Jay Friedenberg
Gordon Silverman

3 EDITION

Cognitive Science

An Introduction to the Study of Mind

9

The Linguistic Approach

Language and Cognitive Science

How very commonly we hear it remarked that such and such thoughts are beyond the compass of words! I do not believe that any thought, properly so called, is out of the reach of language.

—Edgar Allan Poe, 1846

THE LINGUISTIC APPROACH: THE IMPORTANCE OF LANGUAGE

Linguistics is the study of language. There are many different kinds of linguistics studies—each with its own theoretical perspectives and methodologies. Some of these adopt a neuroscience approach and use the case study method; the researchers study the language-related deficits of patients who have suffered brain damage. Others implement various network models of how language information is represented and processed. Some linguists take on a developmental orientation: They examine how language ability grows and changes with time during the development of the individual. Still others who study linguistics are philosophers who ask questions about the nature of language and the relationship between language and thought. In fact, language can be studied from the vantage point of every field of study that has been described in this book and more. What makes linguistics unique, then, is not the perspective or the tools it brings to the table but the subject matter of the investigation—language itself.

Linguistic studies attempt to answer many questions. Language is so complex that much of the research that has been conducted in this area has been directed toward an understanding of the structure of language (in addition to how it is used). These studies have focused on grammatical rules that specify allowable combinations of linguistic elements. Another interesting issue is whether humans are unique in using language or whether some animals possess language ability. Languages, of course, cannot be learned overnight, and many linguists have studied language acquisition—how it is acquired during development.

Linguistics is truly interdisciplinary, and in this chapter, we will survey the approaches to the study of language that are represented by the disciplines of philosophy, cognitive psychology, neuroscience, and artificial intelligence.

THE NATURE OF LANGUAGE

With this in mind, we begin our discussion of the linguistic approach with an exploration of the nature of language. There has been much debate about what language is, exactly—and there is no agreed-on definition. It is easier to list its most important characteristics (in lieu of providing a definition). According to Clark and Clark (1977), language has five characteristics:

1. *Communicative:* Language allows for communication between individuals. **Communication** refers to the production, transmission, and comprehension of information.
2. *Arbitrary:* A language consists of a set of symbolic elements. Symbols, as we noted in Chapter 1, are referential—they stand for or refer to something. Linguistic symbols can be almost anything. Most commonly, these symbols are for sounds, pictures, or words. The defining hallmark of these symbols is that they are completely **arbitrary**. Virtually any sound, picture, or word could be chosen to represent a particular thing. The sound of the word that represents "house" in English is different from the sound of the word that represents the same item in Spanish.
3. *Structured:* The ordering of the symbols in a language is not arbitrary but is governed by a set of rules, or is **structured**. The rules specify how the symbols may be combined. In English, we place the adjective before the noun, as in the phrase "the big house." In Spanish, this same proposition is expressed via a different set of rules: The adjective follows the noun, as in "la casa grande."
4. *Generative:* The symbolic elements of a language can be combined to create a very large number of meanings. Just think of how many six-word sentences one would be able to generate in English. If we start with the sample sentence "The fox jumped over the fence," we can then substitute the words *dog*, *cat*, *deer*, and words for many other animals in place of *fox*. Likewise, we can substitute the words *bottle*, *can*, or *tire* for the word *fence*. So the number of variations on just this one sentence is large. Every day, we utter new sentences that we never have uttered before. The **generative** property of language makes language very powerful, as virtually any idea that can spring to mind can be expressed.
5. *Dynamic:* Languages are **dynamic**, constantly changing as new words are added and grammatical rules altered. Only 30 years ago, there was no word for the concept that *e-mail* represents—because it didn't exist.

There is a fundamental distinction to be made that has to do with the type of linguistic representation: whether it is auditory, visual, or having to do with another sensory domain. Spoken language is naturally produced via the faculty of speech and is understood via listening. Speech and listening to speech can, of course, be transformed into their equivalents within the visual domain—writing and reading. Beyond this, there are also languages that consist of motoric gestures, such as the American Sign Language (ASL), and tactile languages, such as Braille. If we consider spoken language, the most common form of language usage, we must then describe two important elements of spoken language: the phoneme and the morpheme.

A **phoneme** is the smallest unit of sound in the sound system of a language. A phoneme has no meaning. Phonemes correspond in a rough way to the letters of an alphabet; in some instances, multiple phonemes correspond to a single letter. The phoneme for the letter *a* as it is pronounced in *father* corresponds to the "ah" sound, whereas the phoneme for *a* as it is pronounced in the word *cane* corresponds to the sound "ay." There are about 45 phonemes in the English language. Some instances are shown in Table 9.1. The smallest number of phonemes reported for a language is 15. The largest is 85. Phonemes, like letters, are combined to form the spoken versions of words.

Morphemes are the smallest units of spoken language that have meaning. They roughly correspond to words but can also be the parts of words. Thus, the sound of the spoken word *apple* is a morpheme, but so is the sound of *s* denoting the plural form. If we want to change the form of *apple* from singular to plural, we add the letter *s* to form *apples*, which changes the meaning. Similarly, there is the morpheme that corresponds to the sound "ed," which, when added to the root form of many verbs, forms the past tense. Considering that there are about 600,000 words in the English language, the number of morphemes that the language has is quite large.

In addition to the elements of language, there are the rules that allow for their possible combinations. There are multiple sets of rules. **Phonology** refers to the rules that govern

Table 9.1 Selected English consonant and vowel phonemes.

Consonants	Vowels
p (pill)	i (beet)
w (wet)	e (baby)
s (sip)	u (boot)
r (rate)	o (boat)
g (gill)	a (pot)
h (hat)	^ (but)

the sound system of a language; **morphology**, to those that govern word structure; syntax, to those that govern the arrangements of words in sentences; and **semantics**, to those that have to do with word meanings. Collectively, these rules are known as the **grammar** of the language. It is important to distinguish between word meanings as it is used most commonly and its more esoteric meaning (used in linguistics studies). Prescriptive grammar is the formal and proper set of rules for the use of language, in which we all received training at school. Descriptive grammar refers to the underlying rules, which linguistics researchers infer from the way people actually use language. We will come back to the subject of grammar and how it is used to describe the hierarchical structure of sentences in our discussion of Noam Chomsky's theory of language.

Interdisciplinary Crossroads: Language, Philosophy, and the Linguistic Relativity Hypothesis

A principal focus of this book is the nature of thought. If thought is representational, as it most surely seems to be, then what is its form? In Chapter 1, we described the several forms that thought could take. These include images, propositions, and analogies. Assuming that we can think in all these different formats, then thoughts may assume multiple forms—thoughts sometimes may be pictures, sometimes propositions, and sometimes other symbolic representations.

But if we were to vote for the form that we believed thoughts are in most of the time, language would probably win. When we think, it is as if we can hear ourselves talking—what is called implicit speech. In contrast to the imagery that occupies "the mind's eye," implicit speech seems to occupy "the mind's ear." This supposed mental primacy of language has led some to conclude that thought and language are so similar that it may be impossible to express the thoughts generated in one language in another language. This is the strong version of the **linguistic relativity hypothesis**, which also goes by the name of the Sapir–Whorf hypothesis, after the linguist Edward Sapir (1884–1939) and his student, Benjamin Lee Whorf (1956). The weak version denies that such translation is impossible but admits that the language a person speaks influences the way he or she thinks.

Whorf studied the Hopi language, a Native American language, and found that the Hopi experience time as a discrete series, with each unit of time—say, days—considered unique and different from the others. This differs from the Western conception of time, wherein time is experienced as an undifferentiated continuous flow. Thus, a Hopi individual would not say "I stayed five days" but, rather, "I left on the fifth day" (Carroll, 1956, p. 216). The strong version of the linguistic relativity hypothesis would argue that the Hopi are incapable of thinking of time as continuous because they lack the words to express the concept in their language. The weak version would argue that the Hopi can understand this concept of time but that such an understanding would require a reexpression of it that used a completely different set of Hopi words.

So which version of the linguistic relativity hypothesis is correct? Investigation of the hypothesis has proceeded along two avenues. The first has focused on color perception, the second on counterfactual reasoning.

Davies and Corbett (1997) conducted research that tested English, Russian, and Setswana speakers. Setswana is a language spoken in Botswana, Africa. Davies and Corbett gave the participants colored chips and asked them to arrange the chips into groups based on chip similarity in any way they wanted. If the strong version of the hypothesis were correct, the Russians would place light- and dark-blue chips in separate groups, as their language has distinct color terms for light and dark blue. The Setswana speakers would group green and blue chips together, because they have a single term for green and blue. Their study and a follow-up study showed that all the participants tended to group the chips in pretty much the same way, regardless of their linguistic background (Davies, Sowden, Jerrett, Jerrett, & Corbett, 1998). The studies, therefore, fail to support the strong version of the hypothesis.

A counterfactual statement is a hypothetical one. It asks us to imagine what would happen if something were true. Counterfactuals can be expressed in two ways. One way is through the use of the subjunctive mood. The statement "If you bought your ticket now, you would save money" is an example of a statement that uses the subjunctive mood. The use of the word *would* is part of the subjunctive application and signals that it is an imaginary scenario that is being expressed. Other forms of the subjunctive use *were to* or *ought*. The other way to express counterfactual statements is through the use of "if–then" constructions. "If it is hot today, then I will go swimming" is an instance of this construction.

Bloom (1981) set out to test the linguistic relativity hypothesis by taking advantage of the fact that English speakers and Chinese speakers have different ways of expressing the counterfactual (see Figure 9.1). English has both subjunctive mood forms and if–then constructions, whereas Chinese has only if–then constructions. He predicted that the Chinese speakers would, therefore, have greater difficulty understanding the counterfactual. He presented both groups of speakers with a hypothetical scenario and found that the Chinese speakers failed to grasp the nature of the scenario, while the English speakers did not have difficulty. The results of his study support the hypothesis.

Both the color-naming and counterfactual-reasoning studies have been criticized on the grounds that tests for either color naming or counterfactual reasoning represent bad ways of testing the linguistic relativity hypothesis. Color perception is in large part the result of the physiology of the visual system, which is the same in everybody. This means that the way we see determines the names we have for colors—not the other way around. Consistency in color naming across populations and cultures supports this

Figure 9.1 A street in Chinatown, New York. Does speaking Chinese make you think differently?

idea (Berlin & Kay, 1969). Bloom's original study also has been criticized on the grounds that the scenario it employed was not translated well into Chinese (Au, 1983, 1984). When better translations were provided, Chinese speakers' comprehension of the scenario improved dramatically.

Evaluating the Linguistic Relativity Hypothesis

In summary, the two avenues of investigation fail to provide emphatic support for the strong version of the linguistic relativity hypothesis. However, more recent investigations of the mental representations of numbers (Miura, Okamoto, Kim, & Steere, 1993) and the use of classifier words such as *this* and *that* (Zhang & Schmitt, 1998) do provide limited support for the weak version of the hypothesis. More research is needed to establish a definitive answer to this issue. Currently, most investigators believe that languages are powerful enough and flexible enough to express any number of ideas. We can conclude that languages influence, but don't necessarily determine, the way we think.

We should keep in mind that language is but one way of thinking. As mentioned above, there are other forms of mental representation that are not linguistic and that are not governed by linguistic syntactical rules. The formation and processing of visual images seems to constitute an entirely nonlinguistic code for thinking. Mathematical thinking and the mental representation and computation of numerical quantities, although language like, may not rely on language mechanisms to operate and could constitute another distinct format for thought. The same can be said for the mental processing of music. Just because language is powerful and flexible doesn't mean it holds a monopoly on thought. If one idea cannot be expressed in terms of another linguistically, this might be achieved via the use of one of these other formats.

LANGUAGE USE IN PRIMATES

Animals in the wild communicate with one another. A monkey species that lives on the African savannah has a specialized series of cries that signify different kinds of threats. The monkeys use these cries while they are feeding to warn one another of impending danger. If one monkey in a group spies an eagle circling overhead, it emits one type of cry, which sends the members of the group scattering into the trees for cover. If it spots a snake, it emits another cry, which impels the monkeys to stand up on their hind legs and look around so that they can try to locate the snake. Each of the cries has a specific meaning. Each stands for a particular danger to the group. The meaning of the cry is understood by the rest of the group, as indicated by their reactions. This is communication because information about an event was produced, transmitted, and comprehended. But this natural form of communication is not language. The cries are arbitrary, there is no use of grammar to arrange them into any structure such as sentences, and they are not combined to create new meanings.

This raises an interesting question. If animals don't use language naturally on their own, can we teach it to them? Do they have the same capacity for language that we do? Research in this area has focused on primates, such as chimpanzees and gorillas, because of their relatively advanced cognitive capacities. Let's summarize some of this research and evaluate the results.

Early investigations of the linguistic abilities of primates focused on language production. Animals evidenced rudimentary language skills after being trained in one of several linguistic systems that included ASL, as well as a symbolic system employing plastic tokens and one that used geometric patterns called **lexigrams**. Starting in the 1960s, Beatrice and Allen Gardner raised a chimp named Washoe. They taught her to use ASL. Their method was to get Washoe to imitate or reproduce the hand formation that stood for a particular object (Gardner, Gardner, & Van Cantfort, 1989). Washoe learned 132 signs and seemed to show evidence of spontaneous language use. On seeing a toothbrush in the bathroom, she made the sign for it without being prompted. A similar technique was used to teach ASL to a gorilla named Koko (Patterson, 1978). Koko learned an even larger repertoire of signs and was reported to have used syntax and to have made signs spontaneously. Her trainer claims that she even told jokes!

David Premack has used a different approach. He used plastic tokens instead of hand signals as he attempted to teach language skills to a chimp named Sarah (Premack, 1976). The tokens had different shapes and colors and stood for individual words as well as relationships. There were tokens that stood for nouns (*apple*), for verbs (*give*), for adjectives (*red*), and for relationships (*same as*). Sarah produced the "same as" token when she was presented with two "apple" tokens and the "different" token when shown an "apple" and an "orange" token. She seemed to have a rudimentary understanding of sentence grammar, as she was apparently able to tell the difference between two sentences such as "David give apple Sarah" and "Sarah give apple David."

Savage-Rumbaugh et al. (1993) studied a chimp named Kanzi who appeared to have learned the meanings of lexigrams. In addition, Kanzi was apparently able to understand single-word and simple-sentence utterances made by humans. Kanzi's abilities seemed quite advanced. Of his own accord, he would use lexigrams to identify objects, to make requests for food items, and to announce a particular action he was about to undertake. Following more structured language training, Kanzi's abilities were compared with those of a 2½-year-old child named Alia. Both were given novel commands that required them to move objects. In terms of comprehension, the two showed nearly identical abilities: They both demonstrated approximately 70% compliance with the commands. Kanzi's language production skills were more limited—they corresponded to those of a 1½-year-old child (Greenfield & Savage-Rumbaugh, 1993).

Evaluating Language Use in Primates

At this point, we can examine some of the criticisms that have been leveled at this research. Some of the animals described so far were trained via the use of positive reinforcement. They were given a reward, usually a food item, for making the correct sign or using the appropriate token or lexigram. A problem with this is that the animals may have been associating a symbol with a concept because they had been trained to do so and may not have had any awareness that the symbol actually stood for something. If this was the case, these animals fail to demonstrate the arbitrariness aspect of language—that the symbol can be anything and still stand for its referent. An animal's choosing of an "apple" token when presented with an apple does not demonstrate arbitrariness, but using an "apple" token to refer to an apple when an actual apple is not perceptually present does. This aspect of language, in which users refer to something that is removed in space or time, is called **displacement**.

Savage-Rumbaugh (1986) presents some evidence of displacement in chimps. She employed a technique known as cross-modal matching wherein chimps who viewed a lexigram were then required to select the object the lexigram represented via the use of touch from a box filled with objects. The chimps were able to do this, which indicated that they understood what the lexigrams represented. A note of caution is in order, however. Arbitrariness and displacement capabilities were shown for a comprehension task only, where the animals were interpreting the symbols. Earlier studies in which chimps used ASL and tokens have generated less evidence that primates understand the meanings of symbols when they produce them.

What about the structured aspect of language? Do animals understand the syntax that underlies the formation of sentences? The investigations that have been conducted thus far show that primates comprehend and produce very simple sentences—sentences that are in the order of two or three words long. An understanding of the rules of syntax is demonstrated by the ability to rearrange words in new combinations that express new meanings—the generative criterion, defined above. If animals could do this, it would indicate a comprehension of syntactical rules.

The researcher Herb Terrace provides us with evidence that refutes the idea that some animals may have a rudimentary understanding of or the ability to use syntax (Terrace, Petitto, Sanders, & Bever, 1979). Terrace was skeptical that chimpanzees such as Washoe truly understood the meanings of the signs and symbols they used. As alluded to above, he believed that chimps used hand signals or presented tokens because they had been reinforced for doing so. To test the idea, he studied a chimpanzee who he had jokingly named Nim Chimpsky. Nim was raised in a human family and was taught ASL. Rather than use food as a reward, Terrace gave approval that centered on things that were important to Nim. Under this system, Nim did seem to have some grasp of the meanings of his signs, as he was found using them in the absence of their referents. He also appeared to use signs spontaneously to express his desires. For example, he would make the sign for *sleep* when he was bored. However, Terrace concludes that Nim was never able to combine his signs to form sentences and express novel meanings. He did this only when he was directly imitating combinations of signs that had been produced by his trainers.

So, at this point, we can sum up the work on the language abilities of primates. Primates appear to possess some arbitrariness and displacement capabilities because they can comprehend the meanings of a limited number of symbols independent of their referents. This is true whether they are trained directly (with food) or indirectly (with approval) using positive reinforcement techniques. But here is where their linguistic skills seem to come to an end, for primates seem to understand very little in the way of syntax, especially when it comes to language production. They know that some aspects of word order affect meaning. However, most primates tend to repeat the sentences they were taught or produce only small variations on them. They do not come anywhere near to possessing human generative capability. Also, unlike humans, primates—once they have acquired language skills—fail to teach the skills to other members of their species. Unfortunately, the "Doctor Doolittle" scenario of our conversing with animals the way we do with one another just doesn't seem possible.

LANGUAGE ACQUISITION

Clearly, a human being is not born with an ability to speak his or her native language fluently. This ability develops over time. Linguists adopting a developmental perspective have studied the acquisition and development of language skills, from birth through infancy, childhood, and adolescence. They have shown that human beings pass through a series of stages, each one marked by the acquisition of new linguistic skills.

Early in the first year, infants start to utter a wide variety of sounds. All infants do this. At this stage, they begin to exercise their vocal cords and mouths—major parts of the vocal apparatus—the use of which they must master in order to articulate the sounds of the language they are just beginning to acquire. This period of development is known as the **cooing stage**. Figure 9.2 shows an infant communicating with her caregiver.

At around 6 months, the number of sounds a baby produces shrinks. The sounds produced at this stage are consonant–vowel pairs, such as "mama" and "dada." The majority of the utterances made at this time are more phonemic than morphemic in nature. They correspond to sound units rather than to fully pronounced words. However, the intonations of these utterances at this point begin to match those of the language the child is learning. Intonation refers to the rises and falls in pitch and changes in other acoustic properties of one's speech. For example, for most speakers, there is usually a rise in pitch at the end of a question. These abilities arise during the so-called **babbling stage**.

Following the babbling stage and just a few months shy of the child's first birthday, we see the advent of one-word utterances. At this point, children are able to successfully articulate entire morphemes. These morphemes, or words, may not be prescriptively accurate. A child may say "unky" instead of "uncle," but the utterance is being used in a meaningful way. The children are, thus, beginning to use language in a symbolic and semantic fashion. This is the **one-word stage**.

Following this, during the **two-word stage**, children produce two-word utterances. It is at this point that they will say things like "see kitty" or "want toy." Because words are now being arranged into simple sentence-like structures, the two-word stage marks the

Figure 9.2 Early in development, infants will spontaneously babble phonemes.

emergence of rudimentary syntactical skills. After the two-word stage, babies will string together more complicated utterances, speaking out the equivalent of sentences composed of three or more words that convey increasingly complex meanings.

There are no clearly identifiable stages that follow the two-word stage. But this period is characterized by a steady growth in vocabulary and syntax. Also, during this period, children exhibit some interesting patterns of development, especially with regard to their learning of past-tense forms (Kuczaj, 1978; Marcus et al., 1992). Studies of this type of learning show that children first imitate past-tense forms correctly. For the irregular verb *to go*, the past-tense form is *went*. After children have learned the general rule of past-tense formation, they apply it correctly to regular verbs but then overextend the rule to include irregular verbs as well—saying, for example, "goed" instead of "went." Finally, they learn the exceptions to the rule—for example, using *went* only when it is appropriate. This intriguing U-shaped pattern of development indicates the presence of several learning strategies in children: They start out with purely imitative copying, proceed to an understanding of a rule, and ultimately progress to the learning of the exceptions to that rule.

Domain-General and Domain-Specific Mechanisms in Language Acquisition

Domain-specific mechanisms are those devoted to the processing of a single type of information like linguistic information. **Domain-general mechanisms** are generic and can process different types of information. For instance, a domain-general mechanism would be capable of processing both linguistic and visual information. The incredible sophistication of language and the fast rate at which it is learned imply that it is a domain-specific mechanism. However, it is possible that domain-general learning mechanisms contribute too. This issue has been debated in the language research community for decades and, as we will see in this section, has yet to be resolved. However, it is a good example of how evidence in the sciences can support both sides of a position.

The domain-general versus the domain-specific debate is actually unrelated to the nature/nurture or nativist/empiricist debate (Saffran & Thiessen, 2007). Just because a mechanism seems to be domain specific, like the one we see for language, does not mean that it is hardwired or innate. It also does not mean that it is a module, as you will recall that these are some of the characteristics of modules. By definition, all mental processes require an innate structure to operate. Artificial neural networks are domain general. As we saw in the network chapter, they can form the foundation for all sorts of mental processes, but to function, they must have preexisting structure, including the organization of nodes and links into layers, rules by which inputs are summated, and so on. Likewise, domain-specific mechanisms can be learned as seems to be the case for those brain areas like the angular gyrus that underlie reading and writing.

The fact that language is localized to the left hemisphere in most people suggests that it may be domain specific. In addition, identical sounds are actually processed differently at the neural level depending on whether they are perceived as speech or nonspeech (Dehaene-Lambertz et al., 2005). So at both an anatomical and physiological level, we see differences in the way language-like stimuli are interpreted. However, it could just be that the left hemisphere is good at processing stimuli that are short-lived and not language information per se (Zatorre, Belin, & Penhune, 2002).

Infants are predisposed to respond to speech and will pay attention to it in preference to other sorts of sounds (Vouloumanos & Werker, 2004). This preference ensures that speech information "gets into" the system and may help accelerate the rate at which it is learned. If this were the case, then this attentional bias could be domain specific and perhaps innate, while the subsequent learning mechanisms may be domain specific, domain general, or both. This illustrates that learning need not be just one or the other. One mechanism could set the stage for another, or the two different categories of mechanisms could interact with each another.

Determining the boundaries between spoken words is difficult because there are often no pauses. Yet infants can extract words from speech as early as 7 months. One way to do this is by learning which sounds tend to follow one another. The use of this type of information is known as **statistical learning**. If one sound tends to follow another more often within rather than between words, then this can be a cue to indicate which sounds group together into words. Both adults and infants are capable of utilizing statistical learning to determine word boundaries (Saffran, Aslin, & Newport, 1996). Statistical learning, however, seems to be domain general. It is used for visual stimuli and auditory stimuli that are not linguistic in nature. It is also used in several nonhuman animal species.

The syntax or rules that govern language would seem to be clearly within the domain-specific camp. There are several reasons for this. First, it has been difficult to teach animals human language (some would argue any type of language), suggesting that they lack the prerequisite mechanisms. Second, languages around the world do not vary much with regard to certain aspects of syntax (Baker, 2001). Third, syntax cannot easily be derived from the perceptual characteristics of a language. Despite all this, some recent work shows that word categories (noun, verb, adjective, etc.) can be learned by the order in which they occur (Mintz, Newport, & Bever, 2002). For instance, nouns tend to follow the word *the*. So yes, much of the evidence we have encountered so far can be interpreted either way. It is safe to say that the jury is still out on whether language is a domain-general or domain-specific mechanism.

Evaluating Language Acquisition

Saffran and Theissen (2007) make a number of important conclusions regarding domains and learning. They argue that our conception of general and specific are too strict. "Domain-general" mechanisms may not be open to any type of input but are more likely to be open to some but not to others. Similarly "domain-specific" mechanisms imply modularity and innateness, and they need not be either. Modules rather than being innate could emerge from learning within a given domain. McMullen and Saffran (2004) argue that adult cognitive abilities, although localized in the brain to some extent, could be the product of domain-general mechanisms that were developmentally "fed," modality-specific information. Fruitful approaches that may shed light on this debate would be to see what changes (neural, cognitive, or behavioral) occur early on versus later to disentangle the initial states from future outcomes as well as the application of technological advances to see which brain areas are at work for different types of knowledge in early infancy.

LANGUAGE DEPRIVATION

What, then, is the role of environment in language acquisition? One very basic approach to this question is to examine language ability in the absence of exposure to language. If experience of language is necessary and innate linguistic mechanisms are dependent on it, we should see language deficits in the absence of exposure to a normal linguistic environment. If, on the other hand, experience and stimulation have little to do with language development, then language ability should remain relatively intact in their absence. Studies that have investigated this issue have demonstrated the existence of a pivotal time in development during which language must be learned. Children not exposed to language during this time, called the **critical period**, may never acquire it or may suffer severe language impairments.

There is abundant evidence in support of the idea of a critical period. Let's examine some of it. The first evidence comes from studies of birds. It turns out that some birds (as well as human beings) are among the few animals that need to be exposed to the communicative sounds of their own species in order to be able to produce them (Doupe & Kuhl, 1999). Both birds and human beings are better able to acquire communicative ability early in life. Marler (1970) showed that white-crowned sparrows after the age of 100 to about 150 days were unable to learn new songs by listening to a "tutor" bird. This was true for birds that were raised with such tutors or were exposed to taped examples of bird song, as well as those that were raised in acoustic isolation.

Another line of evidence that supports the existence of a critical period comes from studies of the acquisition of second languages. So far, we have been discussing first-language acquisition, wherein a single language is learned. It is often the case, of course, that people learn to speak more than one language. Researchers can study the difficulty with which an individual acquires a second language in relation to the time of onset of exposure to the second language. One such study found that native speakers of Chinese and Korean (for whom English was a second language) received scores on tests of English grammar that bore a relation to the time of onset of their exposure to English: The later their age at time of arrival in the United States, the lower were their scores (Johnson & Newport, 1989). Figure 9.3 shows the results of this study.

The most emotionally compelling evidence that supports the existence of a critical period consists of individual case studies of persons who were deprived of language experience during early development. These cases are tragic but provide a unique opportunity to examine the effects of this kind of deprivation in humans. One famous historical case study is that of the wild boy of Aveyron, who was discovered in a French forest in 1797 (Lane, 1976). The boy, named Victor, had apparently lived much of his childhood completely alone and had very little language ability. He came under the supervision of a physician, Dr. Jean-Marc-Gaspard Itard (1775–1838), who studied him intensively and attempted to teach him language. Despite Dr. Itard's best efforts, Victor never acquired more than the most basic comprehension and production skills.

Figure 9.3 Mean English grammar test scores drop in correlation with greater ages of children who speak another language at the time of arrival in the United States.

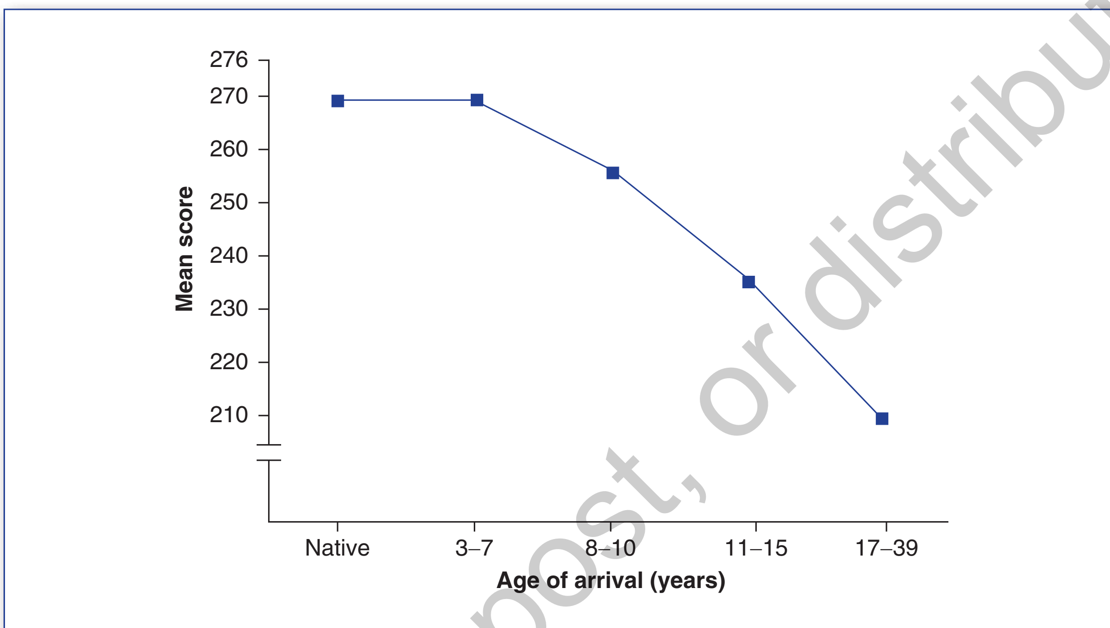

Source: Johnson and Newport (1989).

A more recent case study is that of a girl named Genie. Genie spent much of her early life in social isolation. The evidence suggests that the period of her deprivation began when she was 20 months old and lasted until she was "discovered," at the age of 13 years and 9 months. During this time, Genie was kept in a small room, where she was tied to a potty chair for much of the day and night or confined to an infant crib. The door to the room was almost always closed, and the windows were closed and covered with curtains. Except for her quick feedings, Genie received almost no parental care. She was not spoken to and there was no radio or TV in the household. Thus, she was exposed to little or no spoken language. Furthermore, Genie was beaten for making any sounds of her own.

Given this extreme and prolonged lack of exposure to any sort of linguistic stimulation, what were Genie's abilities like? A number of researchers have chronicled her progress (Fromkin, Krashen, Curtiss, Rigler, & Rigler, 1974; Jones, 1995). On initial examination, she was found not to vocalize at all. Within a few days, she began to respond to the speech of others and to imitate single words. Within a year or so, she was able to understand and produce some words and names. Despite these modest gains, it was clear at the end of the testing period that Genie possessed minimal grammatical ability. Following a period of several years' worth of further evaluation and training, Genie did show signs of simple grammatical comprehension. For example, she was able to distinguish between the singular and plural forms of nouns, and between negative and affirmative sentences.

At 8 months after discovery, Genie uttered two-word phrases such as "yellow car." Later, she was able to produce three- and four-word strings, such as "Tori chew glove" and "Big elephant long trunk." She also demonstrated generativity—she was able to express new meanings by combining words in novel ways.

These capacities show that Genie was able to acquire language in the aftermath of the deprivation period and that the stages of her development—such as her use of progressively longer sentences—paralleled language acquisition in nondeprived children. However, Genie's abilities deviate from those of control children in several ways. Her grammatical ability at the time of early testing was equal to that of a 2½-year-old child, and her speech production capacity was limited. In addition, Genie's rate of language development was slowed in comparison with that of controls. She had difficulty using language to express questions, and many of the hallmarks of language mastery in adults—such as the use of demonstratives, particles, rejoinders, and transformation rules—were absent (Fromkin et al., 1974). In summary, Genie shows that language acquisition following extended deprivation is possible but is severely impaired. To date, Genie has not developed complete adult language skills, and she probably never will.

Evaluating Language Deprivation

Case studies of language-deprived children yield a wealth of information about the individuals under study. They do, however, suffer from a number of problems. To begin with, it is difficult to make generalizations from evidence acquired from a single person or a small number of persons. The findings of case studies do not necessarily generalize to a larger population. Second, the conditions that shaped these subjects are often unknown. In the case of Victor, we do not know the duration of his social isolation, or even if he was isolated at all. Some have speculated that he may have simply had a learning disability or suffered brain damage. With regard to Genie, it is not clear exactly what kind of language information she was exposed to during her formative years, nor the extent to which she may have vocalized to herself.

Research that has investigated the critical period shows that although there may be an innate language-learning mechanism, it is dependent on environmental input for its proper functioning. If this input is absent, the ability to fully use language never appears. Exposure to and practice in the use of a language is, thus, a component essential to the development of language. This is true regardless of the amount or the sophistication of the neural machinery dedicated to language processing that may be in place in an individual from birth.

COGNITION AND LINGUISTICS: THE ROLE OF GRAMMAR

We said earlier that a set of rules governs how words can be arranged in sentences. Grammar is important because it tells us what is a proper way of expressing something in a language and what is not. If there were no rules or constraints on expression, we could string words together in practically any order, and it would be impossible to convey anything. Let's delve a little further into grammar—that is, how it puts constraints on what can be said and how it illuminates several interesting cognitive principles.

Sentences have distinct parts that are hierarchically related. This organization is called a **phrase structure** and can be illustrated via the use of tree diagrams. Figure 9.4 is a tree diagram for the sentence "The big dog chased the black cat." At the highest level, the entire sentence (S) is represented. Moving down one level, the sentence is composed of two parts, a noun phrase (NP) and a verb phrase (VP). Moving down another level, we see that the noun phrase is made up of a determiner (D), an adjective (A), and a noun (N). Determiners are words such as *a* or *the*. The verb phrase is made up of a verb (V) and another noun phrase that itself contains another determiner, adjective, and noun.

Figure 9.4 The phrase structure for a simple sentence.

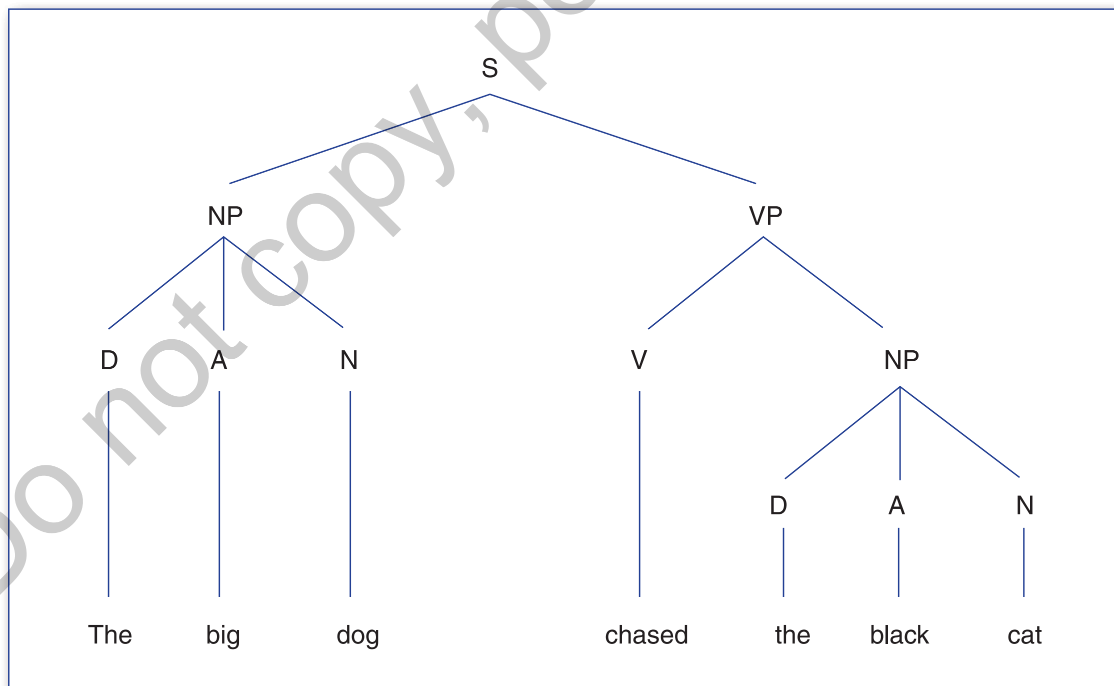

There is a grammar that governs the use of phrase structures. This **phrase structure grammar** imposes certain limitations on how a legitimate sentence can be put together. One phrase structure rule is that all sentences are composed of a noun phrase and a verb phrase. A second rule is that noun phrases consist of a determiner followed by a noun. Verb phrases can be expressed a bit more flexibly—as a verb followed by a noun phrase, another sentence, or other elements.

A phrase structure grammar is useful for understanding the organization of sentences, but it doesn't tell us how we can rearrange a sentence to express new meanings. Noam Chomsky (1957) was the first to point this out. He notes that a given sentence can be changed in three ways. First, we can turn an active sentence into a passive one, as when "The man read the book" becomes "The book was read by the man." Second, we can turn a positive statement into a negative one, by modifying the original sentence to form "The man did not read the book." Third, we can convert the assertion into a question, as in "Did the man read the book?"

To account for these changes, we need a new grammar that allows us to transform one sentence into another. Chomsky's solution was a **transformational grammar**, a set of rules for modifying a sentence into a closely related one. By using these rules, we can reorder "The man read the book" into "The man did not read the book," as follows:

$$\text{NP1} + \text{Verb} + \text{NP2} \rightarrow \text{NP1} + \text{did not} + \text{Verb} + \text{NP2},$$

where NP1 is "The man," the verb is "read," and NP2 is "the book." Similarly, the conversion of "The man read the book" to "The book was read by the man" is denoted as

$$\text{NP1} + \text{Verb} + \text{NP2} \rightarrow \text{NP2} + \text{was} + \text{Verb} + \text{by} + \text{NP1}.$$

An important aspect of a transformational grammar is that one can use it to express two sentences that have different phrase structures but identical meanings. "The man read the book" and "The book was read by the man," above, have different hierarchical organizations, but they have the same semantic content. To account for this, Chomsky proposed two levels of analysis for sentences. The **surface structure** is the organization of the sentence in the form that it is expressed in—that is, how the sentence would be heard if it were spoken or read if it were written. The surface structure is variable and can be rearranged by transformational grammar. The **deep structure** is the underlying meaning of a sentence and remains constant regardless of the specific form in which it is expressed.

You may have been wondering whether our discussion so far applies only to English or whether it applies to other languages as well. It is true that languages have different specific rules, but cross-cultural linguistic analyses have shown that languages have a number of elements in common. These commonalities are summed up in the concept of a **universal grammar**, which comprises the features that are instantiated in the grammars of all natural languages (Chomsky, 1986; Cook, 1988). In this view, each individual language at a fundamental level is not really different from others but represents merely a variation on a theme. Universal grammar is considered as a collection of language rules, hardwired into our brains from birth. In this sense, it is a modular aspect of mind and has all the characteristics of a mental module. It is innate, genetically prespecified, domain specific, and independent of other cognitive capacities.

What are the universal properties of all languages then? One is a phonological rule that specifies the ordering of syllables in a word. According to the **maximal onset principle**, consonants usually precede vowels; more frequently than not, they constitute the onset of syllabic groupings. This feature is found in all languages. Another universal property of language is syntactical and concerns the ordering of the subject and object in sentences. In 98% of the world's languages, the subject precedes the object (Crystal, 1987). Thus, we say, "John kicked the ball," not "A ball John kicked"—even though the latter form is technically acceptable in English.

Universal grammar may be what is responsible for our ability to acquire language so quickly. Language acquisition requires the mastery of a large number of grammatical rules at different levels. There are actually sets of rules, including phonology, to determine acceptable combinations of phonemes; morphology to determine which morphemes go together; syntax for the ordering of words in sentences; transformation rules for changing the forms of sentences, and so on. The ease and rapidity with which this process occurs in humans can be explained if it is true that at least some generic versions of these rules are already present in the head at birth. A child would then adapt these general linguistic rules to the particularities of the specific language he or she grows up in (Bloom, 1994).

Evaluating Universal Grammar

The idea of a universal grammar, or "language organ," as originally formulated by Chomsky has not gone unchallenged. To begin with, there is little evidence to support the notion of specific genes for language. If one looks at other body organs, there are few that owe their existence to individual genes. So it is unlikely that there are specific genes devoted to language processing. There is also doubt about the domain specificity of any proposed language module. The rules governing language use may be more general; they may manifest themselves in other nonlinguistic cognitive capacities. One possibility is that linguistic universals are just the product of general biological mechanisms, implying that language is not "special" in any sense.

NEUROSCIENCE AND LINGUISTICS: THE WERNICKE-GESCHWIND MODEL

Paul Broca (1824–1880) was a French surgeon who worked with patients who had suffered brain damage as a result of stroke or injury. The patients demonstrated various kinds of language deficits, called **aphasias**. Several of his patients had severe difficulty articulating speech. One famous patient was capable only of uttering the word *tan* over and over again. For the most part, these patients could understand what was being said to them, indicating that the faculty of comprehension was intact, but they had problems pronouncing words and producing speech. This deficit is called **Broca's aphasia**. It is also known as nonfluent aphasia.

Patients with Broca's aphasia produce what is called "agrammatic speech." They generate strings of nouns and some verbs but without any filler words, such as *the* or *is*. They also fail to make words plural or to use verb tenses. Their sentences are short and broken by many pauses, which sometimes have earned this kind of speech the nickname "telegraphic" or "nonfluent" speech. The following is an example of the speech of a patient talking about a visit to the hospital for dental surgery.

> Yes . . . ah . . . Monday er . . . Dad and Peter H . . . and Dad . . . er . . . hospital . . . and ah . . . Wednesday . . . Wednesday, nine o'clock . . . and oh . . . Thursday . . . ten o'clock, ah doctors . . . two . . . an' doctors . . . and er . . . teeth . . . yah. (Goodglass & Geschwind, 1976, p. 408)

Postmortem examination of the brains of patients who suffered from Broca's aphasia has revealed damage to the lower portion of the left frontal lobe (see Figure 9.5). This region is believed to be at least partly responsible for language production capacity and has been named **Broca's area**.

A second area, named after Carl Wernicke (1848–1905), mediates language comprehension. This area is located in the posterior region of the left hemisphere (see Figure 9.5). Patients with damage to **Wernicke's area** suffer from **Wernicke's aphasia**. They produce rapid, fluent, and seemingly automatic speech that has little meaningful content. For this reason, this aphasia is also referred to as fluent aphasia. This type of speech sounds normal in the sense that its rate, intonations, and stresses are correct—but it is lacking in content or meaning. These patients have major problems comprehending speech and also demonstrate difficulty reading and writing. Here is an example of the speech of a patient with Wernicke's aphasia.

> Oh sure, go ahead, any old think you want. If I could I would. Oh I'm taking the word the wrong way to say, all of the barbers here whenever they stop you it's going around and around, if you know what I mean, that is tying and tying for repucer, repuceration, well, we were trying the best that we could while another time it was with the beds over there the same thing. (Gardner, 1974, p. 68)

Figure 9.5 Brain areas of the left hemisphere that are part of the Wernicke-Geschwind model of language comprehension and production.

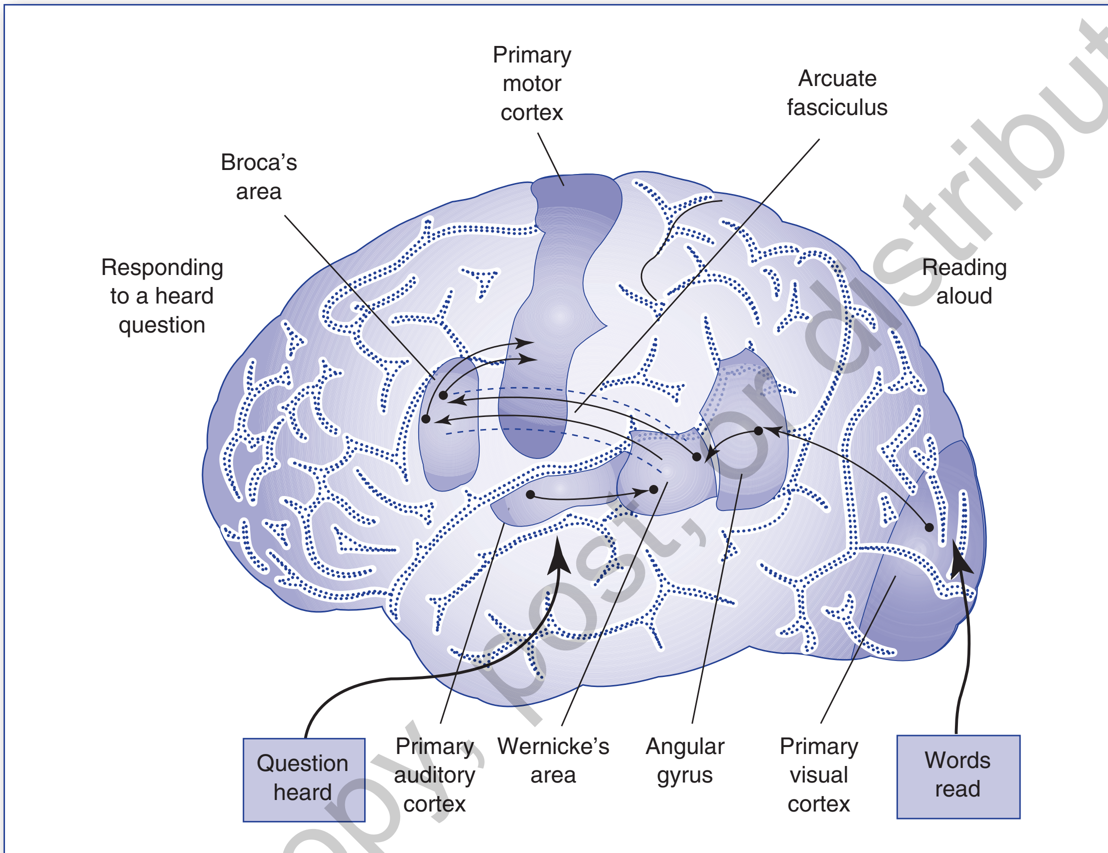

The Wernicke-Geschwind model was first formulated by Wernicke and expanded in the 1960s by Norman Geschwind (1972). It specifies the functional roles of the different brain areas that are involved in language processing, as well as their connections and interactions. Because an understanding of the model relies heavily on an understanding of different cortical areas, we must first introduce these anatomical regions. You may wish to refer back to the neuroscience chapter at this point for a refresher on basic brain anatomy.

Figure 9.5 shows the cortical areas that play the key roles in language processing—as described in the model. For starters, there is the primary motor cortex, located in the frontal lobes in the anterior part of the brain. Commands that originate here send impulses to muscles, causing them to contract and, therefore, initiating movement; this includes the muscles of the mouth, which must be moved as part of the operation of speaking. The primary visual cortex is located at the back of the brain in the occipital region. It is where visual information is first processed. This area becomes active during reading and writing. The primary auditory cortex is situated in the temporal lobes. It is here where sounds striking the ears are first processed. The **arcuate fasciculus** is a pathway that connects Broca's area and Wernicke's area. Damage to this part of the brain results in an individual's difficulty in repeating words that he or she has just heard, known as **conduction aphasia**. Finally, there is the **angular gyrus**, located behind Wernicke's area. Damage to this part of the brain produces **alexia**, an inability to read, and **agraphia**, an inability to write.

According to the model, these areas and the pathways that connect them subsume language comprehension and production with respect to both the auditory and visual modalities. There is an activation of neural pathways that is the basis of listening and speaking. It is as follows: The perceptual characteristics of speech sounds are first processed in the primary auditory cortex. The output of this processing is then passed to Wernicke's area, where the content of what has been said is processed and understanding is born. A reply is then initiated. From here, the information that will become the reply is passed along the arcuate fasciculus to Broca's area. The information is converted into a motor code, or program of articulation, within Broca's area. This code is then passed to the primary motor cortex, where commands to move the muscles of the mouth and produce the speech that constitutes the reply are executed.

A second pathway mediates reading and writing. In this pathway, the primary visual cortex processes inputs that have originated from the words on a printed page. This information is then output to the angular gyrus. The visual representation of what has been read is converted into an auditory code within the angular gyrus, which then sends the code to Wernicke's area. The remainder of this pathway, responsible for producing behaviors such as reading out loud or writing, coincides with the final portion of the pathway described in the preceding paragraph. The information flow would be from Wernicke's area via the arcuate fasciculus to Broca's area, and then to the primary motor cortex, where muscular action is initiated.

Evaluating the Wernicke-Geschwind Model

The Wernicke-Geschwind model has been criticized on a number of counts. It is considered by some to be an oversimplification of the neural basis for language. To begin with, the areas specified by the model are not completely associated with their hypothesized function. Although in most patients damage to Broca's area or Wernicke's area results in the corresponding aphasias, this is not always the case. Lesions to Broca's area alone produce a transitory aphasia—one that presents with only mild symptoms several weeks after the event that precipitated the injury (Mohr, 1976). More troublesome to the theory are the records of patients with Broca's aphasia who have not sustained damage to Broca's area (Dronkers, Shapiro, Redfern, & Knight, 1992). The same is true for patients with Wernicke's aphasia (Dronkers, Redfern, & Ludy, 1995).

The areas specified by the model are characterized as being modality specific, with Broca's area being a motor-only area that codes for speech articulation and Wernicke's area being an auditory, sensory-only area devoted to speech comprehension. However, brain imaging techniques show that these regions are the sites of processing activities that underlie sign language use (Bavelier et al., 1998). This suggests that they may represent more abstract, modality-independent language ability. In other words, these areas may contain knowledge of syntax that can be applied to any language system, regardless of the modalities involved.

Another criticism of the model centers on its original assumption that these areas are devoted exclusively to linguistic processing. Swinney, Zurif, Prather, and Love (1996) have found that lesions in aphasic patients, even those suffering "classic" syndromes such as Broca's aphasia and Wernicke's aphasia, may have disrupted basic processing resources used by the language system. If this is true, damage to the brain areas thought to subsume language only may also lie behind disruptions of other systems that language depends on, such as memory and attention.

Another problem with the Wernicke-Geschwind model is methodological. It was based largely on evidence obtained from clinical case studies of brain-damaged patients, assembled after their deaths. Modern science relies more on brain imaging in live patients. The use of brain imaging techniques has shown that there are many other brain areas that contribute to language function. We can list a few of them here. The insula lies beneath the frontal, temporal, and parietal lobes. Most patients with Broca's aphasia also have lesions in the insula (Vanier & Caplan, 1990). The left inferior prefrontal cortex, just anterior to and ventral to Broca's area, is activated during semantic retrieval (Peterson, Fox, Posner, Mintun, & Raichle, 1988). The basal temporal areas, at the bottom of the left temporal lobe, and the cingulate gyrus are also involved in word retrieval. The anterior superior temporal gyrus, anterior to the primary auditory cortex, is implicated in sentence comprehension. These areas are just beginning to be understood. They are believed to interact as parts of a complex network. There is as yet no overarching theory that can describe this interaction. Until then, the Wernicke-Geschwind model provides a useful, if somewhat outdated, understanding of what goes on in the brain during language processing.

ARTIFICIAL INTELLIGENCE AND LINGUISTICS: NATURAL LANGUAGE PROCESSING

Natural languages are those that have evolved in human societies and are used by human beings. Examples of natural languages are English, Spanish, and French. These are in contrast to formal computer languages such as C++++, or linguistic expressions of logic.

There are two kinds of natural language processing. Understanding a natural language involves an individual's assimilation of linguistic expression in some form, such as speech or writing; extracting its meaning; and then undertaking some action that constitutes a response to this meaning. Understanding is what a computer would need to do if it were to interpret a spoken human command and act on it. Generation is the reverse of this process. It involves taking a formal symbolic representation of an idea and converting it to an expression in English or some other natural language. For example, the idea "It is a sunny day" may initially be stored in a particular format in a computer. A computer would be generating language if it could transform this idea into a spoken utterance that a human being could understand. These two processes are, thus, the computer equivalent of natural language comprehension and production. In this section, we will concern ourselves exclusively with natural language understanding, as that is the area in which research has been concentrated.

Cawsey (1998) outlines four stages of natural language understanding. We will preview each of them, in the order in which they occur:

1. **Speech recognition** is the first step in the process, whereby the acoustic speech signal is analyzed to determine the sequence of spoken words.
2. In **syntactic analysis**, the word sequence is analyzed via the use of knowledge of the language's grammar. This yields the sentence structure.
3. Following this, the sentence structure and the meanings of the words are used to derive a partial representation of the meaning of a sentence. This is the **semantic analysis** stage.
4. **Pragmatic analysis**, the final stage, produces a complete meaning for the sentence via the application of contextual information. This information includes data that have to do with the time and location of the utterance, who was saying it, and to whom it was said.

Speech Recognition

Speech recognition by a machine is a laudable aspiration. Wouldn't it be nice to be able to talk to our computers instead of having to type in commands or use a mouse? Humans use language quickly and effortlessly to communicate ideas to one another. To be able to communicate in a similar way with computers would usher in a new age of efficiency and productivity. Unfortunately, the task of getting a machine to understand speech is much more difficult than it may seem. Let's review some of the steps that speech recognition by a machine would have to include and talk about the problems involved.

Any attempt at speech recognition starts with a speech spectrogram. A **speech spectrogram** is a visual representation of the speech signal; it is a graph that displays the component frequencies of a speech sound over time (see Figure 9.6). From this, a computer program then attempts to extract the phonemes from the segment of speech under analysis. If a phoneme is ambiguous, the segment of the speech signal that it occupies can be matched against similar utterances that have been recorded and analyzed to "fill it in." The phonemes are then assembled into their corresponding words. This is accomplished in part by a statistical analysis that factors in the probabilities that specific words will crop up in speech, that specific phonemes will crop up in specific words, and that specific words will be surrounded by other specific words.

A phoneme-to-word assignment is difficult for two main reasons. The first of these concerns word boundaries. It turns out that there are no pauses between words in spoken speech. This makes it hard to tell where one word starts and another ends. To compound the problem, there are often pauses within words. So pauses cannot serve as reliable indicators of word boundaries. The second major issue is phoneme variability. If each phoneme

Figure 9.6 A speech spectrogram.

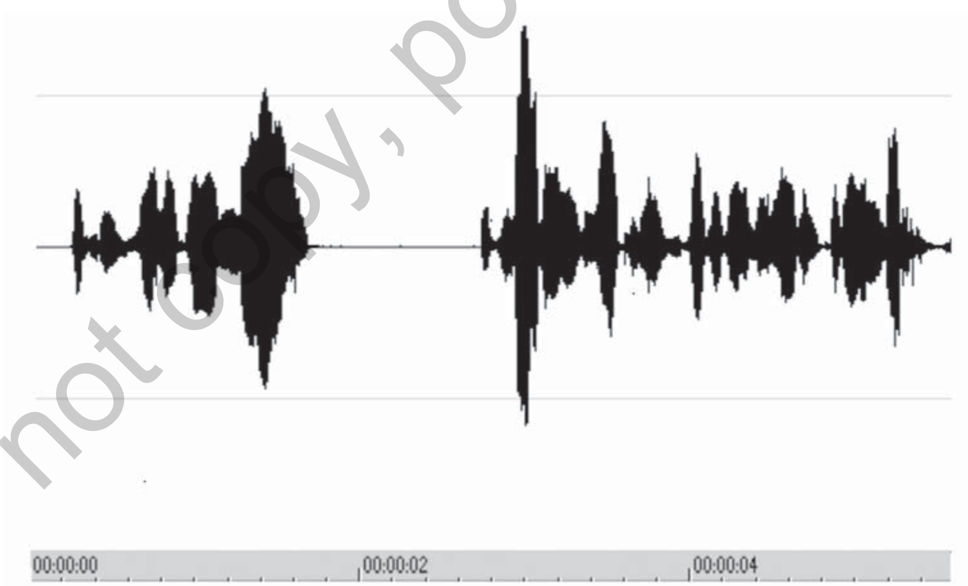

Time (secs)

This is one of our more . . . [pause] . . . exciting lectures that we're going to be having today.

were pronounced clearly and uniformly, speech recognition would be much easier. This is not the case. Speakers vary tremendously with respect to the pitches and durations of their phonemes. They are apt to pronounce a given phoneme variably, and in ways that depend on which phonemes precede it and which come after it. This is known as **coarticulation**. Additional complicating factors are the presence of background noise and the fact that in English a single sound—for example, that is represented by *bear* and *bare*—can belong to more than one word.

People resolve these difficulties by taking into account the overall meaning of a sentence. In one classic study, Warren and Warren (1970) presented participants with recordings of sentences in which a cough sound was substituted for a phoneme. One of the sentences was "It was found that the *eel was on the axle." (The asterisk indicates where the cough sound was inserted.) In other versions of the sentence, the word *axle* was changed to *shoe*, *orange*, and *table*. Asked to interpret the four sentences, the subjects heard the ambiguous word as *wheel*, *heel*, *peel*, and *meal*, respectively. This clearly demonstrates that the meanings of words in a sentence that have already been understood provide a framework for understanding the words that have yet to be understood. It also shows that in human speech perception, recognition is top-down as well as bottom-up, as the meaning of the entire sentence is pulled together simultaneously with the meanings of individual words.

Humans also have the benefit of visual cues when they are taking in speech. We can look at a speaker's lips as he or she is speaking. The positioning of the lips can help us interpret difficult phonemes or morphemes. Because some deaf individuals can understand speech by way of lip reading, there is obviously more than enough information in the visual aspect of speech to enable comprehension. Most computer speech recognition systems in use today must rely on auditory information as their only input and so do not have visual cues.

Syntactic Analysis

Once the individual words and their order have been determined, we can analyze the speech stream at the sentence level. This analysis entails the use of grammar. We already have discussed the various grammars that govern sentence structure. Syntactical analysis programs perform the equivalent of using a phrase structure grammar to evaluate a sentence and to break it down into its hierarchical constituents. An understanding of this structure is necessary if we are to get at the sentence's meaning.

Semantic Analysis

The string of phonemes that make up a word are sometimes enough to reveal the word's meaning. This is achieved by a comparison of the phonemic string with an internal database of sounds. If a match is obtained, the word's meaning is derived. But many times, there isn't a perfect match, and much ambiguity remains as to a word's meaning. In this case, the syntactical structure of the sentence can be useful.

In **compositional semantics**, the entire meaning of a sentence is derived from the meanings of its parts. A syntactical analysis identifies the type of word for each word in the sentence. This gives us information about how those words are related. For example, if a given word is identified as a noun, we know that it can be an agent or an instigator of an action. If a word is identified as a verb, we know that it represents an action. The structure of the sentence can then tell us whether that noun was the agent of that action. If a verb phrase is linked at the next higher level in the phrase structure to a noun phrase that immediately precedes it, there is a good chance that the action represented by the verb is brought into being by that noun. This structure, thus, has told us that the noun is an agent—additional information that has to do with the word's meaning. Similarly, it is clear that the adjective inside a noun phrase is a descriptor that applies to that noun and not any other. The structure has again helped us decipher meaning. We know that this word describes the noun in the phrase and not any other.

Consider the following sentence: Twas brillig, and the slithey toves did gyre and gimble in the wabe. What makes this sentence interesting is that it uses imaginary words yet follows proper grammatical rules. Even though most of these words are meaningless, we can still glean some understanding of what is happening by drawing on our knowledge of grammatical construction. You have probably guessed that *toves* is the noun and *slithey* is the adjective that describes it. You may also have intuited that *gyre* and *gimble* are verbs that describe the actions performed by the "toves" and that they are doing it in the "wabe." Now consider this sentence: Colorless green ideas sleep furiously. Here, we have an example of another sentence that follows grammatical rules. It differs from the one above in that it is composed of meaningful words. This time, however, the words are used in a nonsensical way. Ideas cannot be green, and if they could be, they could not also be colorless. Likewise, ideas can't sleep, and if they could, they couldn't do it furiously. This sentence is even more confusing than the one above because, despite the fact that the words fit the "slots" that make up a proper phrase structure, their meanings conflict with one another.

So we see that grammar does more than tell us what a correct or incorrect sentence is. Grammar specifies the abstract relationships between words in a sentence. These relationships are important clues in our deciphering of individual word and overall sentence meaning.

Pragmatic Analysis

Human language is a social construct. Individuals use it to communicate with one another. We communicate for all sorts of reasons. Sometimes, the purpose of a linguistic utterance is simple conveyance of information, as when we say, "This pillow is soft." But linguistic utterances can serve many other purposes—ones that require an action on the part of the listener, for example. It is sometimes not apparent what action should be undertaken by the listener, as many sentences invite action on the listener's part without directly issuing a command. **Pragmatics** is the social rules that underlie language use, as well as the strategies used by speakers to make themselves clear. Pragmatics helps us understand what actions we should take in response to spoken sentences.

Searle (1979) outlines five different types of spoken statements. Each type demands a different response from the listener.

1. **Assertives** are spoken statements in which the speaker asserts his or her belief. An example would be "It is hot in here." The statement suggests that we should open the window or turn on a fan or an air conditioner.
2. **Directives** are instructions dispatched from the speaker to the listener. They are direct commands and don't require an inference on the part of the listener with respect to what the desired action is. "Turn down the radio" is one such a command.
3. **Commissives** commit the speaker to a later action, as when a child says, "I will take out the garbage later." We would then need to verify that the garbage indeed had been taken out or not, and we would possibly impose a reward or punishment, depending on the outcome.
4. **Expressives** describe the psychological state of the speaker. "I apologize for yelling at you" indicates sorrow or regret and implies that the speaker probably continues to trust the person being spoken to.
5. **Declaratives** are spoken statements in which the utterance itself is the action. "You are fired" means that we have to look for a new job.

In each of these statements, we see that a speaker has used language to get the listener to perform an action. This is the case even in instances in which the sentence has not been phrased specifically as a command. Understanding the meaning of a statement is not enough here: One must infer what action the statement has asked for (directly or indirectly). Social context plays an important role in this process of establishing intent. "Do you know the time?" asked by someone waiting at the bus stop means that the speaker genuinely wants to know the time to find out if the bus is late. This same question asked of someone who has arrived late to an important meeting has a different intent. It is a criticism of that someone's lateness and not a request for the time.

Computer Language Programs and IBM's Watson

Many of you are probably familiar with or have used *Siri*, the language system built into the iPhone that has some limited capabilities to respond to spoken requests. A few more of you may also be familiar with a program called Dragon Naturally Speaking, which is capable of converting spoken language into text. So the age of "talking" computers is already with us. However, both these programs pale in comparison with the abilities of Watson, a computer system developed by IBM that has extensive world knowledge and is thus able to answer a wide variety of questions on different topics.

Watson was created with a specific goal in mind: to beat the human world champions at the television trivia show Jeapordy! In 2011, it did just that, beating Ken Jennings and Brad Rutter and winning a prize of 1 million dollars. How was Watson able to accomplish such a feat? To begin with, it was given access to a large amount of information, including all of Wikipedia, other encyclopedias, dictionaries, news information, books, and articles. When presented with a question, Watson would break it down into keywords and phrases and then search its database for these and related text strings to find an answer. It utilized thousands of algorithms to do these searches, independently and in parallel. The greater the number of algorithms that yielded the same answer, the greater the likelihood that a correct answer was reached. In a final step, Watson would do a secondary match against the database to determine whether the answer made sense.

In the game, human players were able to come up with responses faster than Watson, particularly when there was less contextual information available, but Watson was much faster at making an actual response to activate the buzzer: only eight milliseconds in comparison with several hundred milliseconds for the human participants. Watson is currently being used at the Memorial Sloan Kettering Cancer Center to perform management decisions for lung cancer treatment, where it is reported that most of the nurses using it follow its advice. Other fields for which it has promising applications are in government, telecommunications, and financial services. Although the cost of implementing a system like Watson now is prohibitive, it is probable that costs will come down in the near future and make it accessible to a larger portion of the public.

Evaluation of Natural Language Processing

Speech recognition systems still have problems dealing with noise, differences in pronunciation, and word ambiguity. Also, semantic analysis cannot rely on grammatical structure alone in the decipherment of meaning. As we have seen in the case of Watson, analysis must take into account real-world knowledge. This requires an extensive database filled with facts. Social context is also important. To "understand" what someone is saying, we need to know what their intent is. Are they just communicating facts? Asking a question? Issuing a command? Context expands beyond pragmatics though. To understand someone, it also helps to know who they are, who else may be present, where everyone is, what time of day it is, and a whole host of other situational factors. For instance, knowing that a conversation is taking place outdoors constrains the possible responses one can make (a light switch does not need to be turned on), while knowing who else is present also provides such constraints (it might not make sense to say something bad about Bill if his best friend is present). So if computer programs are to converse with humans in the most natural and sophisticated way possible, they need to have situational awareness, physical as well as social. Language cannot be detached from reality.

OVERALL EVALUATION OF THE LINGUISTIC APPROACH

Hopefully, this chapter has given you a sense of the importance and complexity of language. Linguistics resembles cognitive science as a whole in that it brings multiple perspectives to bear on a single topic. Instead of mind, its topic is language. Linguistics uses a varied set of theories and methods in its attempt to answer questions about language. The common thread that runs through its investigations is the subject matter, not any one technique or orientation.

We have made great strides in expanding our understanding of language in the past few decades. Grammatical analysis has been used to expand our comprehension of language structure. We have elucidated the linguistic abilities of animals, and we know the developmental stages that all children pass through as they acquire language. But there is still much to learn. The relationship between language and thought is still murky, and we await further research that will shed light on the ways in which language influences thinking. The Wernicke-Geschwind model in all likelihood will be revised in favor of a new, comprehensive model that will be more detailed and able to specify the neural basis of processes such as retrieval from the mental lexicon. In all likelihood, there will also be advances in computer-based language comprehension.

SUMMING UP: A REVIEW OF CHAPTER 9

1. Linguistics is the study of language and has been approached by many different perspectives and methodologies.
2. There is no single agreed-on definition of language, but we can say that it has five features. (1) It is used for communication, (2) the symbols used are arbitrary, (3) the ordering of the symbols is structured, (4) the symbols can be combined to form a large number of meanings, and (5) it is dynamic, meaning that it can change over time.
3. A phoneme is the smallest sound unit of a language that has no meaning. There are 45 phonemes in the English language. A morpheme is the smallest unit that does have meaning. Morphemes can be a word or part of a word.

4. The linguistic relativity hypothesis says that language and thought are powerfully intertwined, perhaps so much so that one could not express a thought in one language by using another language. There is, however, no broad support for the strong version of this hypothesis.
5. Animals are certainly capable of communicating, but there is controversy surrounding whether or not certain species possess language skills.
6. Researchers have attempted to teach animals language. The chimpanzee Washoe was capable of using limited ASL. A chimp named Sarah was instructed in the use of a token system. A third chimp by the name of Kanzi was estimated to use words and sentences at the level of a 2½-year-old child.
7. Human children pass through several distinct stages as they acquire language. They start by babbling a wide variety of sounds early in their first year. At around 6 months, they utter consonant–vowel pairs. One-word utterances occur at around 1 year of age. This is followed by two-word, sentence-like statements. There is evidence to support both domain-general and domain-specific theories of language acquisition.
8. Case studies of children deprived of language experience early in life demonstrate that there is a critical period of language acquisition. Genie was deprived of language until she was almost 14 years old. Despite intensive instruction, she was not able to learn adult-level language skills.
9. The rules that govern language use are known as grammar or syntax. Every sentence can be broken down into a hierarchical structure showing the relationship between different word types.
10. Chomsky identified a difference between a surface structure or actual organization of a given sentence and a deep structure containing the semantics or meaning. The surface structure can vary, while the deep structure is more constant. All the world's languages share some features in common. These are known as universal grammar.
11. The neural underpinnings of language ability have been studied for quite some time, especially when examining the language deficits that result from brain damage. In Broca's aphasia, patients have difficulty articulating or producing speech, but their comprehension abilities are intact. In Wernicke's aphasia, patients can produce rapid and fluent speech that is meaningless.
12. In the field of artificial intelligence, the goal has been to create computer programs capable of comprehending and producing speech. This would enable people to interact with machines more easily. Comprehension of natural languages in machines occurs in four stages: (1) speech recognition, (2) syntactic analysis, (3) semantic analysis, and (4) pragmatic analysis.

Explore More

Log on to the student study site at **http://study.sagepub.com/friedenberg3e** for electronic flash cards, review quizzes, and a list of web resources to aid you in further exploring the field of cognitive science.

Suggested Readings

Akmajian, A., Demers, R. A., Farmer, A. K., & Harnish, R. M. (2001). *Linguistics: An introduction to language and communication.* Cambridge: MIT Press.

Jurafsky, D., & Martin, J. H. (2008). *Speech and language processing* (2nd ed.). Upper Saddle River, NJ: Prentice Hall.

Lakoff, G. (1990). *Women, fire, and dangerous things: What categories reveal about the mind.* Chicago, IL: University of Chicago Press.

Rymer, R. (1993). *Genie: A scientific tragedy.* New York, NY: HarperCollins.

Savage-Rumbaugh, E. S., & Lewin, R. (1994). *Kanzi: The ape at the brink of the human mind.* New York, NY: Wiley.