Contents

Preface

Language is an inextricable part of almost every human life. Indeed, language is so interwoven with our humanity that just deciding which thread to pull first to unravel its mysteries is a daunting challenge. Should we study it as part of culture, as part of biology, as part of social psychology, as part of the cognition of individuals, as a complex pattern, or in some other way entirely? Writing this book has amplified my awareness of all the different perspectives on language and its development in human infants and children that I could have taken, but did not. Let me explain why I took the perspective that I did and therefore what you can expect to find if you choose to read beyond this preface.

The perspective most clearly represented in this book is that of cognitive science, the interdisciplinary study of the mind. Cognitive science is a mix of several traditional fields, including psychology, philosophy, computer science, neuroscience, artificial intelligence, and linguistics. I have chosen this perspective for two reasons. The first is that cognitive science is the discipline with which I am most closely affiliated. As an undergraduate, I took courses in psychology, linguistics, philosophy, and computer science. I have continued to pursue some mixture of these fields throughout my career, and currently am the Head of the Cognitive Science Interdisciplinary Program at the University of Arizona.

The second reason for employing a cognitive science perspective is that I wanted to write a book about language development that is theoretically coherent. Rather than following language learners from birth to some point at which we might say they have adultlike language abilities, I have organized this book around a set of contrasting theories. Throughout, I have tried to show how research in different areas of language is driven by researchers' attempts to tease apart these theories and their predictions. Therefore, this book has the potential to provide an understanding of competing accounts of language development, as well as a substantial sampling of the research that these accounts motivate.

Alas, good news seldom comes unaccompanied, and there is a cost to theoretical coherence. With some important exceptions, cognitive science as a field has tended to focus on individual minds and the computational processes used by these minds to make sense of their environments. Many of the enduring controversies about how young humans learn language have construed them in this way—as individual minds trying to extract and generalize the linguistic patterns of their community. This way of viewing language development is certainly missing some central properties of infants and children, such as the emotional factors that motivate them to learn in the first place, what they use language to communicate about and with whom, and how language develops at the same time as many other abilities. Therefore, in striving to paint a coherent picture, I have left out myriad studies and a few areas of research that on their own are priceless gems, but that I found difficult to place into the framework used to organize this book.

For the reader who is already familiar with the field of language development, I apologize if the studies or areas I have left out are some of your favorites. For the reader who is just learning about this field, you will find enough material to keep you occupied for some time, and you will learn about how theoretical debates drive research—certainly a good start. Whatever your reason for reading this book, I hope that you will be able to appreciate the mystery of how young humans come to exhibit a defining characteristic of our species—language.

LouAnn Gerken

1 Introduction

Why Study Language Development?

Our ability to communicate with each other through language, either spoken, written, or signed, is one that many of us seldom consider. However, if you have ever spent several days, or even several hours, in a foreign language environment, you know the feeling of relief and pleasure that comes when you are again able to use your own language. Language is so much a part of most of our lives that losing its use, even for a short time, can feel like the social equivalent of oxygen deprivation. Although we can also communicate without language, using eye contact, gesture, and so forth; these alternatives are crude tools when compared with linguistic forms of communication.

Nevertheless, the state of being unable to produce or comprehend the language of our social environment is the way we all begin our lives. Like adults in a foreign language environment, infants and young children can communicate with the world around them in nonlinguistic ways. However, language becomes the main tool of communication for most children within the first few years of life. What is the nature of human language that makes it the powerful and pervasive communicative tool that it is? How do human infants and children develop language so quickly and with such apparent ease? These two questions, and the research that they stimulate, compose the scientific core of the study of language development, which is the central focus of this book.

But studying language development also has a number of practical, or applied, benefits. For example, many people might want to know the answers to the following questions:

1. Is it good or bad for a child's linguistic development to be raised as a bilingual?

2. What is the best way to teach reading skills?
3. What are the likely outcomes of hearing impairment on language development?
4. How can we build a computer that comprehends and produces language as well as a 3-year-old child?

One group of people who might want answers to such questions is current and future parents. For a parent, the first indication that his or her infant understands a word or phrase is a joy and a relief because it signals the end of a period in which the infant's needs and desires must be guessed from nonverbal cues. Even more momentous is the child's first word. In many cultures, the first word signals the transition from infancy into childhood. For example, some cultures view infants as unintelligent beings until first word production. Indeed, our own English word *infant* comes from a Latin phrase meaning *unable to speak*. For parents, knowing what to expect about the course of language development can make observing the process more interesting. It may be especially interesting to know that, even before your child shows signs of understanding or producing language, he or she is making enormous progress toward joining you as a fellow language user. Knowing what to expect may also allow for better decisions about what to do when language development seems disrupted in some way.

Studying language development can also benefit people who are preparing for a career involving children, including those in daycare centers, schools, hospitals, or clinics for children with speech, language, or hearing disorders. Understanding the normal course of development is useful in these settings in order to identify children who may not follow the typical path. Even more importantly perhaps, people working with children in professional settings should understand the many components of a language system that must develop, and stay updated on the variety of tools being created to assess the separate components of language development. People working with children in professional settings also stand to benefit as our understanding of developing language abilities, such as word comprehension or sentence production, is increasingly paralleled by our understanding of brain development. Some careers in organizations that serve infants and young children also entail making policy decisions, and people who make such decisions should be aware of the likely consequences. For example, the policy that all newborns must

be screened for hearing status that many states have implemented is a direct result of policy makers' understanding the role of intact hearing at the very earliest stages of spoken language development.

Careers in the information technology fields can also benefit from the study of human development in general and language development in particular. For us to create computers that are maximally interactive with humans, understanding how our impressive human abilities developed is invaluable. Conversely, different hypotheses about how language learning works can be implemented in different computer programs in order to determine how well the computer performs in comparison to real human learners. The **computational models** of learning that perform best are taken to be closer approximations to the way infants and children learn. These models of language learning can, in turn, help language development researchers design experiments to ask more fine-grained questions about the language learning process.

In addition to introducing you to the scientific study of language development, one of the goals of this book is to demonstrate some of the currently available tools that might be used to answer practical questions like 1 through 4, above, and to provide basic background to people who interact regularly with infants and children.

The Nature of Language

Many organisms, in addition to humans, communicate with each other, but human language has two properties that, as far as researchers can tell, no other communication system has. First, language has the ability to combine **meaning units** such as words, and different combinations of the same words result in different meanings. For example, *Venetian blind* means something entirely different from *blind Venetian*. Second, the meaning units (e.g., words) of language are not atomic or indivisible units, but are themselves composed from a limited inventory of smaller parts, which we can call **submeaning units**. For example, the words *apt*, *pat*, and *tap* are all composed of the same submeaning units. Similarly, the signs used in American Sign Language (ASL) are composed of smaller components, including the location in which the sign is made and the shape of the hand. (For a more in-depth overview of the material described in this section see Akmajian,

Demers, Farmer, & Harnish, 2001). Although some other animals produce meaning units or meaning unit combinations, and others combine submeaning units, no animal other than humans, as far as we currently know, combines both meaning and submeaning units as part of their communication system. To better understand what the combinatorial ability of language can do, let's consider three systems used by nonhuman animals.

Meaning Units

Vervet monkeys make three danger calls: one for snakes, one for eagles, and one for leopards (Cheney & Seyfarth, 1990). Each call has a particular meaning or referent, and therefore is something like a word in human language. However, there is no way within the call system of the vervets to indicate that the snake that has just appeared is particularly large, or that there is an eagle in the distance that is not a strong threat at the moment. For example, the first message might be conveyed by making the snake-call twice. However, there is no evidence that vervets engage in such call combinations. In contrast, human language allows us to combine meaningful units to produce an utterance with a new meaning. For example, we can say *big snake* and *nonthreatening eagle*.

Combinable Meaning Units

Although there are a few other species that appear to be able to combine calls, the way in which they do so importantly is different from human language. One such species is a New World marmoset monkey called a cotton-top tamarin (Cleveland & Snowdon, 1982). Cotton-tops can make a chirp to indicate alarm and a squeak to indicate alertness. They can combine these calls, always with the chirp preceding the squeak, to indicate vigilance in the presence of danger. At first glance, this system seems very similar to the human ability to combine words. The signals are combined in a particular order, and the combination yields a meaning distinct from the meanings of either signal alone.

However, there appear to be two differences between the cotton-top's call combinations and the combinations of meaning units

used by human language. First, different orders of the same two calls do not appear to have different meanings in the way that *Venetian blind* and *blind Venetian* do. Therefore, although call combinations allow cotton-tops to express a larger number of meanings than they could express by single calls alone, this species does not appear to take advantage of the full communicative power that call combinations can yield. The fact that cotton-tops do not use the full power of a call combination system is probably related to the second difference between their call combinations and those used by human language. Cotton-tops' call combinatons are like their single calls in that they refer to the internal alertness state of the animal making the call. Similarly, the vervet monkey's danger calls refer to a predator present in the environment. Contrast this **referential** use of signals with the English sentence *the eagle is eating the leopard*. This sentence describes a relation between a **predicate** (*eat*) and two **arguments** (*eagle* and *leopard*). The difference between referential and predicative utterances is most obvious when the meaning to be communicated is not about the moment in time at which the signal is produced. For example, it is possible in a predicative system, but not in a referential system, to convey meanings like *I saw an eagle yesterday* or *watch out for eagles*. The ability to use combinations of meaningful units to express predicative relations is a crucial property of human language and appears to have no counterpart in animal communication systems.

Combinable Submeaning Units

A property of human language that appears to have no counterpart in any other communication system is that meaningful units (e.g., words) are made up of combinations of submeaning units that are themselves meaningless (recall *apt, pat,* and *tap*). This ability to combine submeaning units into meaningful units gives human language the power to create a very large number of words. (The number of words could be infinite if there weren't a limit on the length of words we could perceive and produce.) Some species do combine meaningless units into larger sequences. For example, male gibbons (a type of ape) create elaborate songs from a set of acoustically distinct elements, or notes (Mitani & Marler, 1989). The notes are combined according to rules that limit the number of possible combinations. The fact that some combinations are "illegal" in the gibbon communication system

might be seen as similar to human language. For example, recombining the sounds in *apt* to make *pta* is not allowed in English. However, unlike humans combining sounds to make words, one gibbon song with a particular combination of notes does not appear to have a distinct meaning from another gibbon song with a different combination of notes. Therefore, although other species do combine sounds in nonrandom ways, humans appear to be the only species that combines submeaning units into meaning units.

The Combinatorial System Used in Human Language

To fully understand the linguistic system that children develop, we need to consider the combinatorial system of language in more detail. As we have discussed in comparing human language to the communication systems of some other species, language is organized **hierarchically**, with larger units composed of smaller ones. The components of the hierarchy are called **linguistic levels**. In this book, we consider the data on how children master different levels of the hierarchy, with the largest unit under consideration being the sentence. Although adult humans communicate in connected sentences, or **discourses**, which have a structure of their own, we consider children's mastery of discourse only as it relates to their mastery of other levels. The levels in the hierarchy that we discuss in depth are described below.

Phonology (Submeaning Units)

Beginning with the smallest units in the language hierarchy, we have already noted that the meaningful units of language are made up of submeaning units. The submeaning units that are used by a particular language are called the **phonemes** of the language. The inventory of phonemes used in each language and the ways that they can be combined are topics studied in the field of phonology. We discuss how infants and children acquire the **phonology** of their native language in Chapters 2 through 4.

Lexical Semantics (Meaning Units)

Using the principles of phonology, phonemes are combined to make meaningful units, the most familiar of which are **words**. The words a person knows are thought to be stored in a mental **lexicon**, and the study of word meanings is part of the field of **lexical semantics**. We discuss how children acquire the meanings of words and build a mental lexicon in Chapter 5.

Morphology and Syntax

Meaningful units of language can be combined in two ways. One type of combination entails forming new words by combining meaningful units called **morphemes**. A morpheme can be a word itself, like *dog*, or a part of a word like the plural markers on *dogs*. The ways in which morphemes can be combined to make words is called **morphology**. The second way that meaningful units can be combined entails putting words together to make phrases (e.g., *Venetian blind*) and sentences (e.g., *I saw an eagle yesterday*) using principles of **syntax**. We discuss the acquisition of morphology and syntax in Chapters 6 through 8.

Grammar

The entire linguistic system comprising phonology, semantics, morphology, and syntax is called a **grammar**. The notion of grammar that is used in this book refers to a concise description of the patterns of a language that can be found at each level in the linguistic hierarchy. In the field of linguistics, this concise description is called a **descriptive grammar**. This notion of grammar is different from a **prescriptive grammar**, which you may have learned in grade school, in which rules specify the standard ways that language is used among educated people. For example, you may have learned a rule of a prescriptive grammar that tells you not to end a phrase with a preposition, which makes the following sentence **ungrammatical**: *This is the movie I want to go to*. According to such a grammar, the sentence *This is the movie to which I want to go* is **grammatical**. However, people, in

fact, use sentences like the former, especially in casual speech; therefore, the sentence would be grammatical in a descriptive grammar of English. The goal of the study of language development is to determine how learners come to comprehend and produce sentences that are part of the descriptive grammar of their community.

Theories of Language Development

Children enter the world with no obvious linguistic ability. Within less than a year, they can recognize a few words, and a few months later produce words. After only two years, they are able to understand much of what is said to them and produce some simple word combinations of their own. All of this says that there are clear and rapid developmental changes in infants' and children's language abilities. How can we explain those changes? This section outlines four theoretical approaches to language development. Before we consider the four approaches, however, let's consider what we are trying to explain and why theories are important.

What Are We Trying to Explain?

For children to be able to speak or sign like the people in their community, they need to be able to perceive the physical signal that constitutes the utterances produced by others and to move their own articulators (mouth or hands) to make similar signals. Although these are daunting tasks, most theories of language development do not take accomplishing these tasks as their focus. Rather, what is taken to be the central puzzle in language development is how children **generalize** from the utterances that they encounter (the **input**) to new utterances. In other words, it is the ability to use language to convey any thought that might come into the child's mind, even if she has never heard that thought expressed before, that is, the focus of theorizing.

To understand this focus, consider a parent who points to a bear and says, *That's a bear*. Imagine that the child is able to produce some credible imitation of the sentence. Shortly afterward, the parent looks at a carton of orange juice and says *I want some juice*. Imagine that the child is able to produce a version of that utterance as well.

How far along in the process of language development is the child? To ask that question, let's further imagine that the parent never says in the child's hearing *I want a bear*. Now we can ask whether the child, who wants a toy bear, is able to utter *I want a bear*. If the child is able to produce this novel utterance that reflects a novel intended meaning, she is well on her way to becoming a full participant in her linguistic community. For you to find the remainder of this chapter, or for that matter, the remainder of this book, of any interest, you have to believe that children say things that they have never heard before. Although generalizing beyond her input isn't the only task for the child, it is at the heart of the mystery of language development.

Do humans generalize in other domains besides language? Do nonhumans also generalize beyond their experiences? The answer to both questions is *yes*. Let's consider an example of nonlinguistic generalization that we can use as a basis for considering what develops in language development. In the three examples shown in Figures 1–1A through E, Figures 1–2A through E, and Figures 1–3A through E, you will see three bars in a black rectangle. The A through D examples are your input. You need to decide whether the E example is a valid generalization from that input.

The answers to the questions posed in Figures 1–1E, 1–2E, and 1–3E are *no, yes, no*, respectively. In Figures 1–1A through D, the principle used to create the examples was that the three bars together needed to cover more than 50% of the area of the rectangle. This generalization is not particularly natural for humans, although pigeons can learn it over many trials. In Figures 1–2A through D, the principle used to create the examples was that the three bars had to differ in height. This generalization is relatively natural for humans, and even if you didn't get the correct answer from just four input examples, it is likely that you would have made the intended generalization with a few more examples. Pigeons, on the other hand, don't seem to be able to learn this generalization. Finally, the principle used to create Figures 1–3A through D was again that the three bars had to differ in height. However, many people are likely to have made a more narrow generalization—that the three bars had to decrease in height from left to right. Example 1–3E violated that generalization. Note that if you had seen, as part of the input examples, just one example of bars of different heights, but not decreasing in height from left to right, you would have abandoned the generalization about bars decreasing in height from the left.

later (Emmorey & Carina, 1990; Mayberry & Fischer, 1989; Newport, 1990). In summary, the data on age of exposure to first language are remarkably consistent with data from the more typical second language learning situation. The data suggest that there are effects of age and/or of nonlinguistic experience that prevent older learners from achieving the same outcomes as younger learners. We will see similar age effects when we look at people who must create language in order to communicate in the absence of a pre-existing target language.

Creating Language Structure

The studies presented above examined people who were not given a language model until after the normal age of exposure. However, these people were generally linguistically isolated, either because they were indeed alone, or because the people around them did not want to or could not communicate with them. In this section we contrast that situation with one in which a person or people without a full language model either augment a rudimentary input language or create a new language in order to communicate. One such study involves deaf children who are learning ASL from hearing parents who themselves learned sign language in adulthood to communicate with their deaf infant. These parents have been shown to make very inconsistent use of the morphosyntactic markers of ASL. Their children, however, regularize the inconsistent forms, thereby looking more like native signers than their parents (Ross & Newport, 1996; Singleton & Newport, 2004).

A similar situation can be seen in the fascinating emergence of Nicaraguan Sign Language (NSL; Kegl, 2002). NSL was created in the 1970s and 1980s when two schools for the deaf were begun in Nicaragua. Prior to the creation of the schools, deaf children stayed in their home villages, communicating as best they could with the hearing community. When the schools were begun, formal instruction emphasized spoken Spanish, lip-reading, and finger-spelling. However, in order to communicate with each other, the students developed their own sign system that became increasingly standardized across signers. Once the schools had been in operation for a sufficient number of years, new students did not need to create a language, but could learn the system that the previous students (the first cohort)

had created. One interesting line of research has followed the change in NSL as the second cohort of students has begun using the language (Senghas, 2003; Senghas & Coppola, 2001). One study focused on a morpheme that modifies a basic verb (e.g., *pay*) to indicate the object of the verb. For example, a referent *man* might have been indicated as an arbitrary location in space (e.g., to the left of the signer), and to indicate that someone was paying the man, the verb *pay* would be made in the direction of the previously established referent (note that such a system is also used in other independently created signed languages). Members of the first cohort showed little use of this spatial verb marker. Adult members of the second cohort also showed little use of the marker. However, children of the second cohort substantially increased the use of the marker, but only in a set of contexts that were linguistically appropriate. Thus, like children learning ASL from an inconsistent model, children in the second cohort made the language more regular.

Another approach to understanding age effects on learning from inconsistent input can be seen in studies comparing adults' and children's ability to learn an artificial language in the laboratory. Hudson and Newport (1999; Hudson Kam & Newport, 2005) explored the conditions under which language learners are most likely to make changes to their input. They exposed adults and children to an artificial language system in which nonsense nouns and verbs referred to objects and actions in an artificial world presented on a video display. The language contained four types of sentences: intransitive, transitive, negative intransitive, and negative transitive. The sentences were presented with an optional negative marker followed by verb, subject, and an optional object. There were two determiners, and nouns were randomly assigned to occur with just one of them. Participants were assigned to one of four conditions, with consistency of determiner use in the input varying across conditions. Determiners were used either 45%, 60%, 75%, or 100% of the time. In a sentence completion task, adults matched their production of determiners to the condition that they were in. For example, adults in the 45% determiner use condition-produced determiners about 45% of the time, whereas adults in the 60% determiner use condition-produced determiners about 60% of the time. In contrast, 5- to 7-year-old children, when faced with sentences in which a determiner was presented 60% of the time, regularized determiner use, producing determiners on almost all utterances.

These results, like the results from the NSL study, suggest that children are more likely than adults to create a language system that is more regular than their input.

The studies described above all have in common the fact that learners take elements from their input that are not used systematically and make them systematic in ways consistent with existing languages. Other researchers have examined the situation when a deaf child has no language input at all and creates a gestural communication system called "home sign" (e.g., Feldman, Goldin-Meadow, & Gleitman, 1978; Goldin-Meadow & Mylander, 1998). Goldin-Meadow and Mylander (1998) report that four American and four Chinese deaf children, who were never exposed to a signed language, each showed a set of linguistic properties in their utterances that the authors convincingly argue cannot easily be attributed to the gestural input of their parents. First, they produced sequences expressing sentencelike content, as well as gestures for single words. Second, they produced multiclause sentences. Third, all eight children showed a pattern of production that is consistent with a particular grammatical system (ergative) that is found in human languages, such as Inuit. At the very least, the studies of home sign suggest that human children are biased to express themselves in predicates and arguments that have a regular form. That is, unlike the various nonhuman primates described in Chapter 1, children with no language input do not appear to be inclined to produce utterances that simply refer.

Summary

We began Chapter 9 with the hope that studies of language development in atypical populations or situations might augment the studies discussed in Chapters 2 through 8, with the result that we would have a better understanding of the relative contributions of biology, linguistic input, and their interaction. Although the studies presented in this chapter raise at least as many questions as they answer, they combine with the information that we already have to allow us to make several tentative conclusions about the nature of language development.

First, infants and children, regardless of whether they are genetically intact and learning language in a typical situation or not, are

keen observers and impressive generalizers. We began the book with the assumption that generalization is the central question in studies of language development, and study after study demonstrated that infants and children notice and generalize over a vast variety of patterns in their environment. Studies of artificial language learning demonstrate that they even generalize patterns that do not occur in real human languages. Such studies suggest that our human ability to generalize from input may initially extend beyond patterns that we actually put to use in the service of communication. A related point is that the types of generalizations that we make most readily change as we gain experience with our world. We saw in Chapter 3 that older learners fail to make some phonological generalizations that younger learners make. And we saw in Chapter 5 that the shape bias in word learning seems to emerge through experience with many shape-based categories. Thus, whatever propensities to generalize we are born with may quickly be replaced with new ones that are more in keeping with the particular environment we inhabit.

Second, infants and children overgeneralize. The famous example of children producing *goed* instead of *went* has been the subject of at least 50 years of theorizing. Grammarless computational models of overgeneralization are important because they demonstrate that the young child's almost obsessive cleanup of messy input may have its origins in a fairly simple mechanism. Whether such models ultimately can account for the types of overgeneralization we see in real learners has yet to be determined. A number of studies presented in Chapter 9 demonstrate that, although children have a propensity to overgeneralize, adults do not. This observation may in fact lie at the heart of age-related differences in the ability to learn a language. As seen in the studies of Nicaraguan Sign Language, children's propensity to overgeneralize is largely responsible for language change. One question we must now ask is whether children's overgeneralization stems from general cognitive abilities (e.g., the inability to remember the specifics of the input as well as adults can) or whether they show a particular propensity to overgeneralize *languagelike* input.

Third, humans seem to want their communication structured in certain ways. The studies of deaf children's creation of home sign clearly show that what we want to communicate to our fellow humans is predicative, not merely referential. However, the fascinating data on home sign must be seen in the context of the data on Nicaraguan Sign Language. The relation between a human predisposition for predica-

tive communication and the stable communicative system (language) that is the ultimate result is expressed nicely by Ann Senghas (2003), writing about the second cohort of NSL users: " . . . this new version of the language is not unrelated to its model; it is derived from it. Forms that exist in free variation or with some other function in the language of the first cohort were available for the second cohort to use as raw materials for creating new form-function mappings. If the first stage were not necessary, all of NSL would have appeared in a single sweep, instead of being built cohort by cohort." Thus, one way of viewing these studies is that humans are biologically disposed to communicate certain aspects our mental lives: predicates and arguments, or thought of another way, events. The predicate-argument or event structure of our intended communication puts some constraints on the forms that language can take. But those constraints alone ultimately do not determine the grammars that each language settles on. Rather, child language learners, largely through a process of overgeneralization, reshape the input.

Alas, the notion that children reshape their input brings us back to our original question about the nature (and nurture) of language learners. Do the structures that children impose on messy input reflect an innate grammar that is elicited by certain types of environmental stimuli? Or are the linguistic structures that arise from the learning process simply good ways to communicate efficiently? The observation that human languages use a relatively limited set of forms to communicate might be taken as evidence for the view that input elicits pre-existing grammars. However, the data from Specific Language Impairment and Williams syndrome make clear that the entity we call 'language' does not reflect a single ability. Phonology, morphosyntax, and lexical knowledge are differentially affected by different genetic anomalies.

Furthermore, even within these larger linguistic domains, there is considerable variation: Verbal morphology may be affected in SLI differently than some other aspects of morphosyntax. Receptive vocabulary in WS is affected very differently from the ability to produce a label for a picture. In short, when language "breaks," it fractures along lines not easily predicted by either a general purpose learning mechanism or any currently proposed innate grammar. Those of us who care about the biological underpinnings of language and what they reveal about what it means to be human clearly have many years of work ahead of us. This book is an invitation to join us.