

Chapter 1

Basic Concepts

INTRODUCTION

The study of language acquisition raises questions such as these: How do children break into language? How does knowledge of language emerge in early infancy, and how does it grow? What are the milestones of the language acquisition process? What kinds of linguistic knowledge do children display at given points of development?

The framework adopted here to answer these questions is the generative theory of Universal Grammar (Chomsky 1975, 1981, 1986). According to this theory, human beings are innately endowed with a system of richly structured linguistic knowledge, which guides infants in analyzing incoming linguistic stimuli. Such a theory makes possible clear and falsifiable predictions about children's linguistic competence and offers the tools needed to precisely characterize this competence at given points of development. As a first step in this enterprise, this chapter characterizes what it means to know a language and discusses how knowledge of language becomes available. In so doing, it introduces basic concepts underlying the approach taken in the book and presents the general framework of the research to be discussed.

1.1 REFLECTIONS ON THE COURSE OF LANGUAGE ACQUISITION

Mother: Do you want to get dressed to go see piglet?

Nina: I wanna take the play dough to piglet. (2;10)

Diana: Li faccio vedele a Luca la bambola. (2;6)

(I) to + him make see to Luca the doll

'I make Luca see the doll.'

Eve: I ride a funny clown. (1;9)

Diana: C'ha capelli lossi. (2;6)
(she) has hair red
'She has red hair.'

Adult: I don't think you write with pencil on that, Adam.

Adam: What you write with? (3;3)

Adult: You write with some crayons.

Adam: Why d(o) you carry it by de handle?

Rosa: Dov'è un'atta seggiola? (2;10)
where is a other chair
'Where is another chair?'

Rosa: Una seggiola dov'è?
a chair where is
'Where is a chair?'

Mother: È qui?
is (it) here

Human language acquisition is an astonishing process. Let us consider what these children have accomplished in about 3 years. Although their language may still not be perfect, they put words in the correct order. Nina produces quite a complex sentence, putting the complements in the right order (first the direct object and then the prepositional complement) and applying *wanna*-contraction. Scrambling complements is possible in Italian, and Diana shows that she can take advantage of this option, by putting the prepositional complement (*a Luca* 'to Luca') before the direct object (*la bambola* 'the doll'). Eve places the adjective *funny* before the noun *clown*, as required in English, while Diana places the adjective *lossi* (*rossi*) 'red' after the noun *capelli* 'hair', since she speaks Italian. One of Adam's questions features preposition stranding, although it lacks the auxiliary *do* (though Adam does include it on other occasions, as his second question shows). Forming nonadult questions (by failing to use the auxiliary) seems to be specific to English learners; note that Rosa, an Italian-speaking child, forms adultlike questions, putting the subject at the right (*un'atta seggiola* 'another chair') or left (*una seggiola* 'a chair') periphery of the sentence.

For children, acquiring a language is an effortless achievement that occurs

- without explicit teaching,
- on the basis of positive evidence (i.e., what they hear),¹
- under varying circumstances, and in a limited amount of time,
- in identical ways across different languages.

Let us look at each of these accomplishments in more detail.

1.1.1 Acquiring Language without Explicit Teaching

Unlike learning a second language in adulthood, acquiring a first or native language does not require systematic instruction. Parents usually do not teach children the rules of language or tell them what kinds of sentences they can and cannot say. Language develops spontaneously by exposure to linguistic input, that is, on the basis of what children hear. Children are rarely corrected, and even when they are, they resist the correction. For example, for a while English-speaking children say *goed* rather than *went* even though parents may occasionally correct them. For about three years my Italian-speaking child said *facete* rather than *fate* ‘make + 2PL’, although I often corrected him. McNeill (1966, 69) reports the following conversation between a child and his mother:

- (1) *Child:* Nobody don’t like me.
Mother: No, say “nobody likes me.”
Child: Nobody don’t like me.
 (eight repetitions of this dialogue)
Mother: No, now listen carefully; say “*nobody likes me.*”
Child: Oh! Nobody don’t likes me.

The child in this exchange uses double negation (*nobody don’t*), an option that is not allowed in standard English. As the exchange shows, correction does not seem to have helped the child very much: he eventually notices the use of *likes* (though he uses it incorrectly), but he fails to take advantage of the whole content of the correction.

1.1.2 Acquiring Language on the Basis of Positive Evidence

Parents’ corrections should inform children of what is *not* possible in the language they are exposed to; such information coming from correction is called **negative evidence**. As noted, however, corrections are rare, and they do not seem to improve children’s linguistic behavior. Much research has been conducted to establish whether negative evidence is available to children in the form of parents’ disapproval or failure to understand,

parents' expansion of what children say, and frequency of parents' reactions to children's utterances (see Bohannon and Stanowicz 1988; Demetras, Post, and Snow 1986; Hirsh-Pasek, Treiman, and Schneiderman 1984). Although the question is still much debated, the general conclusion is that negative evidence is not provided to all children on all occasions, is generally noisy, and is not sufficient (see Brown and Hanlon 1970; Bowerman 1988; Morgan and Travis 1989; Marcus 1993). Thus, negative evidence is not a reliable source of information. Children have the best chance to succeed in acquiring language by relying on **positive evidence**, the utterances they hear around them—a resource that is abundantly available.

1.1.3 Acquiring Language under Varying Circumstances and in a Limited Amount of Time

Children acquire language under different circumstances, and the linguistic input they are exposed to may vary greatly from child to child (see section 1.5.3 regarding acquisition of American Sign Language and creoles). Nevertheless, they all attain the same competence and do so in a limited amount of time.² By about 5 years of age they have mastered most of the constructions of their language, although their vocabulary is still growing.

1.1.4 Acquiring Language in Identical Ways across Different Languages

Children achieve linguistic milestones in parallel fashion, regardless of the specific language they are exposed to. For example, at about 6–8 months all children start to babble (see chapter 2), that is, to produce repetitive syllables like *bababa*. At about 10–12 months they speak their first words, and between 20 and 24 months they begin to put words together. It has been shown that children between 2 and 3 years speaking a wide variety of languages use infinitive verbs in main clauses (see chapter 4) or omit sentential subjects (chapter 5), although the language they are exposed to may not have this option. Across languages young children also over-regularize the past tense or other tenses of irregular verbs. Interestingly, similarities in language acquisition are observed not only across spoken languages, but also between spoken and signed languages. For example, at the age when hearing babies start to babble orally, deaf babies start to do the same manually (see Petitto 1995). It is striking that the timing and milestones of language acquisition are so similar and that the content of early languages is virtually identical, despite great variations in input and in conditions of acquisition.

1.2 THE LOGICAL PROBLEM OF LANGUAGE ACQUISITION

Looking at the facts described in the last section, researchers have characterized the problem of language acquisition as follows (see Baker and McCarthy 1981):

- Children come to have very rich linguistic knowledge that encompasses a potentially infinite number of sentences, although they hear a finite number of sentences.
- The data that children draw upon consist of positive evidence (sentences that are acceptable in the language they are exposed to).
- Children are not told which sentences are ill formed or which interpretations sentences cannot have in their language, but eventually they attain this knowledge; all mature speakers can judge whether a sentence is acceptable or not (under a given interpretation).
- Although children make “errors,” they do not make certain errors that would be expected if they generalized from the linguistic input. For example, although children hear sentences like *Who do you wanna invite?* and *Who do you wanna see?*, they do not generalize from these to impossible English sentences like **Who do you wanna come?* (see section 1.4); although this generalization would seem reasonable, children never say such sentences.

These points are part of an argument about the mechanisms underlying language acquisition—the so-called **argument from the poverty of the stimulus**. Essentially, this argument starts with the premises that all speakers of a language know a given fairly abstract property and that this property cannot be induced from the evidence available to children (positive evidence).³ What conclusion can we draw from these premises? That is, where does linguistic knowledge come from? After a brief excursion into background assumptions, this is the question we will explore.

1.3 THE NOTION OF GRAMMAR

To know a language means to possess a system of knowledge called grammar. A **grammar** is a finite system since it is somehow represented in the mind/brain. As Chomsky showed in the 1950s, it is a **mental generative procedure** that uses finite means to generate an indefinite number of sentences. The term *grammar*, as used here, refers to a psychological entity, not to an inventory of sounds, morphemes, inflectional paradigms,

and syntactic constructions (e.g., passives, relative clauses). Although we will be using terms such as *passive*, *relative clause*, and *interrogative* and will be discussing the acquisition of the corresponding constructions, it should be clear that we are using such terms only for convenience. They do not have an independent status in the framework adopted here. For example, interrogatives are the result of movement operations that displace constituents in certain ways. These movement operations are not specific to interrogatives, but are shared by other constructions.

Our linguistic knowledge allows us to produce and understand sentences we have never heard before. It also gives us the tools to establish whether a sentence is acceptable in our language or not. For example, although (2) is comprehensible, it is not an acceptable sentence in English. It does not comply with what we know to be licit in English.

(2) Dog a old a bone ate.

It is again our grammar that permits us to say that the sentence in (3) is perfectly sound, but only on the interpretation that Mary washed another female individual. It cannot mean that Mary washed herself.

(3) Mary washed her.

In other words, the pronoun *her* in (3) must refer to or pick out an individual distinct from the individual picked out by *Mary*. As (4) shows, however, pronouns need not always be interpreted in this way.

(4) Mary washes her socks.

The sentence in (4) is ambiguous: it can mean either that Mary washes some other female individual's socks or that Mary washes her own socks. Unlike the pronoun in (3), the pronoun here can be interpreted in two ways: either it refers to the same individual picked out by *Mary* or it refers to another salient individual in the extralinguistic context.

Linguistic ambiguity is pervasive. Sentence (5) is also ambiguous, having the two readings in (6a) and (6b) (example from Lightfoot 1982, 19).

(5) John kept the car in the garage.

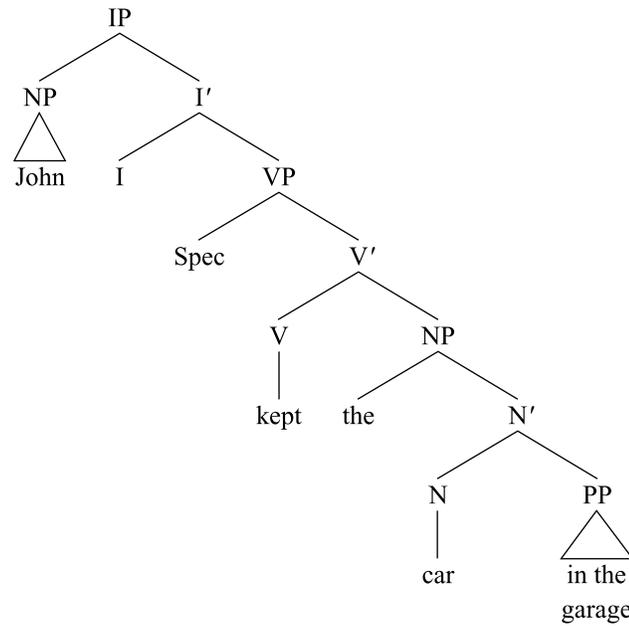
(6) a. The car that John kept was the one in the garage.

b. The garage was where John kept the car.

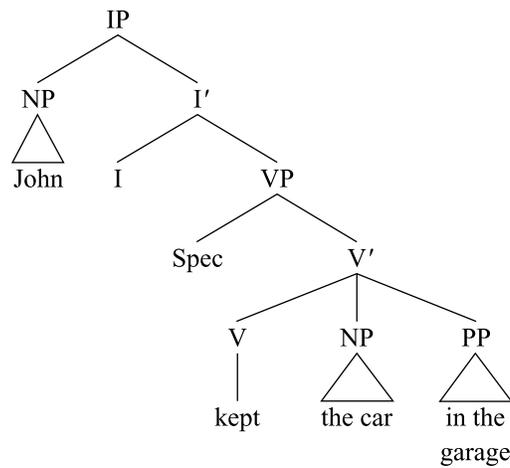
Human beings have the resources to cope with linguistic ambiguity. We know whether a sentence is ambiguous or not, whether we can interpret it in certain ways or not, because our grammar assigns sentences structural

representations constrained in specific ways. The string in (5) can be associated with two structural representations, (7a,b), each corresponding to one of the two legitimate interpretations of this string, (6a,b).

(7) a.



b.



On the interpretation in (6a), whose structural representation is given in (7a), the sequence *the car in the garage* forms a constituent. On this inter-

pretation (5) means that, among the things available, John chose the car that was in the garage. On the interpretation in (6b), whose structural representation is given in (7b), the same sequence is split into two constituents, *the car* and *in the garage*. On this interpretation it means that the garage is the place where John kept the car. The different interpretations that we assign to (5) are based on the different structural representations that our mental grammar associates with it.

In summary, we can do certain things with language because we have a grammar, a psychological entity realized somehow in our mind/brain. This grammar assigns certain structural representations to sentences, and it sanctions certain interpretations while banning others. It does this by means of constraints that establish what is possible and what is not possible in language. In the next section we will look more closely at the notion of constraints.

1.4 CONSTRAINTS

Constraints are linguistic principles that prohibit certain arrangements of words, certain operations, and certain associations of sounds and meanings. They encode properties that hold universally (i.e., in language after language) and are all inviolable (i.e., no violation of any sort is tolerated). In the framework adopted here, constraints are not ranked with respect to one another. (In this sense, this conception of constraints differs from the one advocated in Optimality Theory, where constraints are violable and ranked, and where different constraint rankings are held to underlie differences between languages; see Archangeli and Langendoen 1997; Barbosa et al. 1998.) Sentences must conform to linguistic constraints if they are to be considered well formed or acceptable. For example, the question in (8b), obtained from the declarative sentence in (8a), is judged ill formed by English speakers, as conventionally indicated by the “star” (*), because it violates a constraint of English grammar.

- (8) a. John regrets that Paul behaved badly.
 b. *How does John regret that John behaved?

In this book the term *constraints* will be used as defined above, although in the literature about language it is also used otherwise.

Constraints are of two kinds: **constraints on form** and **constraints on meaning**. Constraints on form encode the linguistic information that certain sentences are ill formed. An example of a constraint on form is the

one operative in (8b). Notice that a minimal variant of (8b)—namely, (9b), obtained from (9a)—is well formed.

- (9) a. John thinks that Paul behaved badly.
 b. How does John think that Paul behaved?

The question of interest here is this: how does the child who has heard (8a) and (9a,b) refrain from abstracting a rule that would yield (8b)? In fact, English speakers all share the knowledge that questions like (8b) are ill formed. Linguists propose that a constraint on grammar is responsible for this knowledge. Moreover, speakers of all other languages investigated thus far also know that the counterpart of (8b) is ill formed in their languages, and that the counterpart of (9b) is well formed. Thus, the kind of knowledge that allows us to say that (8b) or its counterpart in another language is not acceptable cannot be language specific, but must be universal.

Another constraint on form is that governing the optional contraction between *want* and *to* in English.

- (10) a. Who do you wanna invite?
 b. Who do you want to invite?
- (11) a. When do you wanna go out?
 b. When do you want to go out?
- (12) a. *Who do you wanna come?
 b. Who do you want to come?

It is possible to contract *want* and *to* in (10a) and (11a). However, in (12a) the result of *wanna*-contraction is ill formed, something that speakers of English implicitly know, although they may not be able to formally express this prohibition. Essentially, *wanna*-contraction is not possible when the questioned element is the subject of the infinitival clause, for example, *John* in *I want John to come*. Although linguists speak of “a constraint governing *wanna*-contraction,” remember that terms naming specific syntactic constructions are used only for convenience. In its general formulation this is a universal constraint that blocks a certain process from occurring in certain structural configurations.

In summary, a constraint on form underlies our ability to say that certain sentences are ill formed. It is a piece of linguistic knowledge that restrains us from making wrong generalizations—for example, from inducing (12a) from (10) and (11).

Beyond constraints on form, grammars include constraints on the meaning that speakers assign to acceptable sentences. Consider the sentences in (13). Both sentences are perfectly well formed, but while in (13a) the two italicized expressions *he* and *John* cannot pick out the same individual, in (13b) they can (here, “*” indicates that the sentence in (13a) is ruled out when the two italicized expressions pick out the same individual).

- (13) a. **He* danced, while *John* was singing.
 b. While *he* was singing, *John* was dancing.

In other words, (13b) is ambiguous: the pronoun *he* can refer either to the same individual that *John* refers to (anaphoric interpretation of the pronoun) or to another salient character in the extrasentential context (exophoric/deictic interpretation of the pronoun). By contrast, in (13a) the pronoun *he* can only be interpreted as referring to some individual other than John. Interestingly, in all languages investigated so far, the counterparts of (13a,b) work the same way; that is, the constraint governing the interpretation of (13) holds universally. Constraints on meaning can be represented as pairs prohibiting the association of certain sentences (S) with certain meanings (M).

- (14) *⟨S1, M1⟩

The association between sentence and meaning in (13) is governed by a constraint prohibiting an anaphoric interpretation of the pronoun in certain structural configurations. This constraint, to be discussed in chapter 8, is Principle C of the binding theory. Principle C, or whatever subsumes its effects, bans an anaphoric interpretation of the pronoun in (13a), but not in (13b).

In summary, linguistic knowledge about the possible form of sentences and about the possible association of form and meaning is couched in terms of constraints that hold universally (as linguistic research since the 1960s has shown) and that are not violable. Our linguistic behavior is guided by these constraints, witness the fact that we can judge whether or not it is licit to contract *want* and *to* in certain sentences, and that we know how to interpret pronouns depending on the linguistic context.

1.5 WHERE DOES KNOWLEDGE OF LANGUAGE COME FROM?

How do we know that a sentence is ill formed, that it cannot have a given meaning, or that it is ambiguous? Four hypotheses have been advanced,

involving imitation (section 1.5.1), reinforcement (section 1.5.2), association procedures (section 1.5.3), and Universal Grammar (section 1.5.4).

1.5.1 Language Learning through Imitation

One hypothesis holds that children learn language by imitating what adults say, by trying to repeat what they hear. However, several facts, showing that there is no necessary similarity between linguistic input and linguistic output, militate against this hypothesis.

First, studies of parents' speech suggest that children are usually not influenced by caregivers' speech style. Newport, Gleitman, and Gleitman (1977) have shown that a high proportion of parents' utterances are questions (*What do you want?*) and commands (*Get the toy car!*) and only 25% are simple declaratives. By contrast, simple declaratives are the first kind of sentence that children mostly produce.

Second, children continually produce novel utterances, in two senses. For one thing, they hear a finite number of sentences, but they come to be able to produce and understand indefinitely many sentences, including vast numbers they have never heard and therefore cannot be imitating. For another thing (and this is the most compelling evidence against the acquisition-through-imitation hypothesis), children produce utterances that they *cannot* have heard before, because the adult speakers in their environment do not produce them.

It is well known that English learners overregularize irregular past tense verbs and say for example *goed* instead of *went* and *singed* instead of *sang*, although they have never heard these forms, because adults do not use them. In the same vein Guasti, Thornton, and Wexler (1995) have found that English-speaking children aged 4–5 years produce negative questions with the form in (15). No adult utters such sentences; thus, children cannot have learned them by imitation. (Although (15a,b) are not acceptable in the adult language, they are part of children's grammar and therefore are not marked with “*”. This practice is followed throughout the book.)

- (15) a. What does he doesn't eat?
 b. Why could he couldn't wash his hands?

Similarly, Thornton (1990) has shown that English learners produce long-distance extraction questions in which an interrogative pronoun occurs twice, in both sentence-initial and intermediate positions.

- (16) What do you think what the puppet has eaten?

All these examples demonstrate that children go beyond their linguistic input and try to say things they cannot have heard. Children do this because they are attempting to discover the “rules” operating in their language, rules that may vary from one language to another. These facts point toward the conclusion that imitation does not play a crucial role in language acquisition.

1.5.2 Language Learning through Reinforcement

Behaviorist psychologists have claimed that language is learned through the mechanism of reinforcing the contingent association between stimulus and response—the same general-purpose mechanism that is invoked to explain other learning processes in animals and in humans (see Skinner 1957). According to this view, children learn language because they are positively reinforced when they produce correct verbal expressions, negatively reinforced when they make errors.

Although the learning-through-reinforcement hypothesis is simple, it cannot explain how humans acquire language and cannot characterize human linguistic competence, as Chomsky (1959) details in his review of Skinner 1957. First, like the acquisition-through-imitation hypothesis, it cannot explain the fact that children acquire competence over an indefinite number of sentences: they understand and produce sentences they have never heard and produced before, that is, for which no reinforcement was provided. Second, parents generally pay attention to what children say and not how they say it. If a child asks a question, the adult will hardly check for its grammatical correctness, but will simply answer, as the following exchanges illustrate:

(17) *Adam*: Where penny go? (Adam, 2;5)

Mother: I don't know.

(18) *Adam*: Where penny go? (Adam, 2;5)

Mother: Didn't you drop your pennies on the floor?

These exchanges show that the notion of reinforcement is vague. In a sense, by responding, the mother is reinforcing the child: she has understood the question. But if the child were to take this reinforcement as a sign that his question was grammatically correct, he would never converge on the correct grammar. Eventually, children attain adult competence and form adultlike questions, but this does not seem to happen through reinforcement. In fact, as we saw in section 1.1.1, even when

parents correct children's ungrammatical sentences, these corrections go unnoticed. In sum, positive and negative reinforcement do not explain human linguistic attainment.

1.5.3 Language Learning through Association

Another hypothesis about how language acquisition occurs is expressed by an approach called **connectionism**, neural networks, or parallel distributed processing. (See, e.g., Elman 1993; Elman et al. 1996; Rohde and Plaut 1999. See Pinker and Prince 1988 and Marcus 1998 for criticism; and see Pinker 1999 for an introductory discussion.) At the outset it is worth noting, as does Marcus (2001), that the term *connectionism* is ambiguous. Generally it is associated with the idea that brain circuits do not support the representation of symbols and rules; connectionist models are thus usually opposed to models in which symbols are manipulated. However, in addition to symbol- and rule-free models, there exist connectionist models whose goal is to explain how symbolic manipulations can be implemented in a neural substrate (see, e.g., Shastri and Ajjanagadde 1993; see Marcus 2001 for an extensive discussion of these issues). The remarks that follow apply to models that aim at eliminating symbols and rules.

Connectionist models or artificial neural networks are inspired by a coarse metaphor of the brain, in that they consist of several interconnected neuronlike processing units modified by learning associations between input (stimulus) and output (response) patterns. Interactions among these units give rise to behavior that simulates, sometimes very accurately and precisely, actual human behavior. A network consists at least of input and output units connected by modifiable weighted links. During the learning phase the network is presented with examples of both input and output. Given an input, the network modifies the weights of its connections so as to produce the correct output. After learning, the network can generalize to new stimuli provided they belong to the same class of stimuli used in the training phase. Notice that in these models neither nodes nor links correspond to linguistic categories or rules. These are represented in the network by various patterns of activation among links.

Here we will briefly look at some linguistic phenomena connectionists have sought to account for, noting simply that many intricate aspects of language acquisition and of human linguistic competence still await explanation within a connectionist approach (for more detail on debates surrounding connectionism, see the works just mentioned).

Some connectionist models assume that the mental mechanism employed for acquiring language operates on the basis of analogy or similarity. To gain some insight into how these models operate, let us examine an aspect of acquisition they are most frequently used to simulate: the acquisition of the English past tense. We will then consider whether such models can indeed explain how children learn language.

Regular verbs in English form the past tense by adding the morpheme spelled *-ed* to the stem, regardless of the phonetic features of the stem, while irregular verbs are grouped in family resemblance patterns that form the past in various idiosyncratic ways (e.g., *drink/drank*, *sing/sang*, in which a vowel is changed).

Connectionists claim that acquisition of the past tense of regular and irregular verbs consists in learning associations between the phonetic properties of verb stems and the phonetic properties of their past forms and in generalizing these associations to similar-sounding words (Rumelhart and McClelland 1986). In the connectionist view, children learn that verbs ending in *alk* [ɔ:k] (e.g., *talk* and *walk*) are associated with the past tense form *alked* [ɔ:kt]; similarly, verbs having the pattern *consonant-consonant-i-nk* (e.g., *drink*) are associated with a past tense form having the pattern *consonant-consonant-a-nk* (e.g., *drank*). Children are said to exploit these associations to form the past tense of verbs: whenever they hear a new verb with a specific phonetic pattern (input), they will produce the past tense form associated with that pattern. (See Rumelhart and McClelland 1986; Plunkett and Marchman 1993. For criticism, see Pinker and Prince 1988; Kim et al. 1994.)

Connectionist models mimic some aspects of the process of morphological acquisition; for example, they make the overregularization errors that children make in learning the past tense. However, on closer inspection the actual process of language acquisition and these connectionist simulations are not greatly similar, as far as regular verbs are concerned. Marcus (1995) points out that some connectionist models overregularize vowel-change verbs (*sing* becomes *singed*, rather than *sang*) less frequently than no-vowel-change verbs, while children overregularize the former more frequently than the latter; in addition, these models cease to overregularize verbs only after an abrupt change in the training input, while children do the same although there is no change in the input.

Connectionist models regularize irregular verbs on the basis of resemblance to similar-sounding regular verbs—for example, producing *holded* by analogy with *molded*, *folded*. However, Pinker (1994b) shows that

more than similarity of sounds is at work in overregularization: children treat the same phonetic string differently depending on its grammatical status. Stromswold (1990) found that children overregularize main verb *have*, *do*, and *be* just as they overregularize other main verbs, but they do not overregularize the same verbs when they are used as auxiliaries. For example, children say *I doed it* rather than *I did it* or *I haved it* rather than *I had it*, but they do not say *Doed you come?* rather than *Did you come?* or *I haved eaten* rather than *I had eaten*. Evidence also shows that in some groups of people with language impairments the production of regular and irregular verbs is differently affected (see chapter 11; also see Pinker 1999 for a discussion of these issues). These facts cast doubt on the view that a single learning mechanism based on association is responsible for the acquisition of the past tense and of language more generally.

Many connectionist models have attempted to simulate aspects of morphological acquisition; some have also attempted to simulate limited aspects of syntax—for example, sequencing of Noun-Verb, or Noun-Verb-Noun, where the noun can be modified by a relative clause (see Rohde and Plaut 1999; see also Elman 1993).

Although connectionist models can learn some sequencing of words, to date it is unknown whether they can learn the knowledge expressed by linguistic constraints of the kind mentioned in section 1.4 and thus refrain from generating the incorrect sentences discussed there, which children do not produce (see also Marcus 1999). It is also unclear how such models can come to know whether a sentence is ambiguous—indeed, to know all the intricate and abstract aspects of linguistic knowledge discussed in Chomsky's review of Skinner 1957. This criticism is not intended to deny that association and some form of stochastic information are involved in language acquisition. Indeed, some stochastic information may help infants in segmenting speech into word units (see chapter 3). Moreover, an associative mechanism may be appropriate to handle certain linguistic phenomena. Pinker (1997) discusses a theory, called the *word-and-rule theory*, that includes both rule-based and associative components. The rule-based component manipulates symbols and is responsible for the inflection of regular words. It operates on members of syntactic categories (e.g., Verb, Noun) and generates inflected words. For example, it generates the past tense of a regular verb by adding the morpheme spelled *-ed* to the stem; it forms the plural of a regular noun by adding the morpheme spelled *-s* to the stem. The associative mechanism is responsible for the inflection of irregular words and operates on the basis of (sound) similar-

ity. For example, a novel verb such as *spling* is held to have the past tense form *splang* by virtue of similarity to pairs of forms already stored in memory such as *sing-sang* and *ring-rang* (see Prasada and Pinker 1993). Sound similarity plays no role in the inflection of regular words. Thus, according to the word-and-rule theory, two mechanisms—one rule-based, the other associative—are involved in the acquisition of inflection, each subserving different aspects of acquisition (see sections 11.2.1.4, 11.3.2). Association may account for the acquisition of some aspects of linguistic knowledge, but it can hardly answer the entire question of how children acquire language.

Another area where we can compare connectionist models and human learners is the ability to acquire language from radically degenerate input. Human beings clearly demonstrate this ability, as proven by data from creole languages and sign languages (see Bickerton 1988; Goldin-Meadow and Mylander 1984; Kegl 1994). In the nineteenth century people on plantations and in slave colonies often developed a rudimentary form of language to communicate—a lingua franca or pidgin. Once a pidgin has native speakers—the children of the individuals who originate it—it develops into a full language, called a creole. Unlike pidgins, creole languages have function morphemes and a more elaborated structure. Creoles are thus expanded and refined by children on the basis of rudimentary, degenerate input, the pidgin. A similar situation occurs with sign languages. Deaf children born to late learners of American Sign Language (ASL) receive very rudimentary linguistic input, because their parents avoid complex structures and often omit function morphemes. In spite of this degenerate input, these children achieve a more refined competence than their parents, acquiring a sign language that includes complex structures and function morphemes (see Newport 1988).

As Bickerton (1996) points out, connectionist models cannot simulate this ability of human learners. Since a connectionist system learns exclusively on the basis of its input, it will learn a degenerate language if the input is degenerate. If it is to expand and refine the input, it must be endowed with a program that does just that; but such an adjustment does nothing else than supply the model with an innate component, which amounts to recognizing that language acquisition requires innate (possibly language-specific) structures.

In summary, the connectionist models discussed in this section are based on the assumption that language can be acquired through association. However, much linguistic knowledge seems to resist an explanation

in such a paradigm, calling instead for a theory that incorporates innate structure, rule-based mechanisms, and constraints.

1.5.4 The Innateness Hypothesis

Recall the premises of the argument from the poverty of the stimulus: that all speakers of a language know a fairly abstract property and that this property cannot be induced from the evidence available to children (positive evidence). The conclusion that these premises invited us to draw is the answer to the question we started with: where does linguistic knowledge come from? Imitation, reinforcement, and association having failed to answer this question, we must look further. In fact, the answer that Chomsky (1959) gave in arguing against behaviorist views and that conclude the argument from the poverty of the stimulus is that *this knowledge is inborn*.

There is a debate as to how rich the genetic makeup supporting human linguistic abilities is. Researchers in the Chomskyan tradition assume that inborn human knowledge is richly structured and must consist of the kinds of constraints (or of something equivalent in its effects) discussed above. It is very unlikely that these constraints are learned since they hold universally. It would be very curious that all languages conform to these constraints if this crosslinguistic similarity were not somehow dictated by our mind/brain: languages “are all basically set up in the way that human biology expects them to be” (Gleitman and Lieberman 1995, xxi). Thus, children are born expecting that, whichever language they are going to hear, it will have the properties that their genetic equipment is prepared to cope with.

The hypothesis that the language capacity is innate and richly structured explains why language acquisition is possible, despite all limitations and variations in the learning conditions. It also explains the similarities in the time course and content of language acquisition. How could the process of language acquisition proceed in virtually the same ways across modalities and across languages, if it were not under the control of an innate capacity? Of course, not *all* linguistic knowledge is innate, for children reared in different linguistic environments learn different languages. That languages vary is obvious. For example, in Italian the sentential subject can be phonologically silent, while in English it cannot. However, this variation is not unlimited. **Universal Grammar (UG)** is the name given to the set of constraints with which all human beings are

endowed at birth and that are responsible for the course of language acquisition. UG defines the range of possible variation, and in so doing it characterizes the notion of possible human language. A characterization of UG is a characterization of the initial linguistic state of human beings, the genetic equipment necessary for acquiring a language.

According to this nativist view, acquisition results from the interaction between inborn factors and the environment. Language is not learned, but, under normal conditions, it is deemed to emerge at the appropriate time, provided the child is exposed to spoken or signed language. Obviously, children have to learn the words of their language, its lexicon. They also have to figure out what the regularities of their language are, and how innate constructs are instantiated in their linguistic environment (Fodor 1966).

1.5.5 The Principles-and-Parameters Model

Our genetic endowment makes it possible to learn any human language. Children raised in an English-speaking environment speak English, those raised in an Italian-speaking environment speak Italian, and those raised in a Tibetan-speaking environment speak Tibetan. Although all languages have the same basic underlying structure, there are variations. For example, in some languages (e.g., English and Italian) the verb comes before complements; in others (e.g., Turkish and Bengali) it comes after. So, while an English speaker would say *John bought books*, a Turkish speaker would say something equivalent to *John books bought*. Some languages (e.g., Italian and Spanish) allow the sentential subject to remain phonologically unexpressed; others (e.g., English) do not. So, while *Bought a book* is ungrammatical in English, its counterpart is acceptable in Italian or Spanish.

The model of language adopted here makes sense of these variations by holding that UG consists of two types of constraints: **principles** and **parameters**. Hence, it is called the *principles-and-parameters model* (Chomsky 1981). Principles encode the invariant properties of languages, that is, the universal properties that make languages similar. For example, the constraint discussed in section 1.4 governing the interpretation of pronouns is a principle; in any human language this principle regulates the interpretation of pronouns. Parameters encode the properties that vary from one language to another; they can be thought of as switches that must be turned on or off. An example is the pro-drop or null subject parameter governing the phonological expression of the sentential subject.

As a first approximation, we can formulate the pro-drop parameter as in (19).

(19) Can the sentential subject be phonologically null?

Depending on the particular language, the answer to the question in (19) will vary. If a child is exposed to Italian, the parameter in (19) will be set to the positive value; if the child is exposed to English, it will be set to the negative value.

Under the principles-and-parameters model, children are innately endowed with principles and parameters, because both are given by UG. The children's task is to set the parameters to the value expressed by the language of their environment. In this model, then, language acquisition consists (among other things) in selecting the appropriate values of the parameters specified by UG.

The theory of language acquisition endorsed here is a selective theory, rather than an instructive one. "Under an instructive theory, an outside signal imparts its character to the system that receives it, instructing what is essentially a plastic and modifiable nervous system; under a selective theory, a stimulus may change a system that is already structured by identifying and amplifying some component of already available circuitry" (Lightfoot 1991, 2). In other words, under an instructive theory genuine learning takes place; under a selective theory no learning takes place because the stimulus works on what is already inborn.

Selection, rather than instruction, operates in other biological systems besides language. Niels Kaj Jerne has defended a selective theory of antibody formation, whereby antigens select antibodies that already exist in an individual's immune system (for discussion of these issues, see Jerne 1967, 1985; Piattelli-Palmarini 1986). He has also conjectured that certain central nervous system processes might work selectively and has pointed out that in the history of biology selective theories have often replaced instructive theories.

In summary, UG is the human genetic endowment that is responsible for the course of language acquisition. It includes principles and parameters that encode the invariant and variant properties of languages, respectively. Parameters define the range of variation that is possible in language; and together, principles and parameters define the notion "possible human language." Language acquisition is a selective process whereby the child sets the values of parameters on the basis of the linguistic environment.

1.6 THE CRITICAL PERIOD

Innate behaviors are often distinguished by the existence of **critical periods** during which the ability to acquire the competence reaches its peak; thereafter, the ability to acquire that competence declines. For example, visual abilities in animals develop naturally only if animals receive appropriate visual stimulation early in life. Cats seeing only vertical stripes become blind to horizontal stripes and vice versa (Blackmore and Cooper 1970). The same critical period effects are observed in the development of imprinting in ducks and of attachment in young of various species (Hess 1972) and in song learning in birds (Marler 1970). For example, ducklings become emotionally attached to the first moving thing they see, be it the mother duck, a human being, or an object.

Since language is also innate, we may wonder whether its acquisition is subject to critical period effects. Lenneberg (1967) suggested that language can develop fully only if it is acquired before puberty. Since then, evidence has accumulated that a native competence is acquired only if language is acquired before puberty.

One such piece of evidence is provided by children deprived of linguistic and social interaction during their childhood. The girl known as Genie was reared in such conditions until she was discovered at the age of 13 (Curtiss 1977). Even after several years of linguistic rehabilitation, Genie's language abilities were very limited, especially in syntax.

Evidence in favor of the critical period hypothesis also comes from congenitally deaf people who are exposed to a first language, American Sign Language (ASL), at different ages. Singleton and Newport (1994) tested production and comprehension of ASL verb morphology by congenitally deaf individuals exposed to sign language from birth, from 4 to 6 years of age, or after age 12. They found that performance linearly declined with age of first exposure. Individuals exposed to ASL from birth performed better than those exposed from 4 to 6 years of age, and the latter in turn performed better than individuals exposed after age 12.

That early exposure is critical in attaining native competence is also supported by studies showing that a foreign accent can already be detected in individuals first exposed to a foreign language at age 3 and that accents get stronger as age of first exposure increases (Flege, Yeni-Komshian, and Liu 1999). Similarly, Johnson and Newport (1989) have shown that only speakers who have been exposed to American English, as a second language, before age 7 achieve native performance on an examination testing

mastery of morphology and syntax. Speakers who are exposed after age 7 do not acquire native competence.

In summary, the evidence to date shows that there are critical period effects for acquiring the phonology, morphology, and syntax of one's native language. While all human beings are endowed with a richly structured system of linguistic knowledge, this system can develop naturally and fully only if the individual is exposed to appropriate stimuli early in life. This explains the fact that acquiring a language (native or foreign) is a natural achievement for children and becomes more difficult as one gets older.

Key Words

Argument from the poverty of the stimulus
Association procedure
Connectionism
Constraints on form
Constraints on meaning
Critical period
Grammar as a mental generative procedure
Imitation
Negative evidence
Parameters
Positive evidence
Principles
Universal Grammar