Chapter 6

Lexical Entries and Accessing Lemmas

A main thesis of this and the following chapters will be that formulation processes are lexically driven. This means that grammatical and phonological encoding are mediated by lexical entries. The preverbal message triggers lexical items into activity. The syntactic, morphological, and phonological properties of an activated lexical item trigger, in turn, the grammatical, morphological, and phonological encoding procedures underlying the generation of an utterance. The assumption that the lexicon is an essential mediator between conceptualization and grammatical and phonological encoding will be called the *lexical hypothesis*. The lexical hypothesis entails, in particular, that nothing in the speaker's message will *by itself* trigger a particular syntactic form, such as a passive or a dative construction. There must always be mediating lexical items, triggered by the message, which by their grammatical properties and their order of activation cause the Grammatical Encoder to generate a particular syntactic structure.

The crucial role of the mental lexicon in the generation of speech makes it necessary to consider in some detail the internal structure and organization of entries in the mental lexicon. This is done in the first two sections of this chapter. Section 6.1 deals with the structure of lexical entries and their mutual relations. Section 6.2 analyses in more detail the aspect of lexical entries that we called "lemmas" in chapter 1. After these more structural sections we will turn to issues of processing. Section 6.3 reviews some major theories of lemma access in speech. This theoretical section is followed by two more empirical ones. Section 6.4 addresses accessing failures, their taxonomy and their potential causes. The time course of accessing lexical items is the subject of section 6.5.

6.1 The Structure and Organization of Entries in the Mental Lexicon

6.1.1 The Internal Structure of a Lexical Entry

A speaker's mental lexicon is a repository of declarative knowledge about the words of his language. From the point of view of language production, each item in the lexicon is a listing of at least four kinds of features. There is, first, a specification of the item's *meaning*. This is the set of conceptual conditions that must be fulfilled in the message for the item to become selected. For the entry *eat*, the meaning is something like "to ingest for nourishment or pleasure". Second, there is a set of syntactic properties, including the category of the entry (V for *eat*), the syntactic arguments it can take (an external subject and an internal object for *eat*; i.e., the verb is transitive), and other properties. Certain items in the lexicon are activated during grammatical encoding by the fulfillment of their syntactic conditions. There is, third, a morphological specification of the item. For *eat* this is, among other things, that it is a root form (i.e., it is not further analyzable into constituent morphemes), that its third-person present-tense inflection is *eats*, and that its past-tense inflection is *ate*. Fourth, there is a form specification—in particular, the item's composition in terms of phonological segments, its syllable and accent structure. For *eat*, the segment structure is a monosyllabic vowel/consonant sequence, with /i/ as vowel and /t/ as consonant. (See figure 6.1.)

There are, moreover, internal relations among these four kinds of information. In particular, there are systematic relations among the morphology of an item, its meaning, and its syntax. Take, for instance, the word *painter*. Its meaning relates to its morphology *paint-er*, the *er* affix

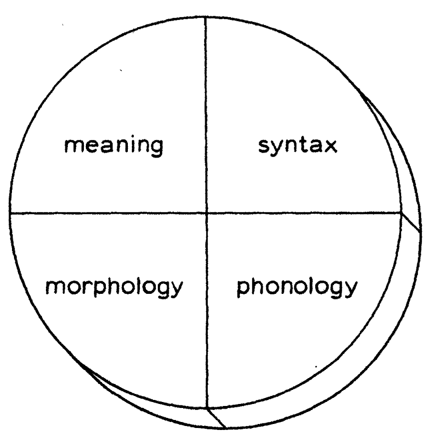

Figure 6.1
Internal structure of an item in the mental lexicon.

expressing the agentive of the action expressed by the verb stem—"one who paints." The *er* affix is, moreover, related to the syntactic category N, and so forth.

There are, probably, additional properties stored with an item. It may have particular pragmatic, stylistic, and affective features that make it fit one context of discourse better than another. The item *policeman* fits better in formal discourse than the item *cop*, which is otherwise very similar in meaning. Certain so-called *registers* (talk to babies, talk between lovers, etc.) seem to select for lexical items with particular connotational properties. Whether such features should be considered as conceptual conditions on the item's use is a matter of much dispute; we will not go into it.

6.1.2 Relations between Items in the Mental Lexicon

Entries in the lexicon are not islands; the lexicon has an internal structure. Items are connected or related in various ways. Let us explore some such relations that are relevant to language production. There are item relations *within* and *between* entries.

Not all lexical items are lexical entries. The various inflections of a verb (e.g., *eat*, *eats*, *ate*, *eaten*, *eating*) are items belonging to the same lexical entry. They are related *within* an entry. The diacritic features for person, number, tense, mood, and aspect will take care of selecting the right item within the lexical entry. We assume that this is generally the case for inflections. Hence, the items *dog* and *dogs* are to be found under the same lexical entry or address, and similarly for *man* and *men*, for *he*, *him*, and *his*, for *big*, *bigger*, and *biggest* (we take comparatives to be inflections), and so forth. But it is not the case for derivations; *act*, *action*, *active*, *activity*, etc. are different lexical entries (Butterworth 1983a).

The relations *between* lexical entries in the mental lexicon are of two kinds: intrinsic and associative. The intrinsic relations derive from the four kinds of features listed for an entry. Items may connect in the mental lexicon because they share certain features. Let us consider each of the four kinds of feature. First, items may have special connections on the basis of meaning. There is evidence in the literature that such connections exist between a word and its hypernym (e.g., between *dog* and *animal*, or between *green* and *color*), between a word and a co-hyponym (e.g., *dog* and *cat*, or *green* and *blue*), between a word and a near-synonym (e.g., *close* and *near*), and so on (Noordman-Vonk 1979; Smith and Medin 1981). Sets of meaning-related items are called *semantic fields*. There is a semantic field of color names, one of kinship terms, and so on (see Miller 1978 for a concise review of this notion). Sometimes speech errors reveal such connections;

for instance, *Irvine is quite clear* involves a blend of *close* and *near* (Fromkin 1973). Second, morphologically determined connections may hold between entries with the same morphological stem. Such entries are said to be *derivationally related*, such as *nation*, *national* and *nationalize*, or *likely* and *unlikely*. Not surprisingly, there are almost always meaning relations involved as well in these cases. There are, third, clear phonological connections between entries in the mental lexicon. Words with the same initial or final speech sounds show connections in speech production, sometimes leading to characteristic errors such as *open* for *over* or *week* for *work* (Fay and Cutler 1977). This suggests (but doesn't prove) that phonologically similar items are, in some way or another, connected in the mental lexicon (other evidence will be discussed in chapter 8). Finally, there is, as yet, no convincing experimental evidence for syntactically conditioned relations between lexical items in the mental lexicon. Are all nouns mutually connected, or all transitive verbs? Surely each such class plays characteristic roles in the generation of speech. This can be dramatically apparent in neurological cases; there are amnesic disorders in which the whole class of nouns has become virtually inaccessible in production. But this does not imply that their members have special mutual connections in the mental lexicon.

The nature of intrinsic relations is an issue in itself. There may be *direct* connections between lexical items, or the relations can be *mediated*. A direct semantic relation, for instance, would be one in which all co-hyponyms of an item were listed with the item. The entry for *green* would contain a listing of co-hyponym addresses—the ones for *blue*, *red*, and so forth. A mediated semantic relation would be one where there is a relation between the *concepts* GREEN and BLUE, but without mutual reference at the level of the lexical entries *green* and *blue*. This distinction is important for the analysis of lexical intrusions in section 6.4. Intrinsic connections are not a *necessary* consequence of feature sharing between lexical entries. Entries connect on the basis of shared word-initial consonants, but a shared first-syllable-final consonant is probably irrelevant.

Associative relations between entries have no necessary basis in their semantic properties; rather, the basis lies in the frequent co-occurrence of the items in language use. *War* and *death* and *truth* and *beauty* are two cases in point. Though these connections are initially mediated by complex conceptual relations, they have become direct associations between lexical items. When the one item is used, the other one will be primed, even when the original conceptual connection is not at issue in the ongoing discourse. It is to be expected that some intrinsic-meaning relations will also develop into strong associative relations, because meaning-related items tend to cooccur in discourse. Antonyms such as *left* and *right* and *big* and *small* are cases in point. (See H. Clark 1970 for a semantic analysis of associative relations.)

6.1.3 Retrieving versus Constructing Words

The mental lexicon is, we assume, a passive store of declarative knowledge about words. It does not contain procedural knowledge, which makes possible the generation of new words. Do speakers generate new words when they speak? The answer is probably Yes for all languages, but the degree of spontaneous new-word formation during normal speech varies drastically between languages.

English is at one extreme of the distribution. English speakers seldom produce words they have never used before. An extreme case cited by Bauer (1983) occurred in the spontaneous utterance *I feel particularly sit-around-and-do-nothing-ish today.* Less extravagant cases are new constructions with *-ful* (such as *bucketful*) or *un-* (such as *unnarrow* or *unobscure*), and new compounds (such as *my lecture-tie*). For English one of the most productive cases is number names, as Miller (1978) observed. Their unlimited amount makes storage in the mental lexicon impossible (we probably have no lexical entry for a number such as 4,257). Still, the use of such new formations is exceptional in everyday language. By and large, English speakers use words that they have frequently used before, and these words are probably stored in the mental lexicon. The other extreme occurs in speakers of certain agglutinative languages, such as Turkish and Finnish. In these languages, words consist of strings of morphemes—a root plus affixes, each adding to the meaning of the word. These strings can become very long—perhaps *arbitrarily* long, as Hankamer (forthcoming) argues for Turkish. In Turkish, the root morpheme for *house* is *ev*. Adding *ler* makes it plural (*ev-ler*), adding *den* creates a word meaning "from the houses" (*ev-ler-den*), and so on. The word for "to the ones of those that are in our pockets" is *ceb-lar-lmlz-da-ki-kar-nln-ki-n-ya*. Hankamer computes that, even if one ignores the existing possibility for affixes to recur in a string, a single Turkish noun can appear in more than 4 million different forms. It is obvious that most of these forms are never used by a Turkish speaker; they are not stored forms in his mental lexicon. But they will be recognized as possible words, and they will be interpreted correctly. This situation is quite comparable to the case of number names for an English speaker. However, for the Turkish speaker this is the normal case rather than the exception. The stored forms in his mental lexicon will probably consist of all stems (such as *ev* for "house"), all possible affixes, and a certain number of frequently used multimorphemic words. In order to use this "passive" store productively, the Turkish speaker must have access to *lexical procedural knowledge*—ways of building new words, given the conceptualizations in the preverbal message. In other words, such a speaker must have a strongly developed processing component dedicated to *lexical encoding*, which produces new words as output.

The conclusions to be drawn from this comparison are that speakers can produce more words than just the ones stored in their mental lexicon and that they have the capacity to construct new words while they are speaking. But languages differ enormously in the degree to which they exploit these word-constructional capabilities. While a Turkish speaker's grammatical encoding consists for the most part of such lexical encoding, an English speaker is extremely "conservative" in the sense that he normally uses words he has used often in the past. For the English speaker, lexical encoding plays a very minor role in grammatical encoding; the action is in syntactic encoding. A theory of the speaker should, of course, encompass both kinds of grammatical encoding. As a matter of fact, however, almost nothing is known about the psychology of lexical encoding. Of necessity the rest of this chapter, as well as chapter 9, will deal with lexical access in the sense of retrieving items stored in the mental lexicon, and these items are all *words* (but see subsection 6.1.4). This does not preclude the possibility that these word items have strong morphemic relations to one another; see Butterworth 1983a and Cutler 1983a for reviews of the evidence.

Little attention will be given to word-constructional processes. This is not a dramatic restriction of my discussion as far as speakers of English or Chinese are concerned, but it will underexpose theoretical issues in the production of Turkish or Japanese, where lexical items are often morphemes, not words.

6.1.4 Phrases and Idioms

Speakers have, over and above a stock of words, stocks of phrases and idioms. Certain concepts map directly onto phrases, such as *Dutch uncle* or *red tape*. That these are phrases, not compound words, is apparent from their stress patterns; they do not have compound-word stress, as do *blAckbird* and *hOt dog*. But they are special in that their meaning is opaque; it does not—as in the case of syntactic phrases—derive from the meanings of their parts. The difference between *red tape* and *green tape* is not in their color, as is the difference between *red apple* and *green apple*. They have *idiomatic* meanings. Such idiomatic phrases also have restricted syntactic possibilities. It is all right to pluralize *hot dog* to *hot dogs*, but it is less good to use it in the comparative degree (*I have a hotter dog than you*). Or take an idiomatic expression like *to kick the bucket*. It has, like the earlier examples, no transparent meaning. It also has rather restricted syntactic possibilities. One can say *he has kicked the bucket*, but one cannot very well say *he is kicking the bucket* (although one can say *he is dying*). Also impossible are constructions such as *the bucket was kicked by John* and *it was the bucket that John kicked*.

There is an extensive linguistic literature on these issues (see e.g. Makkai 1972 and Cruse 1986), but not much is known about the speaker's generation of idioms. We will assume that idiomatic collocations are entries in the mental lexicon. Each entry consists of one or more items. The entry for *kick the bucket* contains items for infinitive and for past tense, but none for progressive. The entry for *red tape* contains the singular and plural item, but no comparative one. Idioms, like words, have their characteristic conceptual conditions. If such a condition is met in the message, the idiom will be accessed.

6.1.5 Lexical Entries, Lemmas, and Morpho-Phonological Forms

The processes of grammatical encoding are, to a first approximation, independent of the phonological information in lexical entries. In addition, for languages without much lexical encoding, a word's morphological composition is, on first approximation, irrelevant for grammatical encoding. Only a lexical entry's meaning and syntax are relevant. There is now a tradition of following the terminology introduced by Kempen and Huijbers (1983) and calling this part of an entry's composition a *lemma*. The entry's morphological makeup and its phonological properties, on the other hand, are essential for phonological encoding; the entry's lemma is, by and large, irrelevant at that stage of processing. Hence, from the viewpoint of language production a lexical entry can be split up into two parts: its lemma and its form information (figure 6.2). This theoretical distinction can be extended to the mental lexicon as a whole: Lemmas can be said to be "in the lemma lexicon," and morpho-phonological forms to be "in the form lexicon." Each lemma "points" to its corresponding form; i.e., it can refer to the address in the form lexicon where the information for that item is stored. This was discussed in subsection 5.1.2.

The partitioning of the mental lexicon in two kinds of store is no more than a spatial metaphor acknowledging the existence of two kinds of

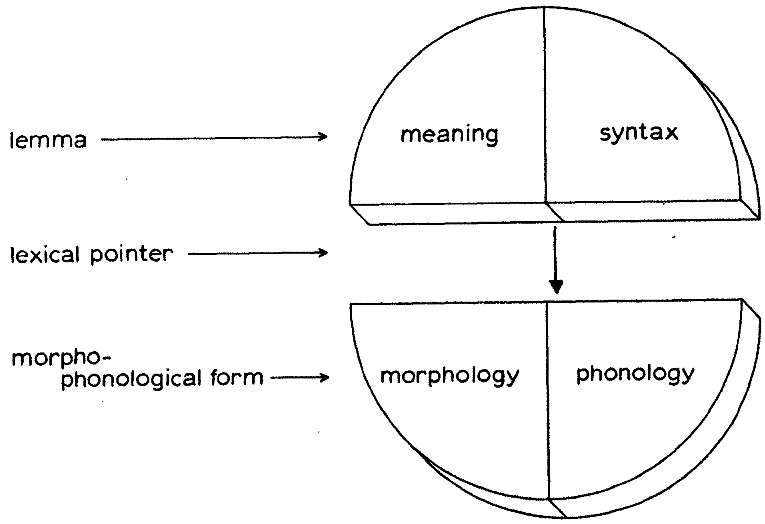

Figure 6.2
A lexical entry consisting of a lemma and a morpho-phonological form. (After Levelt and Schriefers 1987.)

internal organization in the mental lexicon: one according to the meaning of items and one according to their form properties. These two rather independent kinds of organization appear in various production phenomena, such as speech errors and tip-of-the-tongue phenomena (which we will consider in subsequent chapters). The distinction should, however, not be overstated. In particular, we should not conclude that a lexical entry cannot be retrieved as a whole—i.e., that retrieval of the lemma must always precede retrieval of the item's form. This is still an open issue (see subsection 6.5.4).

6.2 The Structure of Lemmas

6.2.1 Semantic and Syntactic Properties

The *semantic* information in a lemma specifies what conceptual conditions have to be fulfilled in the message for the lemma to become activated; it is the lemma's *meaning*. These conditions can be stated in the same propositional format as messages. Let us consider as an example the conceptual specification for the verb *give* as it appears in the sentence *the child gave the mother the cat*. The message structure underlying that sentence was represented in figure 5.1a.

The conceptual specification for *give* is something like the following.

• Conceptual specification:

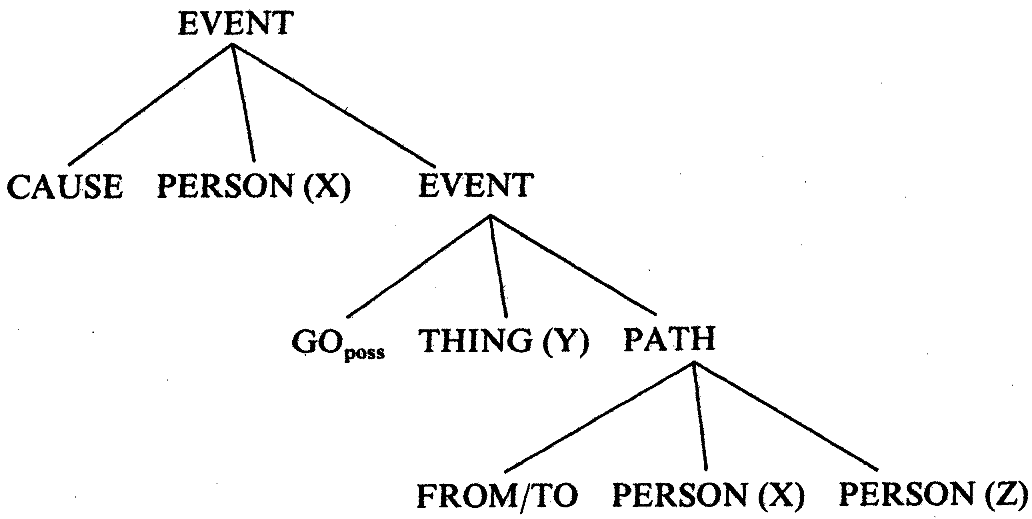

This conceptual specification involves three variables: X, a PERSON, which is both the agent of the causative EVENT and the source of the PATH; Y, a THING, which is the theme, and Z, a PERSON, which is the goal of the PATH. This can be viewed as a conceptual "template" to be matched with substructures of the message. It is irrelevant in what way the variables X, Y, and Z are bound in the message. But the fact that the conceptual specification *has* these three slots for conceptual arguments will turn out to be important. They are the three variables in the conceptual structure to which grammatical functions can be assigned. This must be explicitly stated for the lemma.

- Arguments of the conceptual function: (X, Y, Z).

These arguments fulfill certain thematic roles in the message; they depend on the functions CAUSE, GOposs, and PATH. As we will shortly see, some important grammatical generalizations about lemmas can be expressed in terms of thematic roles.

It is, moreover, important that the three arguments are of the conceptual categories PERSON, THING, and PERSON, respectively. These are called *selectional restrictions* on the use of *give*.*

The message of figure 5.1 binds the conceptual variables as CHILD, CAT, and MOTHER. (We are assuming that an animal is a special kind of THING.) When there is a match between conceptual specification and message (and in this case there is), the lemma is retrieved, which means that

* The selectional restrictions PERSON, THING, and PERSON may, in fact, be too restrictive for *give*. Take for instance the use of *give* in the sentence *The bright lights gave Santa's arrival a colorful appearance*. Selectional restrictions on the kinds of entities that can fill argument slots in the conceptual specification of a lemma are often very hard to define. We will ignore the issue entirely, but see Bresnan 1982.

its syntactic properties become available. This brings us to the syntactic properties of a lemma.

A lemma's *syntactic* information specifies the item's syntactic category, its assignment of grammatical functions, and a set of diacritic feature variables or parameters. Let us consider again the lemma for *give* as an example. It has the following specifications for syntactic category and functions.

- Syntactic category: V.

This means that the lemma will act as a main verb in grammatical encoding.

- Grammatical functions: (SUBJ, DO, IO).

This means that the lemma for *give* requires a subject, a direct object, and an indirect object. In the example these grammatical functions are fulfilled by the phrases *the child*, *the cat*, and *the mother*, respectively. Nothing, however, is said about the ordering of these three. In fact, the sentence *the child gave the mother the cat* has them in the order SUBJ-IO-DO. Still, the order of listing grammatical functions is not arbitrary. It will, by convention, correspond to the order in which the arguments are listed. The first argument (X) has to be realized as the first grammatical function (SUBJ), the second argument (Y) as the second grammatical function (DO), and the third argument (the goal/recipient) as the third grammatical function (IO). It should be remembered that the ordering of the arguments in the message is no more than a notational convention (subsection 3.3.3). But the lemma includes a specification of which conceptual argument is to be mapped onto which grammatical function.

This mapping is not always of the simple one-to-one sort, as for *give*. It is, in particular, not always the case that the number of grammatical functions is equal to the number of arguments; there may be more functions than arguments. This is often the case when the lemma requires a complement of the type V-COMP or S-COMP. Examples of verbs requiring a V-COMP are the so-called raising verbs, such as *believe*. The conceptual structure specification for *believe* has two variables, X and Y, where X is the one who believes, the experiencer, and Y is some state of affairs. But *believe* assigns three grammatical functions in sentences such as *Atilla believed the world to be flat*: SUBJ, DO, and V-COMP. The "additional" direct object *the world* is, in fact, the "raised" subject of the verbal complement that expresses the state of affairs, namely that the world is flat. The lemma for *believe* specifies where the additional grammatical function comes from: DO = V-COMP's SUBJ. It also requires V-COMP to be in the infinitive (*to be flat*). So, for *believe* we have the following additional syntactic specification.

- Relations to COMP: DO = V-COMP's SUBJ
  V-COMP has diacritic parameter "inf".

The lemma for *give*, however, has no such functional relations specified.

Let us now turn to the last two items on the lemma's list. Subsection 6.1.4 discussed the sense in which a lemma "points to" a morpho-phonological form. The lemma relates to specific form information; it "points to a form address." Let the address for the form information of *give* in the speaker's lexicon be 713.

- Lexical pointer: 713.

That address or entry contains several word forms: the inflections of *give*, i.e., *give*, *gives*, *gave*, *given*, *giving*. They can be distinguished only by assigning values to several features or diacritic variables. These parameters are listed in the final item on the lemma's list.

- Diacritic variables: tense, mood, aspect, person, number.

Also to be included here is the lemma's pitch-accent value. The verb *give* may or may not receive focus and hence pitch accent during the generation of the utterance, as was discussed in the previous chapter. The values of all these variables are collected during the process of generating the surface structure. The word-form inventory can be successfully addressed only when all of a lemma's diacritic parameters have been fixed.

The lemma for *give*, before these parameters have been collected, is summarized in figure 6.3. This example was given to itemize the different

give: conceptual specification:
  CAUSE (X,(GOposs (Y, (FROM/TO (X,Z)))))

conceptual arguments: (X, Y, Z)

syntactic category: V

grammatical functions: (SUBJ,DO, IO)

relations to COMP: none

lexical pointer: 713

diacritic parameters: tense
aspect
mood
person
number
pitch accent

Figure 6.3
Lemma for *give*.

types of knowledge a speaker has stored in a lemma. The knowledge *content* under these items will, of course, differ from lemma to lemma. In the next subsection we will consider some further aspects of verb lemmas—in particular, the ways in which verbs relate conceptual arguments to grammatical functions. In subsection 6.2.3, more will be said about other major lemma categories (nouns, adjectives, and prepositions) and about auxiliaries.

6.2.2 Grammatical Functions and Conceptual Arguments

The way in which grammatical functions are assigned to conceptual arguments is not entirely arbitrary in the world's languages. When the verb has an agent as a conceptual argument, it is usually paired with the subject function, as in the example for *give*. If it has a theme or a patient over and above an agent, this argument tends to occupy an object slot, as in *John kicked the ball* (where *ball* is theme) or *George killed the dragon* (where *dragon* is patient). If, however, the subject slot is not occupied, patient and theme are preferably expressed by the subject function. This happens, for instance, for verbs like *fall*, which require a theme but require no agent as argument: *the bottle fell*. If the verb has a source or goal over and above an agent and a theme or patient, this tends to be mapped on an oblique grammatical function. This happens for a verb like *send*; *Henry sent the letter to Japan* has an agent, a theme, and a goal. Here the goal (*Japan*) ends up in oblique function, in a prepositional phrase. But source or goal is "promoted" when object or subject functions are not occupied by other arguments. The verb *leave*, for instance, has a theme and a source, and they map onto subject and object, respectively: *Marcia left Italy*. One could say that there is a preference hierarchy for grammatical functions, from subject via direct and indirect object to oblique functions. In addition there is a "pecking order" for thematic roles, from agent via theme and recipient to source and goal (see Pinker 1985 and Bock and Warren 1985 for a further discussion of these issues).

But not all verbs show this canonical order of assigning grammatical functions to their conceptual arguments. An example is *receive*, where the goal is encoded as subject and the agent as oblique (*the mother received the cat from the child*). More important, the grammar of a language may provide quite regular means for changing the mapping order. The mapping order for *give*, as presented in figure 6.3, is the following:

```
X(agent), Y(theme), Z(goal)
  |         |         |
SUBJ,      DO,        IO
```

But there is another possible mapping for *give*:

X(agent), Y(theme), Z(goal)
| | |
SUBJ, DO, OBL

This mapping is used in the generation of sentences such as *the child gave the cat to the mother*. It has the three arguments in canonical mapping order. This means that the specifications for the lemma *give* in figure 6.3 are not complete. Under "grammatical functions" the alternative mapping (SUBJ, DO, OBL) should also be given. This pair of mappings for *give* is not an exception; there are similar pairs for other verbs, such as *send*, *sell*, and *buy*. In each of these cases there is an alternation between an indirect object function and an oblique function. This regularity—called a *lexical rule* by Bresnan (1978, 1982)—is the "dative shift" rule.

Another lexical rule is "passive". In English most transitive verbs have an active and a passive mapping of conceptual arguments onto grammatical functions. They differ in the way grammatical functions are assigned to two of their arguments, mostly agent and patient:

Active lemma	**Passive lemma**
agent patient	agent patient
| |	| |
SUBJ OBJ	OBL SUBJ

Take the verb *kill*. In the active *George killed the dragon*, agent and patient are in their canonical grammatical functions, subject and direct object. The passive variant imposes a different function assignment—either as in *the dragon was killed*, where the patient has moved up to subject position and the agent is without grammatical encoding, or as in *the dragon was killed by George*, where the agent appears as an object in the oblique *by*-phrase. In the case where the passive variant is chosen for the grammatical encoding of conceptual arguments, the lexical pointer gets the diacritic parameter for perfect-tense morphology. This will select for *killed* in phonological encoding. It will also lead to retrieving an auxiliary verb (*was*), as will be discussed in subsection 6.2.4.

Most languages have passives, and one wonders what they can do for a speaker. What they seem to have in common is that they demote the argument that would otherwise be encoded as subject, i.e., the argument that occupies the top of the pecking order. This will often be the agent, but in the agentless case it may be the theme (*the cat was given to the mother*). Why would a speaker want to deny the top argument its canonical position when assigning grammatical functions? The pragmatic reason is probably that the grammatical subject is universally the preferred way to encode the topic (see Keenan 1976). Usually the sentence topic coincides with the top argument, but it need not. When another argument is assigned the role of topic in the message, a mapping will be needed that allows it to appear as the grammatical subject. This is what passive mappings provide for. They "demote" the top argument from the subject position, freeing that slot for the topicalized argument.

What happens to the demoted argument? The passive does not require its grammatical encoding. If the argument (say, the agent) is absent from the message, it will also be absent from the sentence (*the dragon was killed*). But if it is still an argument to be expressed (though not as a topic), the passive in English allows for at least the agent's encoding at the very bottom of the functional hierarchy: in an oblique *by*-phrase (*the dragon was killed by George*). This is not a very essential property of passive verb forms, and many languages do not have it. What is essential is the lexical possibility of demoting the top argument in order to topicalize a lower one. Normally this results in the reduction of the number of arguments by one. This state of affairs can be pictured as follows (in correspondence with Pinker 1985):

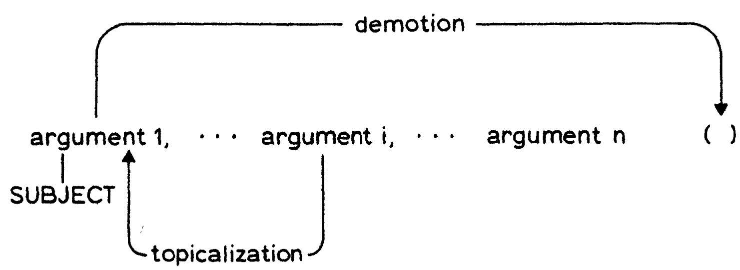

Note that the promoted argument need not be the second in line; the passive variant of *give* can be used to express a topicalized theme (*the cat was given to the mother*) or a topicalized goal (*the mother was given the cat*).

6.2.3 Prepositions, Adjectives, and Nouns

Though the lemmas for verbs play the major role in assigning grammatical functions during speech, formulating is impossible without lemmas for nouns, adjectives, prepositions, adverbials, and other categories such as auxiliaries, determiners, conjunctions. Prepositions, nouns, and adjectives can be heads of PPs, NPs, and APs, respectively, and they may subcategorize for other elements in these phrases. A few words about each of these are presented in this subsection. Subsection 6.2.4 will touch on auxiliaries and other minor categories.

Prepositions

Subcategorization by prepositions is usually obligatory; for example, *toward* requires an NP which expresses the direction argument (*Frederic pointed toward the sun*). The lemma for the preposition specifies the argument and the grammatical function. For *toward* they are the goal and a prepositional object function. The latter becomes expressed as a case feature on the subcategorized NP (*toward him*, not *toward he*). The lemmas of most prepositions have, furthermore, a conceptual specification. For *toward* it specifies that the concept is a DIRECTION, i.e., a PATH that does not contain the goal or reference object (see subsection 3.3.1). Miller and Johnson-Laird (1976) give detailed conceptual specifications for various prepositional lemmas.

But there are also prepositions with empty conceptual specifications. These are sometimes called *idiomatic* prepositions. The *for* in *George waited for the dragon* has an idiomatic relation to the verb and is otherwise meaningless. Still, it does specify case for the NP it is heading. Because it is idiomatic, it is listed in the verb's lemma.

wait:
  conceptual arguments: (X, (Y))
  grammatical functions: (SUBJ, (for OBJ))

This means that *wait* has two conceptual arguments: the one who waits and the entity waited for. The first one is obligatory, the second one optional (as in *John waited*). The optional argument will be grammatically encoded as "for OBJ". But what is the status of this "for"? We will assume that it is the address of the nonidiomatic lemma *for*, i.e., of the meaningful preposition. Activating the lemma *wait* when there is an object argument will automatically involve addressing the lemma *for*. Notice that in this case the conceptual activation conditions of *for* play no role in its activation, and they are also irrelevant in the further generation of the prepositional phrase. Only the syntactic features of *for* and its lexical pointer are relevant.

Adjectives

Adjectives can appear as specifiers in NPs (*hard work*), but also as heads in APs. In the latter case the adjective can subcategorize various elements in the phrase, and the adjective lemma specifies how this is done. One somewhat overworked example should suffice to show this. *Eager* and *easy* can head the adjectival phrases *eager to please* and *easy to please* in *John is eager/easy to please*. Both lemmas are of category A and specify the following grammatical functions:

grammatical functions: (SUBJ, S-COMP)

In both examples the subject is *John* and the S-COMP is *to please*. The difference is in the relations to COMP specified in the lemmas.

eager:
Relations to COMP: SUBJ = S-COMP's SUBJ
S-COMP has diacritic parameter "inf"

easy:
Relations to COMP: SUBJ = S-COMP's OBJ
S-COMP has diacritic parameter "inf"

This makes *John* the subject of *please* when *eager* is the adjective and the object of *please* when *easy* is the adjective. In both cases S-COMP's verb will be in the infinitive. The specifications of relations to COMP in the lemmas function as instructions in the sentence-generation process to create a subjectless phrase to represent the S-COMP of *eager* and an objectless phrase for the S-COMP of *easy*. These are only examples of *grammatical* specifications in the lemmas of adjectives; I will refrain altogether from discussing their conceptual and form specifications.

Nouns

Nouns as heads of phrase can have specifiers such as determiners and quantifiers, but some nouns also subcategorize for complements. The expression of these complements is, however, always optional. Such complements can be PPs (as in *the father of Sylvia*), NPs (as in *Germany's president*), or Ss (as in *the claim that the world is flat*). In all three cases the complement expresses a proper argument of the noun's conceptual specification: a father's child, a president's domain, and the state of affairs the claim is about. They should be distinguished from complements that are not subcategorized by the noun, such as in *the father with money* or *France's cheese*. These complements represent conceptual modifications, not arguments, and they are not specified in the head noun's lemma.

Noun lemmas fall into two major classes: *proper nouns* and *common nouns*. Each proper-noun lemma specifies a conceptual *token*; the conceptual specification need be no more than a pointer to the token's address in memory. Examples of proper names are *Mount Everest*, *Hans Brinker*, and *World War II*. It may well be, however, that a proper name is not purely referential but has additional intentional features as well. It is, for instance, unlikely for *Mount Everest* to be a war, or *World War II* to be a mountain. There is usually some type information in the kind of name given to a token (Carroll 1983a, b).

Common nouns are nouns whose lemmas contain exclusively *type* information as conceptual specification. Examples are *horse*, *furniture*, and *democracy*. The lemma for *democracy* specifies several conceptual properties which a type or token concept in the message must display in order for the lemma to become activated: rule by the people, free elections, or what have you. The problems about the precise structure of these conceptual specifications are horrendous; they will be left undiscussed here, but some of them will be taken up in section 6.3.

The common nouns further subdivide along syntactic lines into *count nouns* and *mass nouns*. The lemmas of English count nouns (such as *dog*, *event*, *belief*) have a diacritic variable for number, which can have the value 'single' or 'plural'. They can also accept either definite or indefinite determiners. Lemmas of mass nouns (such as *sugar*, *happiness*, *furniture*), on the other hand, have the fixed syntactic feature 'singular' and usually accept definite determiners only. Though the distinction between a count noun and a mass noun is correlated with properties of the conceptual specifications of these lemmas (count nouns tend to refer to concepts that are countable; mass nouns refer to substances), it is principally a grammatical distinction. There is, for instance, no conceptual difference in countability or substantialness between *the paintings in my house* and *the furniture in my house*, or between *the blessings of my marriage* and *the happiness of my marriage*.

6.2.4 Auxiliaries and Minor Categories

As we saw above, the retrieval conditions for idiomatic prepositions are not of a conceptual kind. This can also be the case for certain auxiliaries—in particular, English *have*, *be*, and *do*. When a speaker generates the sentence *The child has given the cat to the mother*, the message involves a particular time index (see subsection 2.2.3) which will induce the grammatical encoder to generate a perfect-tense verb phrase. This is done by a special VP procedure (to be discussed in the next chapter), which activates the lemma of the auxiliary verb *have* and provides it with the appropriate diacritic features (in the example: third person singular, present). In other words, the auxiliary lemma is not conceptually activated; rather, it is addressed by a syntactic building procedure. Since there are no direct conceptual conditions for the lemma's activation, its meaning specification is empty.

But this is not so for all auxiliaries. The *modal* auxiliary verbs (such as *can*, *may*, and *shall*) do have independent semantic activation conditions, relating to possibilities and necessities in the message (Lyons 1977; Seuren 1985). They can be treated as main verbs that take verbal complements (Bresnan 1982; Pinker 1985). The syntactic specifications in their lemmas are the following.

modal verb: Syntactic category: V
Grammatical functions: (SUBJ, V-COMP)
Relations to COMP: SUBJ = V = COMP's SUBJ,
V-COMP has diacritic value 'inf'

The surface structure for *Olithia can go* will thus be

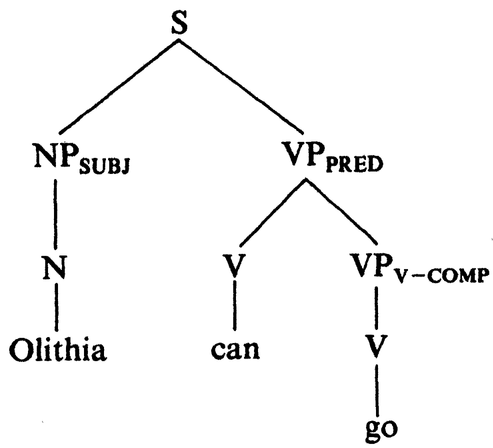

where the subject of *go* has become the subject of *can*, and where *go* is indexed for infinite form.

Most English auxiliaries have rather degenerate inflectional possibilities (see Pinker 1985). The verb *may*, for instance, cannot be used as an infinitive, a perfective, or a progressive (there is no *to may*, *mayed*, or *maying*). The English Formulator cannot do what a Dutch one can, namely produce constructions such as *to may go*. There are no conceptual reasons for this to be impossible; it is due entirely to the syntactic properties of the lemmas involved.

Certain lemmas can (I repeat) be directly addressed by syntactic procedures, whether or not they have a semantic specification. This is the case not only for certain prepositions and auxiliaries, but also for other, especially minor syntactic categories which cannot be heads of phrases. In chapter 7 we will see this happen for certain determiners—in particular, English *the* and *a*. The next section will, however, be devoted entirely to the *conceptual* activation of lemmas.

6.3 Theories of Lemma Access

If indeed the process of grammatical encoding is lexically driven, the way in which lemmas are accessed is a key issue for theories of language production. However, very little is known about how lemmas become activated by fragments of the message. This is not totally surprising, in view of the magnitude of the problem. A speaker with a normal speech rate produces some 150 words per minute (Maclay and Osgood 1959)—on the average, one every 400 milliseconds. Under time pressure the rate can easily be doubled to one every 200 milliseconds. A normal educated adult speaker of English has an active vocabulary—i.e., words he actually uses in his everyday speech—of about 30,000 words.* A speaker makes the right choice from among these 30,000 or so alternatives not just once but, in fluent speech, continuously two to five times per second—a rate that can be maintained without any clear temporal limit. There is probably no other cognitive process shared by all normal adults whose decision rate is so high. Still, the error rate is very low. Collectors of speech errors know that a day's catch is meager. Garnham, Shillcock, Brown, Mill, and Cutler (1982) found 191 slips of the tongue in a text corpus of 200,000 words—about one slip per 1,000 words. Almost half of these (86) were lexical errors. Hotopf (1980) did not find more than 125 whole-word slips of the tongue in the tape recordings of eight conference speakers.

The major issue for a theory of lemma access is how the conceptual specification of a verb, a common noun, an adjective, or another content word comes to "resonate" with some fragment of the message, and how a speedy and accurate choice is made between different lemmas at these high processing rates. The present section will review some theories of access. To begin, let us consider two theoretical issues which are crucial for evaluating any theory of lexical access.

6.3.1 Parallel Processing and Convergence

Parallel Processing

High-speed access requires parallel processing. It would, for instance, be disastrous if, for any concept to be expressed, all lemmas in the mental lexicon would have to be successively checked for their appropriateness until a fitting one was found. This would involve several thousands of tests (from two to five per second) for each new word in the sentence. A touchstone for theories of access will be the degree to which they can reduce such sequential testing. Not only is a parallel account of access necessary for theoretical reasons; there is also convincing empirical evidence for parallel lexical access in speech. Speech errors (blends, in particular) often

* The number will vary greatly from speaker to speaker and is, moreover, hard to measure. Oldfield's 1963 estimate for Oxford undergraduates was a vocabulary size of about 75,000 words, but these were words *understood* by students (i.e., their *passive* vocabularies). Their *active* vocabularies must have been substantially smaller.

reveal the simultaneous activation of two near-synonyms, such as in *stummy* for *tummy* and *stomach* (Fromkin 1973).

There is similar evidence for another kind of parallel processing in lexical access. Different message fragments can trigger lemmas in parallel. If two or more fragments are available at the same time, their activation of lemmas need not be serial. If, for instance, some function/argument structure is to be expressed (say, that the text fits the page), all conceptual fragments—FIT, TEXT, and PAGE—may be simultaneously available for expression in the speaker's mind, each initiating its own search for a lemma. One gains a factor of 3 in speed if there are, on average, three different but overlapping accessing processes running during fluent speech production. This type of parallelness may also lead to a characteristic type of errorproneness. A speaker may happen to say *the page fits the text*, and this type of speech error is indeed not uncommon.

Convergence

Whatever the accessing algorithm, it must eventually converge on a single item, the correct one. With E. Clark (1987, 1988), we will assume that there exists no real synonymy in a language. There is a "principle of contrast" which says that all forms in a language contrast in meaning. If word *a* correctly expresses notion A, then word *b* cannot also correctly express notion A.

Convergence would be guaranteed when concepts and words would entertain a simple one-to-one relationship. This is, to some degree, realized for proper nouns and the individuals they refer to. In these cases, as we saw in subsection 6.2.3, the conceptual specification of the lemma might be no more than a pointer to the token individual's address in memory. If all concept-to-word relations were of this kind, accessing lemmas would be like typewriting: Each concept would hit its own key, printing its own lemma character. In more psychological terms: Activating a lemma is a simple reaction involving no choice. Since Donders (1868), we have known that these mental reactions are the speediest of all. The one-to-one mapping ensures high speed as well as faultless convergence.

The alternative, "componential" view is that a word's conceptual specification (its meaning) is some conglomerate of conceptual components, or features, which have to be checked against smaller or larger fragments of the message. In a procedural theory, such as the one developed by Miller and Johnson-Laird (1976), the components are predicates, and for each predicate there is a testing procedure that can evaluate whether the predicate is true or false for the concept at hand. The conceptual specification for the lemma *give* in figure 6.3 is such a conglomerate of predicates or functions, such as CAUSE, GOposs, and FROM/TO. To test the lemma's applicability, the concept should probably be checked for the presence of each of these components, and for their correct relations, since none of them should contradict the concept. It should be obvious that the more detailed the componential structure of a lemma's cognitive specification, the more tests will have to be executed in order to converge on the correct lemma. And this will be to the detriment of accessing speed if the tests are not run in parallel. Though some degree of componentiality is unavoidable in a theory of word meaning, there is a virtue in carefully testing its empirical necessity, particularly for a theory of access in speech production. Fodor, Garrett, Walker, and Parkes (1980) argue against the componential view and provide some empirical results in support of their position. In discussing his language-production model, Garrett (1982) also proposes a close match between the conceptual inventory of which messages are composed and the words of a language.

There is one particularly nasty convergence problem that has not been solved by any theory of lexical access. I will call it the *hypernym problem*. It appears in different guises for different theories, as we will see, but a short formulation is this:

The hypernym problem When lemma A's meaning entails lemma B's meaning, B is a hypernym of A. If A's conceptual conditions are met, then B's are necessarily also satisfied. Hence, if A is the correct lemma, B will (also) be retrieved.

To give an example, if the speaker intends to express the concept DOG, then all conceptual conditions for the activation of the lemma *animal* are satisfied because the meaning of *dog* entails the meaning of *animal*. Why then does the speaker not say *animal* instead of *dog*? To put it more casually, why do speakers not talk in hypernyms, such as *the person moves* instead of *the man walks*, *the thing travels* instead of *the plane flies*, or *the event caused an event* instead of *his leaving made her weep*?

The hypernym problem is a touchstone for theories of lexical access. After reviewing access theories, I will formulate a processing principle whose incorporation in theories of access will solve the problem. The principle will also shed another light on the issue of convergence in general.

6.3.2 Logogen Theory

Logogen theory, as developed by Morton (1969, 1979), was intended as a general theory of lexical access, i.e., access in both language comprehension and language production. Though most research in this framework has been devoted to language comprehension (particularly word recognition), the theory has been stimulating for production research as well. The present discussion will be limited to the latter. Lexical items are mentally represented as *logogens* in this theory. Logogens are devices that collect evidence for the appropriateness of a word. They are sensitive to information that may indicate the appropriateness of "their" word. All logogens are simultaneously active in collecting their specific information; that is, the logogen system is a parallel accessing device. In normal speech production, the information that activates logogens originates from the so-called Cognitive System, which is the repository of all conceptual, syntactic, and higher-order functions.

Each logogen has a threshold. As soon as the collected bits of evidence exceed the threshold, the logogen "fires" (i.e., makes the word's form available for use). This means that the logogen sends a phonological code to the so-called Response Buffer, where it has a short existence. At the same time, the logogen's activation level drops back to zero, and its threshold is temporarily lowered. The Response Buffer can use the phonological code to initiate a vocal response. It can also return the phonological code to the Logogen System. In that case, the logogen will be reactivated, and since the logogen's threshold is still low, it will normally fire again, which means that the same phonological code is returned to the Response Buffer. This may go on indefinitely. In this way a phonological code may be kept available for use, even when it cannot be immediately uttered. This will often be the case in fluent speech, where words have to be correctly ordered for output. The threshold of a logogen will vary with how frequently the logogen has been activated in the past and how recently that happened. Figure 6.4 is a schematic diagram of this part of the model.

The Cognitive System is equivalent to our Conceptualizer, except that it also generates syntactic information. Logogens are like lemmas, in that they are tuned to specific conceptual and syntactic information. The phonological codes they send to the Response Buffer are the lexical items' form information.

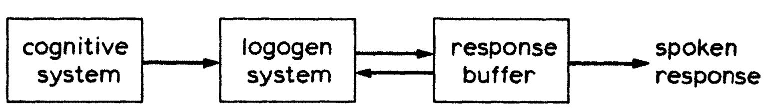

Figure 6.4
Diagram of the Logogen Model. Only parts involved in the generation of speech are represented.

Logogen theory has been used to account for a range of empirical findings in fluent speech production. Here are two examples:

(1) If a logogen's threshold is higher when it has been less frequently used in the past, one expects longer accessing times for low-frequency logogens than for high-frequency ones. This is because more evidence has to be accumulated to surpass a high threshold. Is it the case that low-frequency words are more often preceded by prelexical hesitation pauses than high-frequency words? Findings by Maclay and Osgood (1959) and Martin and Strange (1968) provide some indirect evidence. Function words (articles, pronouns, and other minor category words), which are of high-frequency, were far less often preceded by pauses than content words (such as nouns, verbs, and adjectives), which are, on the average, of far lower frequency. Beattie and Butterworth (1979) argued that it is the *predictability* of a word, not its frequency, that correlates with pausing. Since the predictability and the frequency of words are highly correlated in fluent speech, their effects must be statistically disentangled. Beattie and Butterworth found no frequency effect when predictability was taken into account. However, Levelt (1983) found a strong frequency effect but almost no predictability effect. The data were the prelexical filled-pause hesitations in subjects' descriptions of colored spatial networks (the task was discussed in subsection 4.4.2). Speakers used .. *er* .. more often before a low-frequency color word, such as *purple*, than before a high-frequency color word, such as *red*.

Predictability is again at issue in the second example.

(2) According to the model, a logogen also gathers relevant *contextual* information from the Cognitive System. The word *table*, for example, is easier to recognize when it follows the sentence fragment *The cup was placed on the*— than when it follows *They went to see the new*—. The first context gives a high *transitional probability* for *table*, whereas the transitional probability is low after the second sentence fragment. This latter fragment makes *film* easy to recognize (Morton 1979). Are there similar context effects in the *production* of a word? Lounsbury (1954) predicted that speakers would hesitate more at transitions of low transitional probability, and Goldman-Eisler (1958, 1968) showed this to be the case; it was subsequently confirmed by Tannenbaum, Williams, and Hillier (1965), by Butterworth (1972), and by Beattie and Butterworth (1979). (See Butterworth 1980b for a review.)

Though these two phenomena are in agreement with logogen theory, one should keep in mind that hesitation data can almost always be interpreted in multiple ways. In terms of lemma and word-form access, the word-frequency effect in example 1 above may be due either to accessing of the lemma or to accessing of the word form. This is not a relevant distinction in logogen theory, since the lexical item is accessed as a whole. Arguments for a form-based interpretation of the frequency effect can be found in Garrett 1982 and in Levelt 1983.

The context effect in example 2, moreover, may be due not to lowered thresholds for the logogen *table* or *film* in the appropriate context but rather to the greater availability of the concept TABLE or FILM in contexts where one is talking about placing objects or going out. This was Goldman-Eisler's interpretation of the data. She supported it with the finding that the sentence fragment *following* a target word was equally predictive of hesitation before that target word. Prelexical hesitation indicates, according to Goldman-Eisler, that the speaker is involved in complicated conceptual planning, and this may also cause the following words to be less obvious or expectable continuations of the "hesitant" word. In other words, the context effect on prelexical hesitation originates not at the level of accessing logogens but at the preliminary state of conceptual planning or message construction.

The major attractive property of logogen theory for fluent language production is its distributed control structure. All lemmas are simultaneously testing their characteristic features against the message or message fragments available. The first logogen that reaches threshold activation will "fire." This predicts that, normally, many more logogens are activated to some degree than just the one that is finally used in the sentence. This is all right as long as there is a unique solution for convergence. But does the system guarantee that, eventually, the correct logogen will fire? This property will not fall out naturally. Consider the hypernym problem. If the speaker intends to express the concept DOG, the Cognitive System will make available conceptual features relevant to the activation of the logogen for *dog*. But these features have as a subset those that will activate the logogen for *animal*. Why would *dog* fire but not *animal*? One answer could be that *dog* has a lower threshold than *animal* and so is quicker in reaching threshold stimulation. But then this should hold across the board for all hypernym relations: Hypernyms should have higher thresholds. This, however, cannot be true, since a word's hypernyms can be used with much more frequency than the word itself. Compare *dog* and *collie*. The former is a hypernym of the later; it is also of much higher frequency, and hence of lower threshold. A speaker intending to express the concept COLLIE will necessarily end up saying *dog*. The convergence problem has not been solved in logogen theory.

6.3.3 Discrimination Nets

Discrimination networks were originally proposed in computer models of verbal learning (Feigenbaum 1963), and have since then been widely used in artificial intelligence. Goldman (1975) was the first to construct a discrimination net for handling lexical access in an artificial language-producing system (which was called BABEL). In essence, his discrimination nets are binary tree structures. Each nonterminal node in the tree represents some predicate that is either true or false for the conceptualization at hand. Terminal nodes correspond to lexical items. The access procedure starts by running the test for the tree's root predicate. If it evaluates to "true" for the concept at hand, control moves to the node's left daughter; if it evaluates to "false", it goes to the right daughter node. The next test concerns the daughter node's predicate. The procedure is self-terminating; it iterates till a terminal node is reached. The lexical item at that terminal node is the system's lexical response to the concept.

A made-up example is presented in figure 6.5a. It discriminates among four lemmas which express pieces of furniture: *chair*, *stool*, *table*, and *bed*. If the speaker intends to express the concept STOOL, and the tree in figure 6.5a is his discrimination net, he will first run a test determining whether the concept involves the presence of a seat. This test will evaluate to "true". Control then shifts to the left daughter of the first test node. It represents the test whether the notion to be expressed involves a backrest. In the case of STOOL it will evaluate to "false". This brings the procedure to the terminal node that represents the lemma *stool*. In other words, the lemma is retrieved by executing a sequence of two semantic tests; this also holds for retrieving any of the other three lemmas in the network. A rather less symmetrical case is represented in figure 6.5b. It is a network discriminating among *table*, *chair*, *cabinet*, and *bed*. Here it involves only one test to

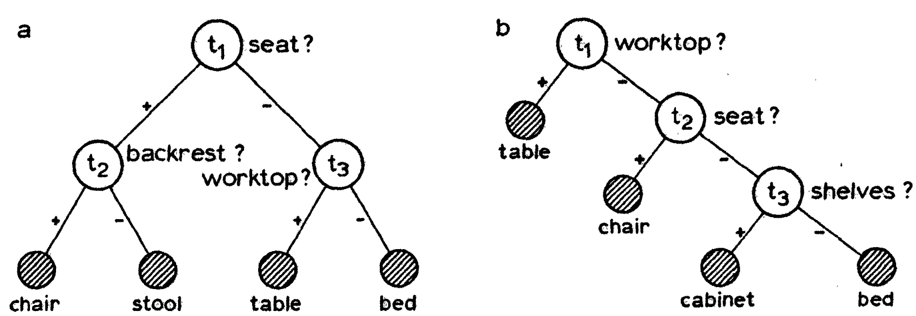

Figure 6.5
An elementary discrimination net for four pieces of furniture.

retrieve *table*, but three sequential tests to come up with either *cabinet* or *bed.*

Goldman did not claim psychological reality for his discrimination-net approach. Still, one might ask whether it is an attractive option for a psychological theory of lexical access in production. Well, it isn't. Let us consider some of the major drawbacks of this approach.

(i) Discrimination nets are sequential testing devices, and won't operate in real time. A few numbers may serve to highlight this problem. If the speaker's lexicon contains 30,000 content words, the speaker's discrimination net will contain 30,000 terminal nodes. How many sequential tests will, on the average, be involved to reach such a terminal node? This depends on the structure of the discrimination net. The optimal case is one where the structure is like that shown in figure 6.5a and where all lemmas involve an equal number of tests. Though this is possible only if the number of lemmas is a power of 2, it can be closely approached in other cases. For such a network, the average number of tests approaches $2 \log N$, where N is the number of lemmas. For a lexicon of 30,000 this number is about 15. In other words, it takes about 15 sequential tests to reach a terminal node in such a network. The worst case is a network like the one in figure 6.5b, where at least one outcome of any test produces a lemma. The average number of tests in this case is $0.5(N + 1) - 1/N$. For our example lexicon, this number is about 15,000. This extreme case may be judged unrealistic, but so is the optimal case. We know that concepts vary in complexity, and there will thus be variation in the number of tests needed to discriminate them; more complexity involves more tests. The average number of sequential tests needed to retrieve a lemma will, in short, be somewhere between 15 and 15,000. To do this between 2 and 5 times per second, a speaker must run his tests at a rate of between 30 and 75,000 per second. This is not very appealing for a real-time model of lexical access.

(ii) The access time for a lemma will be a monotonic increasing function of the number of tests. If, for example, figure 6.5b is accurate for the four lemmas involved, it will take less time to access *table* than to access *bed.* Nothing in the literature supports this prediction. Levelt, Schreuder, and Hoenkamp (1978) and Schreuder (1978) showed that the reaction times for choosing among four verbs of motion (*take along*, *pass*, *launch*, and *pick up*) in the presence of a visual event were not compatible with *any* of the possible discrimination nets. In naming tasks where a subject was free to call the perceived movement whatever he liked, the more generic verbs (such as *move*)—i.e., the verbs with fewer conceptual features—always had longer naming latencies than the more specific verbs (such as *launch*). On any reasonable theory, the former verbs involve fewer conceptual components than the latter ones and should therefore display shorter speech-onset latencies if the lexicon is organized like a discrimination net. In short, the theory makes the wrong kinds of prediction with respect to access times.

(iii) Discrimination nets of the kind proposed fail on the criterion of correct convergence. In particular, they cannot handle the hypernym problem, for quite principled reasons. If lemma *a* is a hypernym of lemma *b* (as are, for instance, *animal* and *bear*), then *a* and *b* cannot be represented on the same discrimination net. Why not? In order for *a* and *b* to be located at different terminal nodes on the same binary tree, there must be at least one test that discriminates between the underlying concepts A and B (ANIMAL and BEAR in the example). But since *a* is a hypernym of *b*, all tests on which A evaluates to "true" will also evaluate to "true" for B (everything that will be true for ANIMAL will be true for BEAR), and what is true for B is either true or irrelevant for A. In other words, there cannot be a test on which A and B receive opposing values. This is, indeed, a serious problem. The mental lexicon abounds with hypernym relations, but these cannot be represented by way of binary discrimination nets. Hence, they cannot be realistic psychological models, for purely structural reasons. For the same reason, they will fail as lexical representations in artificial natural-language generators.

(iv) It is, finally, an illusion to think that the mental lexicon is organized in the manner of a hierarchical taxonomy. Many semantic relations between lexical items are of a quite different nature, as is immediately obvious when one considers semantic relations between kinship terms, color names, names for days of the week, and so on (Miller and Johnson-Laird 1976; Lyons 1977).

6.3.4 Decision Tables

A major improvement over discrimination nets was proposed in Miller and Johnson-Laird 1976 (see also Steedman and Johnson-Laird 1980). Their *decision-table* proposal still involves a speaker who runs a set of tests in order to retrieve a lemma, but the tests can be run in parallel. The example in table 6.1 will serve to introduce the notion of a decision table.

A decision table is a matrix representing a set of IF/THEN or condition/action pairs. The rows in the upper half of the table represent the outcomes of semantic tests. They are the same kinds of tests as those discussed in the

tests (semantic conditions)	outcomes 1	2	3	4	5	6
		IF				
t1	+	+	-		+	-
t2	-	-	+	-	+	+
t3	+	-		+	-	
t4	+	+	-	+	-	+
lemmas (lexical actions)						
a	x					
b ←— THEN		x	x			
c					x	
d				x		
e					x	

Table 6.1
Example of a decision table.

previous section. If, for instance, the decision table concerns pieces of furniture, there will be tests for properties like "worktop", "seat", and "shelves". (A detailed example is worked out in Miller and Johnson-Laird 1976.) If the speaker considers expressing the notion STOOL, he will run all tests in parallel and come up with a value (+ or −) for each test. This pattern of results is called the outcome. Let us assume that the outcome for testing STOOL is the one encircled in column 2 of table 6.1. This outcome is a condition for an action specified in the lower half of the table. The x in column 2 is in the row of lexical action *b*; in the example that action is to retrieve *stool*. In other words, IF the outcome specified in column 2 arises, THEN lemma *b* is to be used. It should be noticed that certain cells in the upper half of the table are blank. This means that the corresponding test values are irrelevant. The action will be performed in case there is a match for all + and − values in the column, whatever the value on the "blank" test. We will return to this important property of decision tables shortly.

The lower half of table 6.1 also has some noteworthy properties. First, can there be two or more crosses in a single row, as is the case for row *b*? Yes; this is a natural way to represent homonymy, the same word expressing different concepts. For instance, columns 2 and 3 could represent outcomes for the notions BANK (of river) and BANK (financial institution), respectively. They would both occasion the lexical action of accessing the word *bank*. What about two crosses appearing in the same column, as in column 5 of the example? This should not be allowed. It violates the theoretical criterion of convergence. A single concept would lead to two different lexical actions. Decision tables can always be arranged in such a way that this does not happen. Can there be outcomes without lexical action, i.e., without a corresponding x in the lower table? This situation is exemplified by the outcome in column 6. It should be noticed that the number of possible outcome patterns increases rapidly with the number of tests in the table. There may, then, not be enough possible lexical actions to take care of each of these outcomes. Miller and Johnson-Laird (1976) talked about the "exponential specter." They tried to curb this threatening explosion by introducing "table-linkage," which will be discussed shortly.

But what does it really mean that there are outcome patterns for which there is no lexical action? Cases where there is a unitary concept without a corresponding lexical item are called "lexical gaps." Such cases are not disastrous for a theory of language production, since most conceptualizations we entertain are not expressible in single words. Under such circumstances the speaker will resort to a syntactic solution, i.e., he will create a phrase or sentence that expresses the concept. If the concept to be expressed is a dead animal, the tests will converge on the lemma *corpse*; if it is a dead tree, they will not converge on a lemma. The speaker will then resort to formulating the phrase *dead tree*. We could accommodate this by adding to the bottom of the table a row *f* standing for "syntactic action". All outcome patterns for which there is no lexical action will have xs in row *f*.

Of course, this leaves unanswered the difficult question of how a speaker partitions a message into "unitary concepts"—i.e., conceptual fragments for which there will normally be words in the language.

Though the "exponential specter" doesn't seem to be much of a problem for a theory of the speaker, there may be other reasons to consider reducing the size of a decision table. Anyone who has ever tried to construct a decision table for a set of lexical items will have noticed that almost every new item in the set requires the addition of a new test. Even if one assumes that, on the average, a new test is needed for every tenth item, one will have 3,000 tests for a 30,000-word mental lexicon. This implies that the IF-statement for the retrieval of each lexical item is a pattern of 3,000 test outcomes, and this is quite far from the ideal case where, for each item, one or a few tests would suffice for its identification.

Table linkage is a way of cutting up the large decision table into a set of smaller ones. There will be a table for each semantic domain: one for furniture terms, one for verbs of motion, one for kinship terms, and so on. Each of these tables involves a small number of tests, so the number of outcome patterns will never be excessive. The tables are linked in that each table can make reference to one or more other tables (called "successor tables"), depending on the outcome pattern of the tests. Testing begins at some privileged table. If an outcome pattern results for which there is no lexical action, control is transferred to a successor table. This will iterate until some table in the sequence comes up with a lexical action. But notice that this is seriality again. How many successive tables a speaker will, on the average, have to go through in order to come up with a lexical action will depend on the way the tables are linked.

How does the decision-table approach handle convergence? Is the correct item indeed found? The touchstone is again the handling of hypernyms. In the previous subsection we observed that discrimination nets cannot represent hypernym relations. But decision tables can. There is an example in table 6.1. Word *d* is a hypernym of word *a* (as in *diplomat* and *ambassador*). How is this visible in the table? Compare the outcomes of columns 1 and 4. They are identical except for the cells in the upper row; whereas column 1 displays a +, column 4 has a blank. This means that outcome 1 implicates outcome 4, but not conversely. That defines the hypernym relation. Consider the example: If someone is an ambassador, he or she is a diplomat; but a diplomat need not be an ambassador. There are diplomats who do not share the special feature of ambassadors of being a country's representative. Test 1 concerns this feature. It is irrelevant for deciding whether a person is a diplomat (the blank cell in column 4), but it is essential for deciding whether a person is an ambassador.

But though hypernym relations are representable in decision tables, they create a violation of our convergence criterion. The obvious reason is that if the column for the specific term (i.e., *ambassador*) is satisfied, the column for its hypernym is by necessity also satisfied. As a consequence, there will be *two* lexical actions if the more specific concept (e.g., AMBASSADOR) is to be expressed; the hypernym will also be accessed. Since many words have a whole hierarchy of hypernyms (*ambassador–diplomat–person–human*–etc.), there will be permanent cluttering of lexical actions. Later we will see that even when there is cluttering (as in blends of words) it hardly ever involves a term and a hypernym.

Decision tables are a major improvement over discrimination nets in that they allow for real-time parallel processing. However, in their present form they defy proper convergence, and this is a serious drawback.

6.3.5 Activation Spreading

Could there be a connectionist account of lexical access? In chapter 1, where we discussed connectionism as a potentially powerful formal language for the description of parallel processing, we went through an example of word-form access (see figure 1.2). The top-level nodes in that example were lemmas, and we traced the activation spreading down to the level of segments or phonemes. Substantive proposals have been made with respect to this process (see especially Dell 1986, 1988); we will turn to these in chapter 9. Our present problem concerns the preceding phase of activation: How are the lemma-level nodes activated? At the time of this writing, the literature contains no serious proposals.

One approach could be to distinguish a level of nodes representing conceptual components. Returning to our furniture example, we could think of such properties as having a worktop or shelves, being a seat, and so on. In a connectionist account, such properties can be represented as input nodes. Let us call them A, B, C, etc. On this account a concept is nothing more than an activated collection of such nodes. This is called a *distributed* representation. The output nodes will be lemmas: *stool, chair,* and so forth. The simplest account is one in which these are the only two levels of nodes, with each input node connected to all the output nodes:

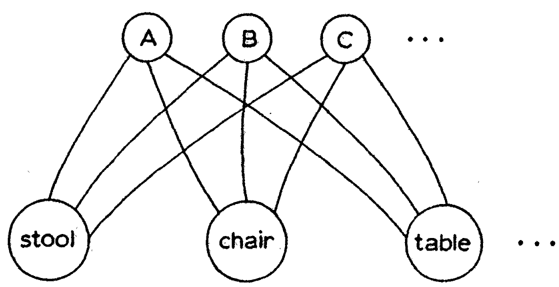

We can now further specify the system by defining the characteristic state of activation of the input nodes for different concepts (this may or may not involve a binary activation scale), the output function of each node (i.e., how the degree of the node's activation determines the level of its output), the input function of the lemma nodes (i.e., how the propagated activity impinging on an output unit summates to activate that unit), and so on. We will not go through this exercise. In fact, there are only casual suggestions in the literature as to how this should be done for the modeling of lexical access.

We will, rather, limit the present discussion to two observations. The first is that, as a parallel-activation account, activation spreading is an excellent candidate for taking care of real-time limitations. The number of steps from input nodes to output nodes is one or (if there is a level of hidden nodes) two. The second observation is that, so far, no one has even begun to account for convergence. Not only is the structure over which the activation should spread not known; there is also no principled approach to solving the hypernym problem.

6.3.6 Toward a Solution of the Hypernym Problem

As was observed above, convergence is a major stumbling block for all theories of lemma access. In particular, none of the existing theories involves a correct treatment of the hypernym problem. In the following, three principles will be proposed that, if incorporated in a processing model, guarantee correct convergence for a term and its hypernyms.

The first principle will use the notion of *core*. A lemma's core meaning or conceptual core condition is a privileged, "most salient" meaning component. There is an empirical basis for assuming the existence of core meanings. Miller (1969) suggested a "negation test" procedure for determining a lemma's core, and this procedure was used by Levelt et al. (1978) and Schreuder (1978) to analyze verbs of motion and by Noordman (1979) to study kinship terms.

The idea underlying the negation test is that negation will, when there is no restriction on the set of alternatives, affect only a word's core. Here is an example from the study of verbs of motion: Subjects were asked to complete sentences such as *They do not ski, but they....* No subject ever completed this sentence with a verb like *breathe* or *think*. The most frequent response was *skate*. Subjects apparently apply minimal negation, in that most of the meaning of *ski* is preserved under negation; like *ski*, *skate* denotes some form of human locomotion over a frozen surface, involving some instrument attached to the feet. Only the character of the instrument is different. The instrument's character is the only conceptual component changed under negation. By this test it constitutes the core component of the lemma for *ski*. Miller called the unchanged part (in the present instance the meaning shared by *ski* and *skate*) the lemma's *presupposition*.

Though the negation test is helpful in demonstrating the notion of core meaning, it is not a foolproof procedure for finding an item's core. When the set of alternatives to the negated item is restricted in the situation of discourse, the test easily fails. If you happen to know that John is either skiing or working, then the completion *John isn't skiing, he is working* is quite natural.

Let us now turn to the first principle. It is a stronger version of what was introduced earlier as E. Clark's (1987, 1988) principle of contrast: "Every two forms contrast in meaning." This principle denies the existence of full synonyms in a language. Restricting this principle to core meanings, we have the uniqueness principle.

The uniqueness principle No two lexical items have the same core meaning.

In other words, there are as many core meanings as there are lexical items. This principle captures the intuition described above in the discussion of decision tables: Every new item added to a table seems to require a new idiosyncratic test. The following two principles formulate the role of core meaning in lexical access.

The core principle A lexical item is retrieved only if its core condition is satisfied by the concept to be expressed.

This is not a very strong principle, but it is a first step toward solving the hypernym problem. The reason is that, by the uniqueness principle, a term and a superordinate or hypernym never have the same core. To keep to the above example: Both *skiing* and *skating* are forms of *gliding*; hence, *glide* is a superordinate of these verbs. When the negation test is applied to *glide*, it will yield verbs like *stick* where there is no smooth continuity in the motion. The core of *glide* is, apparently, this manner feature of smooth continuity. Hence, *ski* and *glide* have different core conditions. If the speaker now intends to express the notion GLIDE, the core principle guarantees that no subordinate term will be used. The speaker will not erroneously retrieve *ski*, since that word's core condition is not met by the concept GLIDE.

But the reverse may still happen. If the notion SKI entails the notion GLIDE, then if SKI is the concept the speaker intends to express, the core condition of GLIDE is necessarily satisfied by that concept. Hence we need an additional principle to prevent retrieval of the hypernym:

The principle of specificity Of all the items whose core conditions are satisfied by the concept, the most specific one is retrieved.

This principle prevents the retrieval of *glide* when SKI is the concept. Since the meaning of *ski* entails the meaning of *glide*, *ski* is the more specific word item. It will be the one retrieved.

These principles are reminiscent of Grice's maxims of quantity. One should not say *glide* when one intends to express the notion SKI. That is like (truthfully) saying "I have two sisters" when one has in fact three.

Lexical access, then, involves essentially recognizing the most entailing predicates in the concept and finding the unique lemmas that have these as their core conditions (when they exist). We began the present section by considering the "ideal" case of a one-to-one relationship between concepts and words. In that case, retrieving a word would be a simple mental reaction. Under the present assumptions this ideal situation is closely approached. Each lexical item can be seen as a testing device for the realization of its own core condition, roughly as in logogen theory. Implementation of the principle of specificity will guarantee correct convergence.

6.4 Failures of Lemma Access

6.4.1 A Taxonomy of Causes

Anthony Cohen, one of the initiators of speech-error research in modern psycholinguistics (Cohen 1965, 1968), called speech errors "blabbers." That speech errors are thoughtless or even indiscrete exposures of the underlying formulating machinery was not only the view of the pioneers in this field, Meringer and Mayer (1895) and Freud (1904); it is the main impetus for what is now one of the most flourishing empirical methods in the study of the formulation of speech. There are three major anthologies in this field: Fromkin 1973, Fromkin 1980, and Cutler 1982a. Cutler 1982b is an exceptionally complete bibliography.

The three classes of speech errors to be discussed in this section have different etiologies in the production process. All are due to derailments in the retrieval of lemmas, but the mechanisms involved are different. We will have to distinguish two major causes of word errors: *conceptual intrusion* and *associative intrusion*.* Conceptual intrusion occurs when lemma selection is disturbed by the simultaneous activity of two or more concepts. Associative intrusion occurs when lemma selection is interfered with by associations between lemmas, better known as *word associations*. We will assume that associative relations between lemmas are *direct* relations, i.e., not mediated by anything else (see subsection 6.1.2). A further distinction should be made: The intruding element may or may not be related in meaning to the intended element. Since the latter distinction is independent of the former, we end up with a four-way taxonomy (figure 6.6).

* It is not a matter of concern here *why* a person happens to have an alternative thought in mind, or *why* he built up a particular association in his lifetime. See Cutler 1982a for a distinction between these more remote causes of error and the proximal causes and mechanisms we are dealing with here.

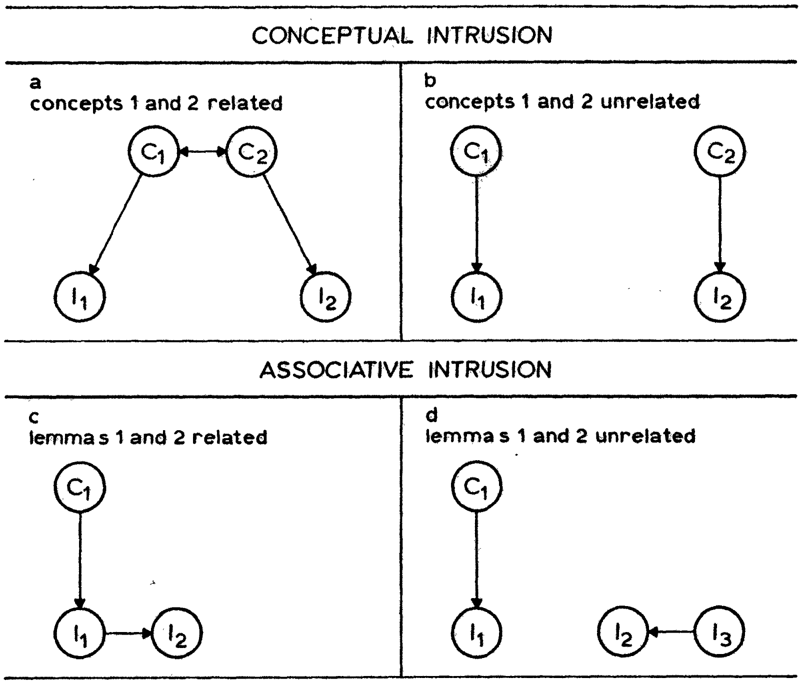

Figure 6.6
Conceptual and associative intrusion in lemma access. The intruding concept, C_2, is either (a) meaning related to the intended one, C_1, or (b) not. Also, the intruding lemma, l_2, may (c) or may not (d) be meaning related to the intended lemma, l_1.

In the figure the intended concept, the one to be expressed, is always C_1. The appropriate lemma to express this concept in the context of use is always ℓ_1. Conceptual intrusion is caused by the presence of another concept, C_2. It activates its own lemma ℓ_2, and this interferes with the activity of ℓ_1. Concepts C_1 and C_2 can be closely related in meaning, as in figure 6.6a, or they can be unrelated, as in figure 6.6b. Associative intrusion is caused by one lemma's directly activating another one. In figure 6.6c the intended lemma, ℓ_1, activates a closely associated lemma, ℓ_2, which may then interfere with ℓ_1's activity. In figure 6.6d the intruding lemma, ℓ_2, is unrelated to ℓ_1; it is activated by the independent but still active lemma ℓ_3, to which it is closely associated. Let us now make these notions more concrete by applying them to the three classes of speech error: blends, substitutions, and exchanges of words.

6.4.2 Blends

In a word blend two words are fused into one. Two lemmas are retrieved, which compete for the same syntactic slot. Semantically speaking, there are two kinds of word blends. The first kind involves words of similar meaning;

the second kind involves what we will call distractions. Here is a set of examples of the first kind:

(1) The competition is a little stougher [stiffer/tougher] (Fromkin 1973)

(2) Irvine is quite clear [close/near] (Fromkin 1973)

(3) To determine watch [what/which] (Fromkin 1973)

(4) At the end of today's lection [lecture/lesson] (Garrett 1975)

(5) and would like to enlicit [enlist/elicit] your support (Garrett 1980a)

(6) "Lexikon" ist ein Terminus, der sich eingebildet [herausgebildet/eingebürgert] hat (*"Lexicon" is a term that has emerged/has become familiar*) (Bierwisch 1982)

The major observation about this class of word blends is that the two words are roughly equivalent in meaning in the context of the message as a whole. Is this relatedness of meaning a conceptual relation, as in figure 6.6a, or a lemma association, as in figure 6.6c? The empirical evidence is in favor of the former interpretation. In the case of lemma association one would expect the two blended words to be close associates. Among the closest associates in word-association tables (see, for instance, Palermo and Jenkins 1964) are antonyms. But antonym blends are highly exceptional. Meringer and Mayer's (1896) sample contains only two: *wehr* [*weniger/mehr*] and *Beneinung* [*Bejahung/Verneinung*]. Fromkin's (1973) list of 65 blends contains no antonym blends, and Hotopf (1980) also presents numbers testifying to this exceptionality. The few cases observed may have originated at the conceptual level. The speaker who produced *wehr* may have had two closely related concepts in mind, namely LESS (activating the lemma for *weniger*) and NOT MORE (activating the lemma for *mehr*). Antonyms are semantically related, but only as far as their presuppositions are concerned; their core meanings are mutually exclusive. Only equivalence or similarity of core meaning seems to count for this category of blends.

The latter restriction also excludes blends of a term and a hypernym. In subsection 6.3.6 it was proposed that a term's core meaning is always different from the core meanings of its hypernyms. And indeed, as Hotopf (1980) shows, such blends hardly ever occur. There are, for instance, only three potential cases of word/superordinate merges among Fromkin's (1973) 65 published blends: *dealsman* [*dealer/salesman*], *hegraines* [*headaches/migraines*], and *swifting* [*shifting/switching*]. I say "potential" because the two core meanings may have been *functionally* equivalent in the contexts where these blends occurred (Bierwisch 1982).

It is not characteristic of blended words to be close associates of one another. It may just happen to be the case, but it may equally well not be the case. The etiology of blends of related words, therefore, seems to be that the message (fragment) itself contains a certain ambivalence with respect to two equally appropriate concepts (there are "alternative plans," as Butterworth 1982 called this kind of interference). These closely related concepts trigger their lemmas almost simultaneously. Both lexical items, whether word associates or not, are retrieved, and they are both inserted in the same surface-structure slot. This induces their word forms to become blended at the level of phonological processing. In short, these blends are due to conceptual intrusion.

The same can be said about blends-by-distraction. The speaker who is in the train of expressing concept C_1 through lemma ℓ_1 may get an intruding thought by distraction or whatever Freudian cause (Butterworth called this the "competing plan" kind of interference). A concept C_2, not part of the message, will independently activate a lemma ℓ_2. If both lemmas fire simultaneously, a blend can arise. Example 7 is of this sort.

(7) Dann aber sind Tatsachen zum Vorschwein [Vorschein/Schweinereien] gekommen (*But then facts came to appearance/filthinesses*) (observed by both Mayer and Freud, and cited in Meringer and Mayer 1895 and Freud 1904)

The speaker of these famous words agreed to the two observers that he was thinking of "filthinesses" when he was speaking his rather neutral sentence. The intruding thought (C_2) activated an intruding lemma (ℓ_2), and the two word forms merged. Clearly, the two concepts involved were unrelated, and there was no association between the two words. Blends-by-distraction are of the type shown in figure 6.6b.

Two blended words are usually of the same syntactic category, as is the case in all of examples 1–7. On first view this is surprising; if lemmas of content words are triggered solely on the basis of their conceptual specifications, syntactic category information should be irrelevant, and lemmas of different categories but similar meaning might be prone to amalgamation as well. One should, however, imagine what would happen if two such lemmas were to compete. They cannot appear in the same phrasal environment because of their category difference, and at this point one of two things may happen. The first one is that the already existing phrasal restrictions are such that only a lemma of a particular category can be accepted for insertion. In that case, the other lemma will be unfit for use and without effect for further processing. One can say that it is "filtered out" by subsequent syntactic processing. If the phrasal environment is not sufficiently restrictive to filter one lemma out right away, a second thing may happen: Each lemma may trigger the construction of its own phrasal environment, and some phrasal blend may occur. These cases are exceptional; the following example, reported in Fay's (1982) corpus of sentence blends, may be due to such a process:

(8) And that is how I got interested into it [into/interested in]

The notion of filtering out is also important for the explanation of other aspects of word blends. Many word blends are words themselves; this is, for instance, the case for examples 2, 3, 4, and 6 above. Dell and Reich (1981) provided statistical evidence that there is a slight overrepresentation of real words among blends. This may be the result of filtering at a later stage, namely when word forms are generated. Selective filtering at that stage may also be the cause of blends between words that are similar in word form, as in *herrible* [*terrible*/*horrible*] (see Butterworth 1982 for an analysis of this phenomenon). These issues are discussed further in chapter 9.

Though the blends discussed in the present subsection are due to conceptual interference of one kind or another, there are also blends with different (for instance, phonological) etiologies. See Butterworth 1982, Harley 1984, and Stemberger 1985a for a wider discussion of blending.

6.4.3 Substitutions

Word substitutions can be of various sorts. The following are some examples of the selection errors that concern us here:

(9) He's a high–low grader [low → high] (Fromkin 1973)

(10) Don't burn your toes [fingers → toes] (Fromkin 1973)

(11) Der hat so'n Ding geheiratet–ich meine geerbt [geerbt → geheiratet] (*He married–I mean inherited–this thing* [*inherited* → *married*] (Bierwisch 1982)

(12) I just put it in the oven at a very low speed [temperature → speed] (Garrett 1982)

(13) Met een blauw vlakje, een blauw rondje aan de bovenkant [rondje → vlakje] (*with a blue spot, a blue disk at the upper end* [*disk* → *spot*] (Levelt 1983)

This set of examples is somewhat biased, in that no substitutions in which the target word and the substitution are unrelated are included. They do occur. However, we will first look into the most frequent type of substitution, which involves some semantic relation between the target and the error. This means that we are first dealing with cases a and c in figure 6.6. The question then is whether the intrusion is conceptual or associative.

The most frequently observed cases involve antonyms or other cases of semantic opposition, as in example 9, or co-hyponyms (i.e., words from the same semantic field), as in example 10. This is a pattern one typically encounters with word associations: similarity in presupposition, but difference in core. Fromkin's (1973) list of word substitutions is again a goldmine here; one finds substitutions such as *last → first*, *wine → beer*, *later → earlier*, *come → gone*, *little → much*, and *sun → world*. There are also more remote co-hyponyms where the common class is less immediate, as in examples 11 and 12. Hotopf (1980) claims that these cases more or less exhaust the semantic relations to be found in word substitutions.

These facts show that word substitutions, unlike blends, often reflect associative relations. This betrays an etiology like that in figure 6.6c. In *don't burn your toes*, for instance, the concept FINGERS activates the lemma *fingers*, which has *toes* as a close associate, which it activates. For some reason, the activation of *toes* reaches threshold before *finger* "fires" and the accident is created.

What could cause the associate lemma to become available before the target lemma? An obvious suggestion is that the associate has a lower threshold in the sense of logogen theory (subsection 6.3.2). According to logogen theory, high-frequency words have lower thresholds than low-frequency ones, and one wonders whether the substituting words in association-caused substitutions are of higher frequency than the target words. This would predict that a high-frequency word is more likely to substitute for a low-frequency one than inversely. I substantiated this by analyzing the word substitutions listed by Fromkin (1973), using the word-frequency tables of Kucera and Francis (1967). In the list of substitutions there were 23 cases where the substitution was clearly an associate of the target word (e.g. *question → answer*, *east → west*) and for which word frequences were available. Of these, 17 cases were in the predicted direction; only 6 went the other way. This difference is significant at the 0.02 level. No frequency difference between target and substitute can be discerned if *all* substitutions are taken into account. It only holds for those errors that can be interpreted as intrusions by a word that is *associated* to the target word. The demonstrated frequency effect does not contradict Hotopf's (1980) statistical finding that the substituting word tends to be of the *same* frequency class as the target word. That finding is not surprising, insofar as, quite generally, a word and its first associates are of the same frequency class. The demonstrated frequency difference holds *within* such frequency classes.

The next question is whether we can exclude *conceptual* intrusion as a cause of word substitutions. Surely examples 11–13, though of a less frequent sort, do not involve very close associates. What about words that are conceptually equivalent in the context? Synonyms and superordinates (e.g. *animal* for *dog*) do not appear in Hotopf's data. Garrett (1982) also called attention to these missing cases. As we have seen, synonyms (or, better, equivalent terms in the context) are the normal case for blends; why are they missing in substitution data? This may not be more than an artifact. If A and B are near-synonyms in the message context, it really doesn't matter which lemma is triggered first. In both cases the selected word will be appropriate, and nobody will ever notice that there was a "race" between the two activation processes. In blends, however, A and B get merged, and this will strike the attentive ear as a slip. In substitutions nothing so striking happens; we cannot exclude that a near-synonym may substitute for the target word.

Can a similar argument account for the near absence of superordinates (hypernyms) in word substitutions? At least in part. How can one know that a speaker should have said *dog* instead of the actually used word, *animal*? It will often go unnoticed, but the speaker may reveal it by making a correction. An anaysis of speaker's spontaneous repairs (Levelt 1983) shows that it is not at all unusual for speakers to replace a word by a more specific one. Example 13 is such a case. The speaker was describing visual patterns composed of colored disks connected by black arcs. In this context a disk (*rondje* in Dutch) is clearly a spot (*vlakje*), but the inverse need not be the case. *Spot* is a superordinate of *disk*. Still, even in these cases it is undecidable whether the concept triggered a hypernym. One cannot know whether the concept DISK was already active at the moment that *spot* was used; it could have been an afterthought. In that case, example 13 is not a substitution at all. In short, we can neither confirm nor deny that hypernym substitutions occur. But if, as we suppose, word association is a major cause of substitutions, one would expect substitutions by superordinates to occur. Word associations like *dog* → *animal* are not infrequent.

So far, we have considered substitutions involving some semantic relation, and we have tried to interpret them as either of type a or of type c in figure 6.6. We found that there is good evidence for an etiology as in c, associative intrusion; there may or may not be conceptual intrusion. What about semantically unrelated intrusions? Can one find cases like b and d among word substitutions? Two cases from Fromkin 1973 that seem to be of type d are the following:

(14) a branch falling on the tree [roof → tree]

(15) Q: when are you going to have the ale?

A: with the beer [dinner → beer]

In example 14, *tree* is not meaning-related to the target word *roof*, but it is a close associate of *branch*. The constellation is as in figure 6.6d, where the intended concept (C_1) is ROOF and the intended lemma (ℓ_1) is *roof*. The intruding lemma (ℓ_2) is *tree*, which is a close associate of still another lemma *branch* (ℓ_3). Similarly, in example 15 *beer* (ℓ_2) is associated to the interlocutor's *ale*, which became also activated in the speaker's mind (ℓ_3). So, recently activated lemmas other than the target word can also cause intrusion by a nonintended but associated lemma, and this case is not at all infrequent in the published data (see especially Harley 1984).

Example 15 shows that the priming may be due to a word *perceived* by the speaker. Recall the experimental data collected by Levelt and Kelter (1982) and presented in subsection 4.2.5. When shopkeepers were asked the question *At what time do you close?*, the answer was more likely to contain *at* (e.g., *At six*) than when the question was *What time do you close?* (typical answer: *Six o'clock*). This is a case of *identity priming* in production; the interlocutor's use of a word (*at*) increases the probability that the speaker will use that word in the next turn.

Though associative intrusion of the sort presented in figure 6.6d is clearly a cause of many word substitutions, one cannot exclude the possibility that other cases are due to *conceptual* intrusion, as in figure 6.6b. Example 16, also from Fromkin's sample, is probably of this sort.

(16) he's not that happy in Hawaii [Illinois → Hawaii]

The speaker was already thinking of Hawaii, which was going to come up in the next sentence. This example forms a natural bridge to the last kind of access errors to be discussed here: word exchanges.

6.4.4 Exchanges of Words

It was suggested earlier that word exchanges result from different message fragments' being active at the same time. Two characteristic examples from the many that have been published suffice as demonstrations:

(17) Well you can cut rain in the trees [rain ↔ trees] (Garrett 1982)

(18) This spring has a seat in it [spring ↔ seat] (Garrett 1980a)

In word exchanges it is no more than accidentally the case that the two words are close associates. They typically express different concepts that are both about to be formulated, mostly as parts of the same sentence. (Butterworth 1982 calls this kind of interference "plan internal.") Example 16 may have been a case of a beginning exchange between two consecutive sentences. This puts word exchanges in the category represented in figure 6.6b. Word exchanges are the clearest evidence available for parallelness of the second type discussed in subsection 6.3.1, i.e., the simultaneous accessing of different lemmas by different fragments of the message. Such parallel processing probably contributes to the speed and fluency of speech; it apparently also creates some accident-proneness.

In word exchanges of this type the two words are always of the same syntactic category. This is what one would expect if the insertion of the lemma in the developing surface structure were to require a fitting syntactic category. Garrett (1980a) observed that the exchanging words belong to different phrases and play similar roles in their respective phrases. It is hard to define what a similar role is, but something functional seems to be at stake. The elements are often both heads of phrase, and they tend to be similar in the thematic arguments they express or in the grammatical functions they fulfill. We will return to these issues in the next chapter.

6.5 The Time Course of Lexical Access

6.5.1 Stages of Access

How rapidly is an object's name retrieved? When speakers are shown pictures of objects and asked to name the objects as quickly as possible, it takes some 600 to 1,200 milliseconds from picture presentation to the initiation of the vocal response. This speech-onset latency is the result of several stages of processing, which may or may not show some overlap in time. The speaker will first process the picture visually. This involves extracting visual features, borders, corners, shades, overlaps, foreground/background relations, and so on. The next stage will be to categorize the visually emerging object as a car, a table, a clock, or whatever the kind of object presented. This category or concept is the occasion for retrieving the corresponding lexical item: *car*, *table*, etc. We are not sure whether the latter stage involves one or two steps, but at any rate both the lemma and the form information will be retrieved. Finally, the articulatory response has to be prepared; only then can the overt naming be initiated.

It is, of course, not the case that all accessing of content words in speech passes through all of these stages. The visual processing and categorization is not an essential part of lexical access. There are other ways of conceiving of a concept to be expressed. Still, studies of the time course of lexical access have almost exclusively been studies of naming—the naming of objects or of visual relations. The present section is concerned with these studies.

The speech-onset latency in naming depends on several factors which exert their effects in one or another of the conjectured stages. There are visual effects, conceptual effects, word-frequency effects, and so on. We will first discuss latency effects, which are largely caused during the visual and conceptual stages of processing; then we will turn to effects that are due to lexical accessing proper. Next we will consider the issue of lemma versus form retrieval; are they two successive steps or rather two aspects of the same retrieval operation? Finally, a few remarks will be made about potential overlap between the stages of categorization and lexical access.

6.5.2 Visual Processing and Categorization

How is a visually represented object categorized? This question deserves a book in itself and will not be extensively discussed here. But naming latencies do reveal preferences of categorization. Rosch and colleagues (see especially Rosch, Mervis, Gray, Johnson, and Boyes-Braem 1976) introduced the notion of *basic object level*. It is the preferred level of categorization of visual objects: apple, shoe, chair, car, etc. There is a superordinate level of categories (fruits, clothing, furniture, vehicles, and so on), and there are subordinate categories for each of the basic-level terms (cooking apple, pump, Ottoman, convertible, etc.). Several studies, nicely reviewed by Seymour (1979) and including some of his own, have shown that naming latencies are shortest when a speaker uses basic-level terms for the objects he is presented with. When a subject receives a mixed set of fruits, furniture items, vehicles, and so on, and is required to say *fruit*, *furniture*, *vehicle*, etc. (that is, to use superordinate categories instead of basic-level terms like *apple* and *chair*), the naming latencies are substantially longer. But latencies are also longer when subordinate terms have to be used. This finding can probably not be explained by the lower frequency of usage of non-basic-level terms. It shows, rather, that there are preferred ways of categorizing visually presented objects.

The process of visual object categorization can further be traced by the so-called priming paradigm. The picture to be named is preceded by another visual stimulus (the prime) which can entertain some relation to the target object, and this relation may effect the naming response to the target. Flores d'Arcais and Schreuder (1987) used a picture-naming task where the prime was another picture. It appeared 450 milliseconds before the presentation of the target and was visible for 200 milliseconds. The prime could be related to the target in various ways, among them the following

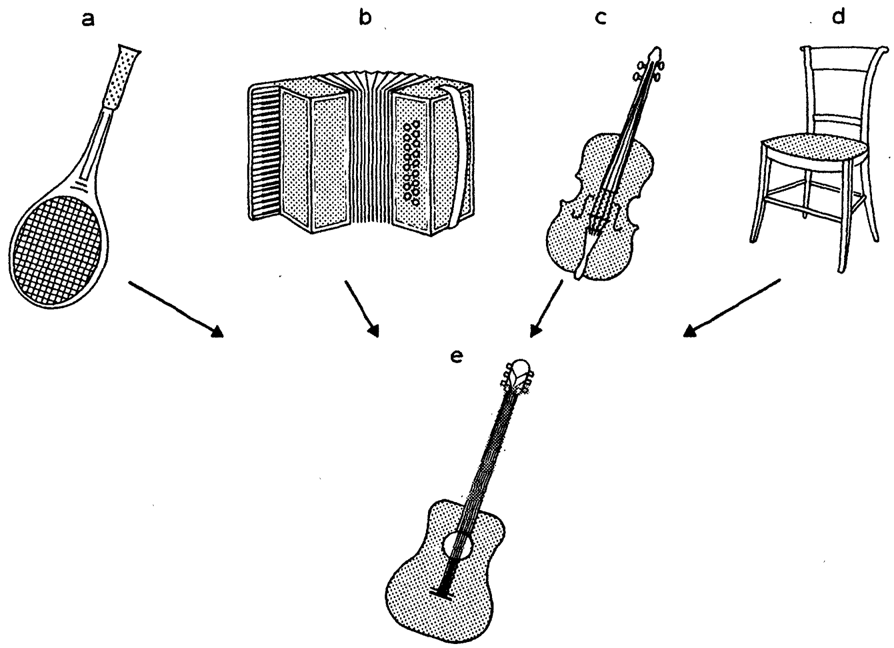

Figure 6.7
A target picture (e) and four prime pictures entertaining different relations to the target: only perceptually related (a), only functionally related (b), both perceptually and functionally related (c), and unrelated (d).

two: (i) the prime and the target could be *perceptually similar*, i.e., could share certain visual features; or (ii) they could be of the same superordinate-level category, making them *functionally similar*. By realizing both, one, or none of these two, four kinds of prime result. An example of a target picture and instances of these four kinds of prime are presented in figure 6.7.

Flores d'Arcais and Schreuder conjectured that there are two "priming routes," depicted here in figure 6.8. The first route involves visual perceptual features shared by the prime and the target picture; the second one involves functionally shared features. The bottom part of figure 6.8 is similar to figures 6.6a and 6.6b, but the figure represents the additional levels of perceptual and functional feature nodes. The target picture activates a set of visual features, among them P_1 and P_i. The prime picture activates another set of features, among them P_i and P_m. Hence, P_i is a perceptual feature activated both by the prime and by the target. If the two pictures are very similar, they share a large number of perceptual features. The target picture's features will activate the corresponding concept node, C_1, and the prime's perceptual features will activate C_2. The first, visual priming route goes from prime picture to shared perceptual feature (P_i in the figure) to target concept (C_1). When the prime preactivates P_i, the

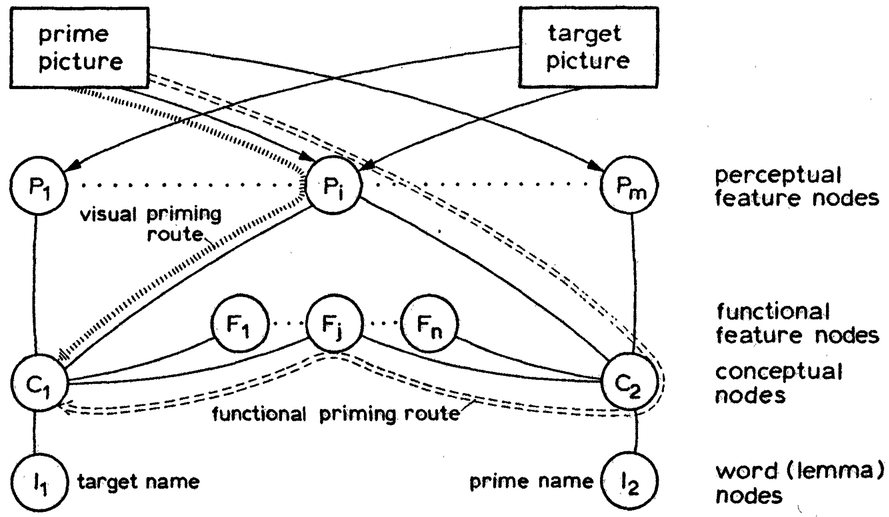

Figure 6.8
Two priming routes. There is direct visual priming, not involving the prime's concept node, and there is functional or conceptual priming which does involve the prime's concept node.

speaker will be quicker in categorizing the target object. This, in turn, will speed up the naming response. Notice that this route does not involve C_2; the prime concept need not be activated.

This is different for the second, functional priming route. The prime's visual features activate the prime's concept node, C_2. This, in turn, shares a functional (or conceptual) feature, F_j, with the target concept (both being musical instruments, for instance), through which activation spreading may take place. The prime/target pairs in figure 6.7 make it possible to test which of these routes is the more effective in speeding up the naming response.

The results are clear-cut. When a prime like that in figure 6.7a was used (i.e., one that shares visual features with the target), the naming response was sped up by some 35 milliseconds as compared to when an unrelated prime was used (figure 6.7d). The visual route, therefore, is quite effective. A prime like 6.7b, however, was not very effective. Naming latency decreased by a nonsignificant 16 milliseconds. The best prime was 6.7c, which has both perceptual and functional features in common with the target. It sped up the naming response by 125 milliseconds.

Huttenlocher and Kubicek (1983), in a very similar experiment, found substantial priming by functionally related pictures (e.g., a drum and a guitar). Some of these pairs were also perceptually related (e.g., a violin and a guitar). In those cases, however, they had been depicted from a different visual perspective (or angle) so that low-level visual processing of the target would not be aided by the prime. These cases were independently studied by Flores d'Arcais and Schreuder, who found that there was substantial priming (47 milliseconds) between functionally related objects when they were physically similar (e.g., a guitar and a violin), even when the visual perspective was different.

Many authors have not used a prime *picture*, but a prime *word*. The printed word doesn't share visual features with the target picture; this excludes the visual priming route. This work is well reviewed in Glaser and Düngelhoff 1984. Here we will consider a major set of results obtained by these authors. An important variable in their study was the delay between presentation of prime word and target picture. This is called the *stimulus-onset asynchrony* (SOA). If the prime precedes the target, the SOA is negative; if it follows the target, the SOA is positive. In all cases the subjects' task was simply to name the picture, and their speech-onset latencies were measured. It is, of course, possible that a prime word can have an effect on naming latency even if it is presented shortly after the picture (i.e., with a positive SOA).

What kinds of prime words did the authors use? First, the prime could be conceptually related to the target. If the picture showed a church, the prime word could be *house*. This is a *related* prime word. Second, the prime word could be *unrelated*. If the picture was a church, an unrelated prime could be *car*. Third, the prime word could be the picture's name. If the picture to be named displayed a church, the *identical prime* was the word *church*. And fourth, the prime could be *neutral*. In this case the "word" presented was just the row xxxxxx. The latter condition was taken as the baseline condition. If a prime word caused picture naming to be faster than in the neutral condition, there was *facilitation*. If naming was slower than in the neutral condition, there was *inhibition*. Figure 6.9 shows the inhibition and facilitation observed for related, unrelated, and identical primes given at different SOAs.

It is clear from figure 6.9 that the identical prime generally caused facilitation, even if it came as late as 200 milliseconds after the target picture. However, the other two kinds of prime—the related and unrelated primes—generally produced an inhibitory effect; their presentation slowed down the naming response. The strongest inhibition resulted from the related prime. The latter result seems to contradict the above-mentioned findings of Huttenlocher et al. and Flores d'Arcais et al. They found that a related picture prime caused shorter naming latencies than an unrelated picture prime. However, both these studies used negative SOAs with an

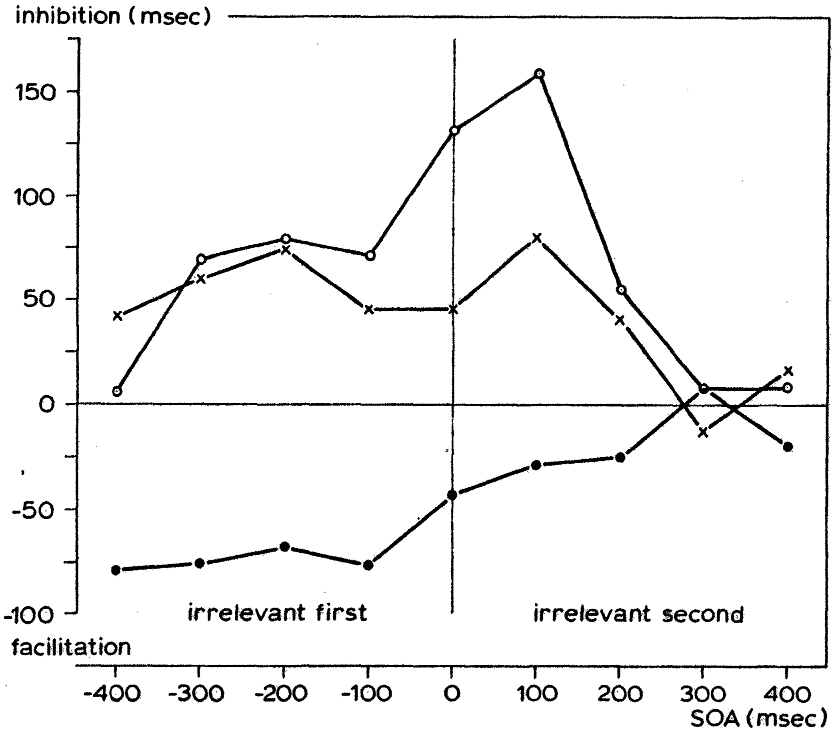

Figure 6.9
Facilitation and inhibition scores for three kinds of prime—related (o), unrelated (x), and identical (•)—and for different stimulus-onset asynchronies. (After Glaser and Düngelhoff 1984.)

absolute value of more than 400 milliseconds, the extreme leftmost value in figure 6.9. There may well be a crossover of related and unrelated curves in that region. The interpretation of the results in the figure is a matter of much dispute (see especially Vorberg 1985). For our present purposes it suffices to conclude that if the prime activates a concept other than the target, whether related or unrelated, there is interference with the naming response. The naming response is facilitated only when the identical prime is given—which shouldn't be surprising, because it activates the target concept.

In what on first view seems to be a very different set of experiments, Schriefers (1985) also found evidence for interference by a second concept, but now in the naming of a relationship. He presented the subjects with pairs of figures, such as in figure 6.10.

Consider the pair of triangles in figure 6.10a. The subject is first presented with a +, either in the left or in the right side of the field. After 1.5 seconds the pair of triangles is presented. The subject's task is to say "bigger" or "smaller", indicating whether the triangle in the position of the

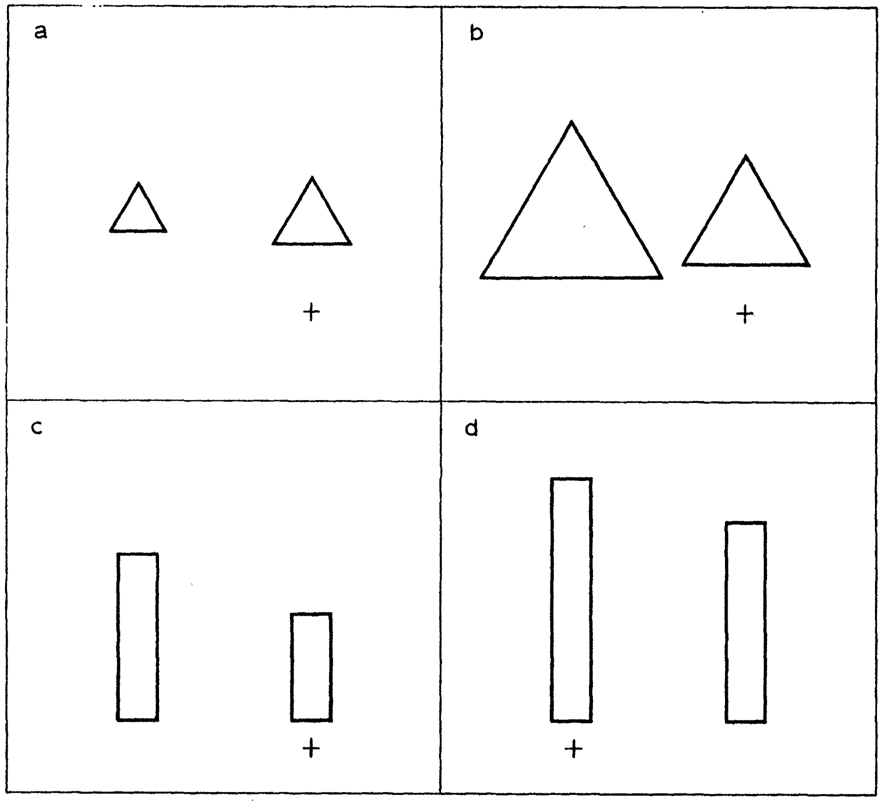

Figure 6.10
Pairs of triangles (a, b) and of sticks (c, d) used in comparative-judgment experiments by Schriefers (1985). The subject's task was to say whether the figure marked by + was *bigger* or *smaller*, *taller*, or *shorter* than the other one.

cross is the bigger or the smaller one of the pair. Schriefers complicated this relation-naming task by mixing pairs of relatively small triangles (as in 6.10a) with pairs of big triangles (as in 6.10b). The obtained naming latencies showed a strong *congruency effect*: When there was a pair of small triangles the subjects were quicker in saying "smaller" than in saying "larger", but with pairs of large triangles they were quicker to say "larger" than to say "smaller". Exactly the same pattern of results was obtained for pairs of sticks that varied in tallness, as in figures 6.10c and 6.10d. Saying "taller" was easier when both sticks were tall, as in 6.10d; saying "shorter" was the speedier response when both were short, as in 6.10c. In other words, the latency of the naming response is shortest when the to-be-judged *relative* size of a figure is congruent with the *absolute* size of the pair. Schriefers explained this result by assuming a conflict in the decision-making (or categorization) stage. When there is a pair of big figures, the concept BIG is automatically created in the subject's mind. This will interfere with the preparation of the judgment SMALLER, but it will facilitate the preparation of the judgment BIGGER.

This congruency result is very much like the previously discussed findings. When, together with the target concept, a different concept is created in the speaker's mind, there will be interference with the naming of the target. An essential question for all these findings is whether this is, as I have suggested, a matter of interference at the conceptual level, or whether the retrieval of the lexical item for one concept may suffer from activation of the other lexical item. Schriefers tested this explicitly. If it is really a matter of interference at the level of lexical access, the congruency effect should disappear when no verbal response is required. To test this, Schriefers repeated the experiment with tall and short sticks (figures 6.10c and 6.10d), but now requiring the subject to press one button when the taller of two sticks was marked and another button when the shorter stick was marked. The outcome showed an undiminished congruency effect, revealing that the interference occurs really at the judgmental, conceptual level.

6.5.3 Lexical Access

In the same series of experiments Schriefers also found a *semantic-markedness effect*. When triangle pairs of only one size category were presented, as in figure 6.10a, the response "bigger" was systematically given with shorter latency than the response "smaller". Similarly, for pairs of tall and short bars, as in figure 6.10c, the response "taller" came reliably more quickly than the response "shorter".

Adjectives such as *big* and *tall* are called *semantically unmarked*. These terms can neutrally denote the whole dimension of size or vertical extension. For example, the question "how big is your dog?" does not presuppose that the dog is big, whereas "how small is your dog?" does presuppose that the dog is small. (See Clark and Clark 1977 for a detailed discussion of the markedness notion.)

What Schriefers found, therefore, was that it takes more time to retrieve a semantically marked relation name than it takes to retrieve an unmarked one. Is this a matter of judgment, like the congruency effect? To test this, Schriefers repeated the experiment with tall and short bars, but now with push-button reactions. Would the markedness effect stay, as did the congruency effect? The result was unequivocal: There was no longer any trace of the markedness effect. To obtain the effect, it is apparently necessary to have a *verbal* response. In other words, the effect is due to finding or articulating the word. Schriefers ruled out the latter possibility by measuring pure articulation latencies independently. The semantic markedness effect is due entirely to accessing of the lexical item.

Another effect that is also generally considered to arise at the level of lexical access is the *word-frequency effect*, already mentioned in the discussion of logogen theory (subsection 6.3.2). On a picture-naming task, Oldfield and Wingfield (1965) found a high correlation between naming latency and the frequency with which the object name occurs in language use. Speech onset for *basket*, for instance, took 640 milliseconds from picture onset; *syringe* took 1,080 milliseconds on the average.

Is this difference due to visual processing and categorization, i.e., to the time needed to *recognize* the object? Wingfield (1968) showed that this was not the case by measuring recognition latencies for the same objects. To measure the recognition latency for a basket, for instance, subjects were repeatedly given the same set of pictures, but now with the task of pushing a button every time they saw a basket and giving no response otherwise. Such a recognition test was done for each object in the set. Word frequency showed no correlation with these recognition latencies. Like the semantic-markedness effect, the frequency effect arises in the phase of lexical access, not earlier. The two effects may, moreover, be connected; semantically marked adjectives are usually of lower frequency of usage than unmarked ones.

The word-frequency effect was also found in the above-mentioned study by Huttenlocher and Kubicek, who found that pictures with high-frequency names were responded to more rapidly (by some 100 milliseconds) than ones with low-frequency names. This effect was independent of the priming effect; i.e., the same frequency effect was found whether or not the priming picture was related to the target. This again testifies to the stage view of lexical access: The categorization stage is sensitive to semantic priming, the following lexical accessing stage is sensitive to word frequency, and the two effects are additive. Schriefers's congruency and markedness effects were also additive, the same two stages being involved.*

Huttenlocher and Kubicek showed, further, that the 100-millisecond word-frequency effect they found was not due to the initiation of articulation. The effect vanished almost completely when the same words that had been given as names in the original experiment were simply read. Like Schriefers, these authors excluded an articulatory explanation.

---

*A recent study by Humphreys, Riddoch, and Quinlan (1988) provides some evidence that the frequency effect is not *always* independent of semantic priming. Some interaction may arise when very long (negative) SOAs are used (−5 seconds in their experiment).

6.5.4 Accessing Lemmas and Word Forms: Two Stages?

So far we have recognized the following stages or levels of processing in picture naming: There is, first, a stage of visual processing. At this level there is a facilitating effect of visual similarity between prime and target. There is, second, a stage of conceptual categorization. At this level of processing, conceptual or functional closeness of prime and target has an effect on naming latency. Third, there is the level of lexical access proper. Here word frequency and semantic markedness have their effects. The semantic conditions for the retrieval of an item at this stage are created at the second level of processing, during conceptual categorization.

A question deserving further scrutiny is this: Are an item's lemma and form properties retrieved *simultaneously* or *successively* during the phase of lexical access? The issue is, of course, not whether lemma and word-form information are distinct kinds of information in the lexical entry, nor whether these kinds of information are relevant in subsequent phases of the formulating process (viz., during grammatical and phonological encoding, respectively). Rather, the issue is whether the lexical retrieval stage has to be further partitioned into two subsequent retrieval steps. Let us anticipate the conclusion: We do not know.

One phenomenological argument for a two-step retrieval process is the "tip-of-the-tongue phenomenon," which will be discussed in chapter 9. Everyone has had the experience that—even in fluent speech, and often to the speaker's surprise—an intended word just doesn't come. Still, one knows that something has been accessed—that the word is on the tip of one's tongue. One may even know the initial consonant or the number of syllables. It is tempting to say that in such cases the lemma has been accessed. The word's syntax has become available; the troublesome word is correctly embedded in the syntactic environment. The phonological form, however, resists retrieval.

But no critical experiments have been done to show that under tip-of-the-tongue conditions the item's syntactical properties, such as its gender, have indeed been retrieved. Does the French speaker who is in the tip-of-the-tongue state know whether the word is a *le* word or a *la* word? It might still be the case that as soon as an item's conceptual conditions are met, syntax and word form are simultaneously retrieved. In the tip-of-the-tongue state this *joint* retrieval would then be blocked.

Butterworth (forthcoming) presents an excellent review of the evidence for the two-stage theory of lexical access. There is indeed convincing experimental support (Kempen and Huijbers 1983; Levelt and Maassen 1981) for the notion that semantic activation precedes form activation in lexical access. The problem is that we do not know the status of this semantic activation. Is it "mere" stage-2 activation, i.e., activation of the concepts to be expressed? Or is it stage-3 activation, i.e., lemma activation involving a word's semantic/syntactic composition? This question is, as yet, unanswered.

6.5.5 Are Categorization and Lexical Access Nonoverlapping Stages?

A related question still under scrutiny is whether the second and third levels of processing, categorization and lexical access, are not only successive but also temporally *nonoverlapping* stages. Does lexical access begin only after the categorization has been completed, or is it rather the case that lexical activation occurs as soon as there is *any* categorization response? It could, in particular, be the case that a first, vague categorization response activates a semantically related set of lexical items—a *semantic cohort*—which is further narrowed down as the categorization becomes sharper and more definite. We considered this possibility to explain blends such as *clear* [*close/near*] in subsection 6.4.2.

Evidence that there is temporal overlap between the phases of categorization and lexical access can be found in Humphreys et al. 1988. Levelt et al. (in preparation) found that, on an object-naming task, phonological activation follows the categorization response with a very short delay. There may be a short period during which there is both semantic and phonological activation. The latter kind of activation, however, is maintained till the moment of speech onset, whereas the categorization response fades out quickly after the picture's presentation. In other words, the two phases are successive, but there may be a short period of overlap.

The fact that it is possible to demonstrate the existence of a "late" phase in object naming, in which there is activation of an item's form but no longer any activation of its meaning, is of course supportive of the idea that there can be form activation without lemma activation.

Summary

The mental lexicon plays a central role in the generation of speech. It is the repository of information about the words the speaker has available for production. This information involves, at least, the meaning of each item and its syntactic, morphological, and phonological properties. Some authors distinguish between an item's lemma information, which is essential for grammatical encoding, and its form information, which is used in the subsequent phase of phonological encoding. Entries in the mental lexicon are mutually related in various ways, according to both meaning and form.

Not all the words a speaker uses are stored in the lexicon. Words can be newly constructed on the spot. But languages differ enormously in the amount of lexical encoding they require. Speakers of English rely to a great extent on their store of frequently used words and idioms.

A lexical item's lemma information consists of the conceptual specifications for its use (including pragmatic and stylistic conditions), and of various kinds of (morpho-)syntactic properties. Among the latter are the lemma's syntactic category, the grammatical functions it imposes, and the relations between these functions and the conceptual variables or thematic roles in the conceptual structure. Also, lemmas have various diacritical variables whose parameters have to be set during grammatical encoding. Among them are such variables as person, number, tense, aspect, mood, case, and pitch accent.

There is a preference hierarchy for grammatical functions, with the subject function being the most highly valued one. Similarly, there is a "pecking order" for thematic roles, with agent as the highest-ranking one. The highest-ranking thematic role tends to be mapped onto the most preferred grammatical function. This alignment of functions and roles is acknowledged by most verbs, though not by all. Often verbs have alternative mapping orders, such as active and passive. The latter makes it possible to assign the subject function to a lower-ranking argument (e.g., for purposes of topicalization).

How are lexical items accessed during fluent speech? There are two essential criteria on which theories of access should be judged. The first one is how a theory accounts for the speed of access. Theories that are strongly sequential in character fail on this criterion. The fast access that occurs in fluent speech, as well as various speech-error phenomena, are revealing of parallel processing. The second criterion is convergence. Can a theory account for accessing of the correct item, given a speaker's intention to express a particular concept? The touchstone for correct convergence is a theory's solution of the hypernym problem: Given a concept to be expressed, how does the theory ensure that the corresponding item is accessed but none of its hypernyms? A review of various theories showed that some fail on the speed criterion (discrimination nets, maybe table linkage), and that all fail on the criterion of convergence (i.e., also logogen theory, decision-table theory, and activation spreading). A step toward the remedying of this lack of convergence was the formulation of three principles that, when implemented in a processing theory, guarantee correct convergence without creating insurmountable problems for meeting the speed criterion.

Speech errors such as word blends, substitutions, and exchanges are not random phenomena. The tongue slips in patterns, revealing aspects of the machinery involved in lexical access. Blends have their prime cause at the conceptual level: Two or more concepts compete for expression. The intruding term is often a near-synonym, but a blend can also be due to an intruding (Freudian) thought. Substitutions, on the other hand, are caused mainly (though not solely) by word association—i.e., direct associative connections between lemmas. Exchanges, finally, can arise when two different concepts in a message simultaneously trigger lexical items that are of the same syntactic category. This last kind of error causation will be taken up extensively in the next chapter. In all cases, speech errors reveal the parallel activation of more than a single lexical item.

The final section of the present chapter discussed the time course of lexical access, in particular the processing stages involved in the naming of objects or of visual relations. Three stages or levels of processing were distinguished. The first is the visual processing of the picture. This stage is sensitive to interference by visually similar material. The second stage is concerned with categorizing the visual pattern. There is a preference for using "basic-level" categories at this level of processing. This stage is differentially sensitive to interference by primes of the same and of different semantic categories. The third stage is lexical access proper, the speed of which is highly dependent on the frequency of usage of the addressed item. There is reason to suppose that this last process evolves over time in such a way that an item's lemma information decays before its form information.

Chapter 7

The Generation of Surface Structure

According to our lexical hypothesis, grammatical encoding is lexically driven. Conceptually activated lemmas instigate a variety of syntactic procedures to construct their proper syntactic environments. The procedures build the phrasal, clausal, and sentential patterns that express the required grammatical functions by means of phrasal ordering and case marking. Grammatical encoding is the process by which a message is mapped onto a surface structure.

In the present chapter we will first consider the kind of architecture needed for a Grammatical Encoder that is not only lexically driven but also incremental in its operation. The requirement of incrementality, as we saw in chapter 1, implies that a surface structure is, by and large, generated "from left to right" as successive fragments of the message become available. Wundt's principle requires that the generation of surface structure occur without much lookahead or backtracking, so that each surface unit produced can immediately be processed by the Phonological Encoder. The only explicit or computational theory of grammatical encoding that is both lexically driven and incremental in its operation is the Incremental Production Grammar (Kempen and Hoenkamp 1982, 1987; Kempen 1987; De Smedt and Kempen 1987; Hoenkamp 1983; for a review of other computational models of language generation, see Kempen 1988). The sketch of the Grammatical Encoder in section 7.1 will follow the main lines of this theory, which is a natural companion to Bresnan's Lexical Functional Grammar. The first aim of this sketch is to show that an incremental lexically driven Grammatical Encoder is a coherent and possible notion. The second aim is to demonstrate the notion's empirical potential. I will do this by mixing the presentation of the architecture with analyses of speech errors. Errors of grammatical encoding can be quite revealing of the underlying mechanisms. Are they consistent with the proposed architecture, or are they occasionally problematic?