

With such a treatment of *wh*-questions and similar constructions, a question that naturally arises is whether the displaced elements (often referred to as "fillers") move from their initial positions (known as "gaps") to their final positions in one fell swoop or by means of a sequence of smaller movements. That is, in an example like (18), does the filler *which dog*, move from the gap position adjacent to *scratch* in one long movement, as in (19a), or in several smaller movements, as in (19b).

(19) a. $\boxed{\text{FILLER which dog}}$ do you think we saw $\boxed{\text{GAP \_\_\_\_}}$ scratch you

b. $\boxed{\text{FILLER which dog}}$ do you think we saw $\boxed{\text{GAP \_\_\_\_}}$ scratch you

The issue can be formulated in a more theory-neutral way by asking whether the relationship between a gap and its filler is a direct one, or is instead mediated by intervening material. This was a hotly debated topic within generative grammar in the 1970s (sometimes labeled the "swooping vs. looping" controversy). A real measure of progress in the field is that this debate has been definitively settled in favor of "looping." All generative grammarians now recognize that long-distance filler-gap dependencies are mediated by the intervening material.

The key evidence for this comes from languages which require some sort of marking of clauses that intervene between fillers and gaps. Quite a number of such cases have been discovered, from a wide range of language families (see Zaenen 1983 for presentation of a few). Exactly where in the intervening clauses the marking occurs, and what form it takes varies from language to language (though there seem to be some regularities).

A clear and relatively simple example is the relative clause construction[1] in Irish. Irish relative clauses, like those in English, immediately follow the noun they modify, and must contain a gap. The filler for the gap is the noun the clause modifies. Now consider the following examples (adapted from McCloskey 1979):

(20) a. Mheas mé gur thuig mé an t-úrscéal
thought I that understood I the novel
"I thought that I understood the novel."

b. an t-úrscéal a mheas mé a thuig mé
the novel that thought I that understood I
"the novel that I thought I understood"

c. Shíl mé go mbeadh sé ann
thought I that would-be he there
"I thought that he would be there."

d. an fear a shíl mé a bheadh ann
the man that thought I that would-be there
"the man that I thought would be there"

e. Dúirt mé gur shíl mé go mbeadh sé ann
said I that thought I that would-be he there
I said that I thought that he would be there."

f. an fear a dúirt mé a shíl mé a bheadh ann
the man that said I that thought I that would-be there
"the man that I said that I thought would be there"

g. an fear a shíl go mbeadh sé ann
the man that thought that would-be he there
"the man that thought he would be there"

h. an fear a dúirt sé a shíl go mbeadh sé ann
the man that said he that thought that would-be he there
"the man that said he thought he would be there"

i. an fear a dúirt gur shíl sé go mbeadh sé ann
the man that said that thought he that would-be he there
"the man that said he thought he would be there"

Underlining indicates the regions of these sentences that are between gaps and their fillers. That is, the word immediately preceding each underlined piece is a filler for a gap located immediately after the underlining. Now look at the words that have been translated as "that." Where there is no underlining, the Irish equivalent of "that" is either *go* or *gur* (the difference between them is not relevant to the present discussion). But wherever "that" translates an underlined word, the word it translates is *a*. These words are known as "complementizers" (see Baker, chapter 11, this volume, for more discussion of complementizers), because they introduce clausal complements to verbs like *mheas* ("thought"), *shíl* (also translated as "thought"), and *dúirt* ("said"). Examples like those in (20) indicate that Irish employs different complementizers in the region between a filler and a gap than elsewhere.

Modern transformational analyses of filler-gap relationships posit movement through a series of intermediate positions. This fits well with the Irish data, if the complementizer *a* serves as a special gateway through which long-distance movements must pass.

## *4.3 Island constraints*

The notion of gateways for filler-gap dependencies has also been useful in discussions of another much-studied set of phenomena. Although there is no bound on the distance between fillers and gaps, there are a number of constraints on the relative positions in which fillers and their corresponding gaps may appear. These are known as "island constraints," following Ross (1967).

One such restriction on filler-gap dependencies is that the gap may not be in a relative clause if the filler is outside of it. Thus, for example, *wh*-phrases in

English questions cannot fill gaps inside of relative clauses, as illustrated in (21). The relative clauses are enclosed in square brackets, and the gap positions are marked "___".

(21) a. *Which dog did you criticize the person [who kicked ___]?
b. *How many sources does the prosecutor have evidence [which was confirmed by ___]?
c. *Who did everyone envy the writer [whose book claimed ___ was the real Deep Throat]?

If a *wh*-phrase has to pass through intervening complementizer positions on its way from its initial positions (the gap) to its surface (filler) position, then it seems natural to block examples like (21) on the grounds that the relative clauses already have *wh*-phrases (*who*, *which*, and *whose book*) in their complementizer positions. Such an analysis would also rule out gaps internal to embedded questions, as in (22).

(22) a. *Which dog did you ask [who had kicked ___]?
b. *How many sources does the defense wonder [why the prosecutor asked for ___]?
c. *Who did everyone inquire [whose book claimed ___ was the real Deep Throat]?

Not all island constraints are covered by this. For example, a gap cannot be in coordinate conjoined structures not containing its filler, unless all conjuncts have gaps filled by the same filler:

(23) a. *What did they [buy ___and forget their credit card at the store]?
b. What did they [buy ___ and forget ___ at the store]?

A great deal of research has gone into island constraints: classifying them, checking their cross-linguistic variation, and, most extensively, seeking explanations for them. The question of explaining island constraints will be addressed again below.

## 5 Varieties of Generative Grammar

As noted earlier, generative grammar is not so much a theory as a family of theories, or a school of thought. The preceding sections have focussed on common elements: shared assumptions and goals, widely used formal devices, and generally accepted empirical results. (For convenience, the idiom of transformational grammar has been employed in the descriptions of tag questions, imperatives, and filler-gap dependencies, but the discussion in section 5.2 below shows that this was not essential.) This section explores some of the ways in which generative theories differ from one another. There are too many such theories to provide a comprehensive survey (see Sag and Wasow (1999: appendix B) for a brief overview of fourteen theories of grammar), but the following sections characterize some of the major divisions, beginning with a brief description of the historical development of transformational grammar.

## *5.1 Transformational theories*

Transformational grammar has evolved considerably over the decades (see Newmeyer 1986). The earliest work (Chomsky 1957) was concerned largely with showing the inadequacy of context-free grammar for the analysis of natural languages, and with providing precise, explicit transformational descriptions of particular phenomena (largely from English). In the 1960s, transformational grammarians began paying more attention to the relationship between syntax and semantics, leading to heated debates over the best way to incorporate a semantic component into transformational theory. At the same time, the emphasis turned away from providing rule systems in careful detail to exploring the wider implications of transformational analyses. This was when questions about universal grammar and the relevance of linguistic theory to psychology came to the fore (Chomsky 1965). Since the early 1970s, the primary focus of transformationalists has been on developing a highly restrictive theory of grammar – that is, one that narrowly constrains what kinds of descriptions are possible (Chomsky 1981). The goal of this enterprise, as articulated by Chomsky, is to account for language learnability by making the theory so restrictive that a descriptively accurate grammar of any language can be inferred on the basis of the kind of data available to a young child.

As the goals and style of transformational grammar have evolved over the years, the technical details have changed, as well – many almost beyond recognition. Through all these changes, however, this line of research has maintained the idea that sentences are derived by means of a sequence of operations that modify tree structures in prescribed ways. Inherent in this conception is a directionality: derivations proceed from underlying structures to surface forms. This directionality found its way into analyses sketched in this chapter wherever one rule or principle was said to operate "before" another. Examples are the treatment of imperative tags, in which the tag formation transformation had to operate before the imperative rule, and in the account of island constraints in terms of one *wh*-phrase getting into the complementizer position before another one needed to move there.

Many linguists find this sort of talk troublesome. Grammars are supposed to be characterizations of linguistic competence – that is, the knowledge of language that underlies both speaking and understanding. Speaking involves articulating thoughts – going from meanings to sounds; understanding involves extracting meanings from sounds. So, in an intuitive sense, these processes operate in opposite directions. The knowledge of language that is common to both should be process-neutral and hence non-directional. It is possible to regard the talk of operations and directions as strictly metaphorical, a move that has sometimes been advocated. But translating from procedural formulations into more static ones is not always straightforward.

The problem is not just that readers tend improperly to read some psychological significance into the directionality inherent in transformational derivations (though this tendency certainly exists). But psycholinguists and computational linguists who have tried to use transformational grammars as components in models of language use have found that transformational derivations are typically not easily reversible. Precisely worked out systems to parse sentences – whether they are intended as models of human performance or as parts of computer systems for understanding languages – have almost never incorporated the transformational analyses proposed by theoretical linguists. These analyses do not lend themselves to being used in going from the surface form of a sentence to its meaning. Moreover, as noted by Fodor et al. (1974), psycholinguists have been largely unable to find behavioral evidence for the psychological reality of the intermediate stages of transformational derivations. While the nature of the intermediate stages posited by transformational grammarians has changed radically since Fodor et al. made that observation, the observation itself remains accurate.

## *5.2 Non-transformational analyses*

A variety of alternatives to transformational grammar have been developed (see, e.g., Gazdar et al. 1985, Bresnan in press, Steedman 1996, Pollard and Sag 1994). Some grammatical theories have questioned the basic conception of phrase structure embodied in tree diagrams (e.g., Hudson 1984), but most are less radical departures. Instead, they build on context-free grammar, providing enhancements designed for the description of natural languages. This section offers a sample of what such descriptions are like by revisiting some of the phenomena discussed in earlier sections.

Consider first the imperative construction. Imperatives behave as though they had a second-person subject (i.e., *you*), based on evidence from tags and reflexives; but no subject appears in imperative sentences. The transformational analysis offered above posits two distinct trees for imperative sentences, one with a subject and one without. An alternative approach is to posit a single tree without an overt subject phrase, but with the information necessary to get the facts about tags and reflexives right.

Suppose that the category of a word is a complex feature structure (see section 3.3 above) that contains within it at least the following: (i) information about what other kinds of elements it can appear with; (ii) information about its semantics; and (iii) information about how the syntactic information in (i) is linked to the semantic information in (ii). For example, the lexical entry for the verb *protects* should indicate (i) that it requires a third-person singular

NP subject and an NP object; (ii) that it denotes the protection relation; and (iii) the roles played in that relation by the NPs' referents, namely, that the referent of the subject protects the referent of the object. One possible formalization of this information is the following:[2]

(24)
$$\left[\begin{array}{ll} \textit{protects} & \\ \text{POS} & \text{verb} \\ \text{FORM} & \text{present-tense} \\ \text{SUBJ} & \left[\begin{array}{ll} \multicolumn{2}{c}{\text{NP}_1} \\ \text{PER} & \text{3rd} \\ \text{NUM} & \text{sing} \end{array}\right] \\ \text{OBJ} & \text{NP}_2 \\ \text{SEM} & \underline{\text{protect}} < \text{NP}_1, \text{NP}_2 > \end{array}\right]$$

In most cases, the arguments of the semantic relation (that is, the elements between the angle brackets) are linked one-to-one to the syntactic arguments, such as the subject and object. That is the case in (24). In imperatives and some other constructions, however, there may be a mismatch. So, for example, the lexical entry for the imperative use of the verb *protect* might be something like (25).

(25)
$$\left[\begin{array}{ll} \textit{protect} & \\ \text{POS} & \text{verb} \\ \text{FORM} & \text{imperative} \\ \text{SUBJ} & \text{none} \\ \text{OBJ} & \text{NP}_2 \\ \text{SEM} & \underline{\text{protect}} < \left[\begin{array}{ll} \multicolumn{2}{l}{\text{NP}_1} \\ \text{PER} & \text{2nd} \end{array}\right], \text{NP}_2 > \end{array}\right]$$

This representation incorporates both the information that imperative *protect* has a second-person argument and that it has no subject. Further, the second-person argument is the one that plays the protector role in the semantics.

Now, in order to get facts like (15) right, it is necessary to interpret the binding principles as making reference to semantic argument structures. That is, the term "local" in the binding principles, which was left undefined in the earlier discussion, can now be taken to mean "in the argument structure of the same predicate." Thus, principle A now says that a reflexive pronoun must have an antecedent that is an argument of the same predicate.

This characterization of locality makes an interesting new prediction: a reflexive pronoun in object position may not have an antecedent that is only part of the subject. That is, examples like (26) are correctly ruled out.

(26) a. *Your mother protects yourself.
b. *A picture of them upset themselves.

A definition of "local" in terms of simple proximity (based on either word strings or trees) would very likely not cover (26).

Filler-gap dependencies can be handled in a way that is at least partially analogous. A feature – call it GAP – taking another feature structure as its value can encode what is displaced. That is, the value of GAP provides the information that there is a phrase missing in the environment, and specifies what properties the missing phrase should have. This information is represented on every node in the tree between the position of the gap and that of the filler. For example, in a sentence like *What would you like?* the category of *like* would include the information that it has no object, but that it has a GAP value that is linked to the second semantic argument of *like*, as in (27a). This GAP information would be shared by the VP and S nodes above *like* in the tree, as in (27b).

(27) a.
$$\begin{bmatrix} \textit{like} & \\ \text{POS} & \text{verb} \\ \text{FORM} & \text{infinitive} \\ \text{SUBJ} & \text{NP}_1 \\ \text{OBJ} & \text{none} \\ \text{GAP} & \text{NP}_2 \\ \text{SEM} & \underline{\text{like}} < \text{NP}_1, \text{NP}_2 > \end{bmatrix}$$

b.

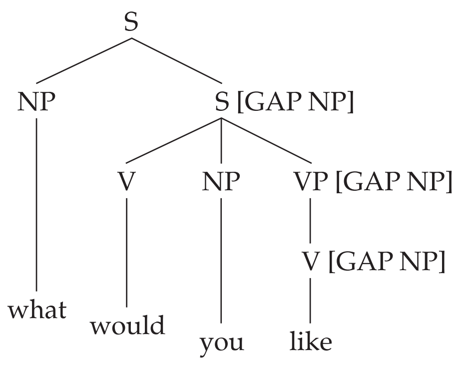


The phrase structure rule licensing the top part of this tree (where the gap is filled) must specify that the features on the NP must match those in the value of the GAP feature. The rule, then, is something like (28), where the identity of the subscripts is intended to indicate identity of all features.

(28) $S \rightarrow X''_1 \quad S[GAP\ X''_1]$

Informally, what (28) says is that a sentence containing a gap may be combined with a phrase of the appropriate type on its left to form a complete sentence (where the appropriate type of phrase is one that has the properties of the missing element that are encoded in the GAP value).

In addition, a principle is required that will guarantee that GAP values are shared between a node and the one immediately above it in a tree, except where rule (28) fills the gap. A GAP value on a node says that there is a gap somewhere within that phrase, and the filler for that gap is outside the phrase; the value of the GAP feature gives the syntactic details of the displaced element.

Many details have been left out of this account, but this minimal sketch is enough to address some of the phenomena discussed in sections 4.2 and 4.3. First of all, the fact that local requirements can be satisfied by a distant filler follows from the fact that the filler must share all its features with the GAP value. Any local requirements on the GAP value must therefore be met by the filler. Consider, for example, (18). If a reflexive pronoun appears as the object of *scratch*, then principle A requires an antecedent that is also an argument of the *scratch* relation. The feature structure for *scratch* in this sentence identifies the GAP value with the first argument of the *scratch* relation, and this GAP value must match the features of the filler, *which dog*. Hence, *which dog*, but not *you* can be the antecedent of a reflexive in (18), despite their positions in the sentence.

Turning now to the swooping vs. looping controversy, it is evident that no such issue arises in this non-transformational analysis. The information about gaps that must be available at the position of the filler is transmitted through the intervening structure. Hence, the presence of a gap in a phrase is necessarily encoded in the category of the phrase. Phenomena like the Irish data in (20) are easy to account for: the choice of complementizer differs depending on whether the clause introduced has a GAP value.

Similarly, island constraints can be straightforwardly formulated in terms of the GAP feature. In fact, if GAP is formulated as suggested above, the island constraints discussed here are almost automatic consequences. Relative clauses and embedded questions are constructions that involve filler-gap dependencies. As long as GAP can have only one value, this makes it impossible to introduce a second gap inside one of these constructions. For example, in an embedded question like *who had kicked that dog* in (29), rule (28) licenses the combination of the filler *who* with the S[GAP NP] *had kicked that dog*.

(29) You asked who had kicked that dog.

If one tried to question the object of *kicked*, yielding (22a), the phrase *had kicked* would need to have two different GAP values.[3]

The facts illustrated in (23) – that filler-gap dependencies in coordinate constructions (i.e., phrases conjoined by *and* or *or*) are impossible unless they involve all conjuncts – are natural consequences of the analysis in terms of GAP. Coordinate conjuncts must share most syntactic features. For example, words with different parts of speech cannot usually be conjoined, and a VP whose FORM value is present-tense cannot be conjoined with one that is [FORM infinitive]. This is illustrated in (30).

(30) a. *Pat became famous and away.
b. *Everyone wishes for comfort and happy.
c. *Chris eats snails and drink wine.

If GAP is one of those features that must be identical across conjuncts, then facts like (23) are an immediate consequence. In a coordinate structure, either all conjuncts lack any GAP value, or they all have the same GAP value. That is, they either all are gap-free, or they all have a gap with the same filler.

## 6 The Future of Generative Grammar

Despite the variety of generative theories of grammar that have been put forward, the field has been dominated throughout its history by the work of one individual, Noam Chomsky. He was its founder; he has been its most prolific innovator; and the mainstream of generative research has always followed his lead. Even the proponents of alternative theories (such as the non-transformational approach sketched in the previous section) generally take work of Chomsky's as the point of departure for their proposals.

In the early years of generative grammar, the field was constituted largely by Chomsky and his students and collaborators. Over the decades, however, the number of generative grammarians has grown exponentially. Under these circumstances, it is remarkable that Chomsky has retained his dominant position. It seems likely that this will eventually change.

Given a saturated academic job market, increasing numbers of linguists are seeking employment in industry. This puts pressure on the field to give more attention to potential applications of its theories. The most obvious type of application for work in generative grammar would be in the development of natural language technologies – that is, computer programs that deal with human languages, e.g., doing machine translation, information retrieval from text files, summarization of texts, and the like. To the extent that such applications motivate theoretical work, considerations of computational tractability are likely to play an increasingly important role in theory construction. Likewise, such applications call for looking at how people actually use language, rather than focussing exclusively on what is grammatically possible. The investigation of real usage data is greatly facilitated by the availability of large on-line text files, which can be sampled and analyzed with computational tools that did not exist until quite recently. This is already having a noticeable effect on the sorts of data used by generative grammarians in their theoretical arguments.

These potential changes should not be worrisome. The history of generative grammar is one of numerous upheavals, as Chomsky has modified the foundations of the theory. These upheavals have been accompanied by vigorous debates and lively competition from alternative frameworks. The result has been – and promises to continue to be – a robust line of research that has greatly enriched our understanding of human linguistic abilities.

## NOTES

1 Relative clauses are noun (or noun phrase) modifiers, such as the bracketed portion of the following:

(i) The student [that you rely on] isn't here yet.

2 This representation glosses over a great deal, including how the formalism is to be interpreted. Italics have been used in place of what should probably be a phonological representation, and underlining is used to designate a semantic relation, with the arguments in the relation listed immediately following, enclosed in angle brackets. The information in (24) also needs to be augmented by characterizations of subject and object in terms of tree configurations, but this is straightforward, at least for English.

3 Both this explanation for these island constraints and the transformational one based on the idea of a blocked gateway rely on the presence of a filler-gap dependency in the embedded structure to establish its status as an island. This seems plausible for English, since, overlapping filler-gap dependencies are not in general possible. Hence, in questions with multiple *wh*-words, only one can be a filler:

(i) What did Pat give to whom?
(ii) To whom did Pat give what?
(iii) *What to whom did Pat give?
(iv) *To whom what did Pat give?

But there are other languages (e.g., Polish and Greek) that permit analogs to (iii) and (iv). In those languages, GAP would need to be allowed to take multiple values (or a value that is a list of feature structures). Unfortunately, the correlation between restrictions on overlapping filler-gap dependencies and island constraints is not perfect, so these matters remain issues of ongoing research.