AN INTRODUCTION

TO LINGUISTICS

III

THIRD YEAR/BASIC EDUCATION

FACULTY OF EDUCATION

SECOND TERM

2022-2023

DR. REWAA ANWER

Introduction

RALPH FASOLD AND JEFF CONNOR-LINTON

"History is universal and basic," a history professor said during a faculty meeting, "It's about every event that involves all people at all times and in all places." "Yes," observed his colleague from linguistics, "but how would you record and interpret that history without language?" Indeed, it is hard to imagine how there could even be history without language, without a means to pass a record of what has happened from one generation to the next through retold stories and sagas, even before written records. Much of the history (and prehistory) of the human species consists of the development and adaptation of various tools to meet a broad range of needs: think of the wheel, the domestication of animals, the steam engine, computers and the internet. The development and refinement of these and all other tools could not have been accomplished without language.

The human capacity for self-awareness and abstract thought is facilitated by language, if not dependent upon it. The ability to transfer complex information, to discuss the meaning of events and possible outcomes of alternative actions, to share feelings and ideas – all these are impossible without language. The origins of language are shrouded in obscurity, but archaeological records suggest that communication with language emerged about 200,000 years ago. The ability for an individual to model the world for him/herself and to communicate using language was probably the single most advantageous evolutionary adaptation of the human species.

Universal properties of language

Over thousands of years of evolution, the human species developed a vocal tract flexible enough to make a wide range of distinguishable sounds and the ability to perceive differences among those sounds. But most important, the human species developed the ability to use these sounds in systems which could communicate meaning. No one knows just how this happened. Perhaps mental capacities that had evolved for a variety of other adaptive purposes (like fine motor hand–eye coordination) were "re-purposed" to support a complex symbolic and communicative system. Perhaps some mental capacities are exclusively dedicated to language and evolved more gradually along with the increasing complexity of human communication. Or perhaps once they reached a certain level of neurological

and cognitive complexity, the synapses of the brain "reorganized" themselves, making the development of language possible. In any case, language is a distinctive attribute of the human species.

Although languages differ in many ways, they are all made possible by the same genetic information, they are all processed by the brain in basically the same ways, and, not surprisingly, they all share certain fundamental "design features" and structural characteristics that enable them to work the way they do. For example, although different languages use different sets of sounds, their sounds are organized and combined according to just a few principles. If there were no shared, universal features of language, we would expect the sounds of languages and their combinations to vary randomly. Instead, the sounds of languages and their combinations are limited and systematic. Likewise, all languages follow similar constraints on how they can combine words into phrases and sentences.

Understanding and explaining the properties which are universal to all languages – as well as those which vary across languages – is the fundamental job of the linguist.

Modularity

Most linguists believe that language is a modular system. That is, people produce and interpret language using a set of component subsystems (or modules) in a coordinated way. Each module is responsible for a part of the total job; it takes the output of other modules as its input and distributes its own output to those other modules. Neurolinguistic studies show that different regions of the brain are associated with different aspects of language processing and, as the following chapters show, dividing language into modules facilitates linguistic analyses greatly.

Some modules have been central to linguistics for a long time. Phonetics is about production and interpretation of speech sounds. Phonology studies the organization of raw phonetics in language in general and in individual languages in particular. Larger linguistic units are the domain of morphology, the study of structure within words – and of syntax, the study of the structure of sentences. Interacting with these modules is the lexicon, the repository of linguistic elements with their meanings and structural properties. In recent decades, philosophers have developed the formal study of semantics (the detailed analysis of literal meaning), and linguistics has incorporated and added to semantics as another module of language. Still more recently, discourse – organization of language above and beyond the sentence – has been recognized by most linguists as another important subsystem of language.

Constituency and recursion

All languages are organized into constituents, allowing more complex units to enter structures where simpler ones are also possible. So we can say in English, "*She* sat down," "*The smart woman* sat down," "*The tall, dark-haired, smart woman with the bright red sweater and pearl necklace* sat down." Being composed of constituents gives language a balance of structure and flexibility. Constituents can be replaced by other constituents, but you

can't replace a constituent with a series of words that is not a constituent. So you can't replace *she* with *smart with the bright red sweater* ("*Smart with the bright red sweater* sat down" doesn't work). Constituents can be moved, but you can only move a complete constituent. *She is very smart* is possible and so is *Very smart, she is*, but not *Smart, she is very*.

Being composed of constituents also allows languages to be recursive. Recursion is the property of language which allows grammatical processes to be applied repeatedly, combining constituents to produce an infinite variety of sentences of indefinite length. Recursion is what allows someone to expand a short sentence like *He was tall* into longer sentences like *He was tall and strong and handsome and thoughtful and a good listener and . . .* or to embed clauses, as in *This is the mouse that nibbled the cheese that lay in the house that Jack built.* The recursiveness of language has profound implications. It means that no one can learn a language by memorizing all the sentences of that language, so there must be some other explanation for how human beings are able to learn them. The human brain is finite, but recursiveness means that it is capable of producing and understanding an infinite number of sentences.

Discreteness

Another property of all languages is **discreteness**. The range of sounds that human beings can make is continuous, like a slide whistle. For example, you can slide from a high "long *e*" sound (as in *feed*) all the way down to a low "short *a*" sound (as in *bat*) and then slide back to a "long o" sound (as in *poke*) – all in one continuous glide. But all languages divide that continuous space of sound into discrete, incremental territories, just as most western music divides the continuous range of pitch into discrete steps in a scale. Sounds that are discrete in one language may not be discrete in another. In English, for example, we distinguish "short *a*" from "short *e,*" so that *pat* and *pet* are different words. The same is not true in German, so German speakers have trouble hearing any difference between *pet* and *pat*. At the same time, German has a vowel that is like the English "long *a,*" but with rounded lips, spelled *ö* and called *"o*-umlaut." The distinction between the vowel that is like English "long *a*" and this rounded vowel is responsible for the meaning difference between *Sehne* ('tendon') and *Söhne* ('sons'). This distinction is as easy for German speakers as the *pet* and *pat* distinction is for English speakers, but it is hard for English speakers. Precisely *what* is discrete varies from one language to another, but all languages have the property of discreteness.

Discreteness also shows itself in other modules of language, for example, meaning. The color spectrum is a clear example. Color variation is a continuum – red shades through red-orange to orange to yellow-orange to yellow and so on through the spectrum. But all languages divide the color spectrum into discrete categories, although languages differ in how they divide those continua into words. In some languages there are only two basic color terms, roughly meaning 'light' and 'dark;' others add red, yellow, and green, whereas still others, including English, have developed words for many more colors. Likewise, although the claim that Eskimos

have hundreds of terms for snow is overstated, the languages of Native Americans living in the far north do distinguish more kinds of snow than do languages which have developed to meet the needs of peoples living in warmer climates. Similarly, American English has a range of words for different types of automotive vehicles (sedan, sports utility vehicle, minivan, convertible, wagon, sports car, for example) related to the importance of the automobile in that culture.

Language is composed of separate sounds, words, sentences and other utterance units. The fact that we hear speech as a sequence of individual sounds, words, and sentences is actually an incredible accomplishment (and all the more incredible for how instantaneously and unconsciously we do it). Acoustically sounds and words blend into each other. (If you have tried to learn a second language as an adult, you know how hard it can be to separate words spoken at a normal conversational pace.) Remarkably, babies only a few weeks old are able to distinguish even closely related sounds in the language of their home from each other and to distinguish the sounds that belong to the language they are learning from the sounds in other languages at a very early age. Furthermore, children in the first year or two of life learn to pick out words from the stream of speech with no instruction.

Productivity

Another key feature of language is **productivity**. When people hear a word for the first time, they often ask, "Is that a word?" If they ask a linguist, the answer is likely to be, "It is now." If the novel word is formed according to the morphological and phonological rules of its language and it is understandable in context, it is a bona fide word, even if it's not found in a dictionary. Languages can systematically combine the minimal units of meaning, called morphemes, into novel words, whose meaning is nonetheless deducible from the interaction of its morphemic components. Imagine each speaker in the world coining just one new word, and you'll have some idea of just how productive a language can be. Most of these spontaneous coinings – inspired by a particular context – are not used frequently enough to ever make it into a dictionary, but some coinings *do* become part of the lexicon because they meet a new need. Productivity is one way in which languages change to meet the changing communicative needs of their speakers.

The productivity of language comes from more than just the ability of speakers to coin new words. Sentences can become indefinitely long, by adding modifiers (*A great big huge beautifully designed, skillfully constructed, well-located new building . . .*) or by including one sentence in another, over and over again (*He said that she said that I said that they believe that you told us that . . . .*). Since languages place no limits on the use of these recursive processes, all languages are potentially *infinitely* productive.

Arbitrariness

The productivity of languages derives, in large part, from the fact that they are organized around a finite set of principles which systematically

constrain the ways in which sounds, morphemes, words, phrases, and sentences may be combined. A native speaker of a language unconsciously "knows" these principles and can use them to produce and interpret an infinite variety of utterances. Another fundamental property of language is its **arbitrariness**. With few exceptions, words have no principled or systematic connection with what they mean. In English, the first three numbers are *one*, *two*, *three* – but in Chinese they are *yi*, *er*, *san*. Neither language has the "right" word for the numerals or for anything else, because there is no such thing (Bolton, 1982: 5). Even onomatopoetic words for sounds, like *ding-dong* and *click,* that are supposed to sound like the noise they name, actually vary from language to language. The linguist Catherine Ball has created a website listing the sounds various animals make in different languages around the world (http://www.georgetown.edu/faculty/ballc/animals/dog.html). In English, for example, a dog says *bow wow* or maybe *woof woof*, but in Hindi it says *bho: bho:*. Greek dogs say *gav* and Korean dogs say *mung mung*. People perceive these sounds through the arbitrary "sound filters" of their respective languages, so even something as seemingly objective as a dog's bark is in fact represented arbitrarily in language.

Even the speech sounds of a language are arbitrary. English is spoken using only 36 different sounds (a few more or less, depending on how the English sound system is analyzed). But, as you will learn in detail in Chapter 1, the sounds used in English are not all the same as the sounds needed to speak other languages, nor are they put together in the same way. These 36 sounds are in turn arbitrarily represented by 26 letters, some of which stand for two or more sounds (like *g* in *gin* and in *gimp*) while other sounds are spelled in two or more different ways (consider *c* in *center* and *s* in *sender* or *c* in *cup*, *k* in *kelp,* and *qu* in *quiche*). The patterns into which words and sounds are arranged are also arbitrary. We know perfectly well what *tax* means but any English speaker knows without a doubt that there is no such word as *xat*. Adjectives go before nouns in English – so it's *fat man*; in French nouns go before adjectives, making it *homme gros*. Arbitrariness is a property of sign languages as well as spoken language. Some visual signs in sign languages are iconic – they look like what they mean – but most signs give not the slightest clue to their meaning.

It's important to remember that arbitrariness doesn't mean randomness. It means that, for example, the sounds that one language uses and the principles by which they are combined are inherently no better or worse than those of any another language. Likewise, it means that the principles of one language for arranging words are inherently no better or worse than those of another.

Reliance on context

A corollary of arbitrariness – of association between sound sequences and meanings or in the order of words in phrases – is **duality**. Because there is nothing about the pronunciation of the word *one* (transcribed phonetically – as it sounds – it would be [wʌn]) that necessarily associates it with the

numeral 1, that same sequence of sounds (but spelled *won*) can also be used to mean something entirely different, the past tense of the verb *win* (Bolton, 1982: 5). But if the same sequence of sounds can represent different concepts in the same language, how are you able to figure out which meaning I intend when I say [wʌn]? The answer – which is as complex as it is obvious – is that you rely on its context. If I say [wʌn] *before* a noun, as in "[wʌn] dog," your knowledge of English grammar will lead you to guess that I mean *one*. On the other hand, if I say [wʌn] *after* a noun (or pronoun), as in "Mary [wʌn]," that same knowledge will lead you to guess that I mean the past tense of *win.*

Reliance on context is a crucial property of languages, not just in figuring out the meaning of words like *one* and *won*, but in interpreting the meaning of entire utterances. The meaning of a sentence depends crucially on the context in which it is uttered. That context could be the sentence or sentences that immediately precede it, or it could be the broader physical or social circumstances in which the sentence it uttered. If someone says "One," the meaning of that utterance is only clear in the context of a preceding utterance – for example, "Do you want one lump of sugar or two?" Similarly, "It's cold in here" could be a complaint, a request to close a window, or even a compliment (about a freezer, perhaps). Who or what a given pronoun (like *she*, *it*, *us*, or *them*) refers to may rely on prior sentences or the immediate physical environment. Languages rely on the connection between form (what is said) and context (when, where, by whom, and to whom it is said) to communicate much more than is contained in a sequence of words.

Variability

Although all languages share some universal characteristics, languages also differ in many ways. The language that people use varies depending on who's speaking and the situation in which they're speaking. In fact, variability is one of the most important – and admirable – properties of language. Variation (also known as difference and diversity) is the essence of information. Without variation in light frequencies there would be no sight; without variation in sound frequencies, there would be no language and no music. (And as we are beginning to realize, without a certain minimum level of genetic diversity, our ecosystem is threatened.) Variability in language allows people to communicate far more than the semantic content of the words and sentences they utter. The variability of language is indexical. Speakers vary the language they use to signal their social identities (geographical, social status, ethnicity, and even gender), and also to define the immediate speech situation.

People let the world know who they are by the variety of their language that they use. They reveal their geographical and social status origins after saying just a few words. People also use their variety of language to signal membership in a range of overlapping social groups – as male or female, as a teenager or an adult, as a member of a particular ethnic group. They keep their speech, often despite the best efforts of teachers

to change it, because at an unconscious level, maintaining their ties to their origin is more important than any reason to change.

People also use language variation to communicate the situation and purpose in which they are talking, as well as the roles they are playing in those situations. A priest uses different forms of language during a sermon than during the social hour after a church service, playing different roles (and projecting different roles on the churchgoers he addresses). At work, people speak differently to subordinates than to superiors, and differently during coffee breaks than in meetings. Parents speak differently to their children than to other adults (or even to other people's children). The language used in writing typically differs from the language used in speaking, reflecting and communicating the different conditions under which language is produced and its various purposes.

A large part of a speech community's culture is transacted through the medium of language variation. Norms of appropriate language use help speakers to construct and negotiate their relations to each other. The unwritten and unconsciously applied rules for the various forms and uses of language can vary from one cultural milieu to another, within and between societies, and even between genders. This raises the risk of misunderstanding when speakers unknowingly are behaving according to different cultural norms, but enriches our ways of seeing the world when those differences are understood.

The descriptive approach

The fact that language is a universal characteristic of human beings means that all languages (and language varieties) are equal. That is, they all come from the same genetic blueprint, and they all are equally "human." Language varieties differ because over time they have adapted to the differing needs of their speech communities. Each language does things differently: some languages morphologically mark several verb tenses (English marks only two); some languages organize nouns into many "gender" categories (English does not). Each language is equally "functional" at meeting the communicative needs of its own speech community. But sometimes when two or more speech communities come into contact, one group will have more power, status or economic resources than the others. Not surprisingly, the language variety of that dominant group is often perceived as having higher status as well, especially if speaking it affords increased access to power or wealth. By comparison, the language varieties spoken by the less powerful groups often are stigmatized as "incorrect" or "bad" language.

Linguists approach language in the same way that astronomers approach the study of the universe or that anthropologists approach the study of human cultural systems. It would be ridiculous for astronomers to speak about planets orbiting stars "incorrectly" and inappropriate for anthropologists to declare a culture "degenerate" simply because it differs from their own. Similarly, linguists take language as they find it, rather than attempting

to regulate it in the direction of preconceived criteria. Linguists are equally curious about all the forms of language that they encounter, no matter what the education or social standing of their speakers might be.

The fact that, in most societies, some varieties of language are perceived as "correct" while others are considered "incorrect" is, for linguists, a social phenomenon – an aspect of language use to be explored scientifically. Since "correct" language is *inherently* no better or worse than the varieties that are considered "incorrect," linguists eagerly seek to discover the reasons for the conviction that some part of language variability is superior to the rest, and to examine the consequences of those beliefs.

One consequence of these kinds of language attitudes – in which one language variety is considered better than others – is the corollary belief that speakers of "incorrect" varieties are somehow inferior, because they will not or cannot speak "correctly." Their "incorrect" language is then used to justify further discrimination – in education and in employment, for example. Discrimination on the basis of language use is based on two false propositions: that one variety of language is inherently better than others, and that people can be taught to speak the "correct" variety. However, so powerful are the natural forces that guide how a person learns and uses spoken language that explicit teaching on how to speak is virtually irrelevant. If a person is not very good at mathematics, we are probably justified in assuming that he or she did not learn mathematics in school. The same may well be true of reading and writing; if someone cannot read or write, it is likely that something went wrong with that person's schooling. But the same is not true with spoken language. A person who uses double negatives, as in *She can't find nothing* (**negative concord**, as linguists call it), or says *knowed* for *knew*, may have received the best instruction in the rules of traditional grammar from the most skilled teachers available. However, just knowing what the rules are, or even practicing them for a few minutes a day in school, will be as effective in influencing how someone speaks as a meter-high pine tree would be in stopping an avalanche. The most powerful feature influencing spoken language is its ability to mark a person's identity as a member of the group closest to him/her in everyday life. This power trumps grammar instruction in classrooms every time.

Even the best-educated speakers of American English will not say "For what did you do that?" (which is formally correct); they'll say "What did you do that for?" Nor will they say "Whom did you see today?"; instead it will be "Who did you see today?" For exactly the same reason, a speaker of nonstandard English will say "I ain't got none," knowing that "I don't have any" is considered correct – in either case, to use "correct grammar" would make the speaker sound posh or snobbish and cost him/her the approval of his/her peers. There is an enormous disincentive to use language in a way that makes it seem that you are separating yourself from the people who are most important to you.

In fact, people who speak in close to the approved way did not learn to do so in school. They are just fortunate to come from the segment of society

that sets the standards for correct speech. This segment of society also controls its schools – and the language variety used and taught in its schools. Ironically, when children learn to use the socially approved variety of spoken language in school, it is not from what their teachers explicitly teach in class, but rather from adjusting their speech to match the speech of the other children in the halls, on the playground, and outside of school, and thus gain their approval.

Defining language

A precise definition of language is not easy to provide, because the language phenomenon is complex. Slightly modifying a definition provided by Finegan and Besnier (1989), we might define language as *a finite system of elements and principles that make it possible for speakers to construct sentences to do particular communicative jobs*. The part of the system that allows speakers to produce and interpret grammatical sentences is called **grammatical competence**. It includes the knowledge of what speech sounds are part of a given language and how they may and may not be strung together. Grammatical competence also includes knowing the meanings signified by different sound sequences in a language and how to combine those units of meaning into words, phrases and sentences. Grammatical competence is what allows a speaker of English to string together 21 sounds that sound something like "The dog chased the cat up the tree" and allows another speaker of English to understand what dogs, cats, and trees are, what chasing is, and which way is up. Further, grammatical competence is what allows these speakers of English to share the understanding that it was the dog doing the chasing and that it was the cat that went up the tree. Of course this does not apply only to English. Grammatical competence contributes similarly to comprehension in all human languages.

But people use language to do far more than just communicate the literal meanings of grammatical sentences. The sentence "The dog chased the cat up the tree" might be used to accomplish a wide variety of jobs: to narrate part of a story, to complain to the dog's owner, to help the cat's owner find his pet. The second part of the definition, "to do particular communicative jobs," refers to the notion **communicative competence**. The most frequent "job" that people do with language is communicating with other people.

Grammatical competence is almost useless for human interaction without communicative competence. In fact, a lot of the actual use of language is not in sentences at all, but in discourse units larger and smaller than sentences, some grammatical (in the technical sense used in formal linguistics), some not. To be effective, speakers have to combine grammatical competence with the knowledge of how to *use* grammatical sentences (and other pieces of linguistic structure) *appropriately* to the purpose and context at hand. The two taken together comprise *communicative competence*. Communicative competence – the knowledge included in grammatical

competence plus the ability to use that knowledge to accomplish a wide range of communicative jobs – constitutes *language*.

The diversity of linguistics

Unlike other linguistics textbooks, each chapter in this book has been written by a linguist who teaches and does research in that area. The field of linguistics, like the phenomenon of language which it studies, is broad and diverse, and although linguists share some beliefs – in the descriptive approach, for example, and in the functional equality of all language varieties – they differ in some of the assumptions they bring to their analyses. Some chapters in this book – particularly those on phonetics and phonology, morphology, syntax, semantics/pragmatics, and historical linguistics – assume, to varying degrees, that the forms of language can be understood separate from their use. These chapters are primarily about language form. The topics of these chapters constitute what was considered the essential core of linguistics in the mid twentieth century. Since then, the discipline has expanded considerably, and this book is designed to represent that increased scope.

Today the field of linguistics studies not just the nuts-and-bolts of forms and their meanings, but also how language is learned (both as a first and second language), how it plays a central role in reflecting and creating the interactive and cultural settings of talk, how computers can be designed to deal with language, and how language is represented in our very brains. Because much of this expanded scope of the field involves intense study of people actually *using* language, some doubt has arisen about how separate, or "dual," grammatical and communicative competence actually are. The degree of distinction between grammatical and communicative competence is understood differently by different linguists, even among the contributors to this book.

Chapter 3, "The structure of sentences," for example, takes a strongly dualist stance, presenting syntax as essentially about form, not use, while recognizing the supplementary contributions that more communicatively oriented approaches might make. Chapter 4, "Meaning," is largely about the boundary between semantic meaning (grammatical competence) and pragmatic meaning (communicative competence), all the while assuming the validity of duality. A somewhat similar contrast is seen in Chapter 1, "The sounds of language," involving the boundary between phonetics – the observable and observed phenomena of pronunciation – and the more abstract phonological aspects of pronunciation systems – a part of grammatical competence.

Chapter 8, "Language change," and Chapter 12, "Writing," implicitly present both language form and communicative use, but with less concern for strict boundaries between them. In Chapter 5, "Discourse," in Chapter 6, "Child language acquisition," and in Chapter 13, "Second language acquisition," the dual viewpoint is respectfully considered, but the authors' sympathies are with the primacy of communicative competence. In Chapter 10, "Language and culture," we find the clearest departure from a dualist

perspective, as the author suggests replacing the terms "language" and "culture" with "languaculture" to emphasize the inseparability of language from its context – in essence, the inseparability of grammatical and communicative competence.

How are we to understand these differences in emphasis? Perhaps linguists are like the blind men in John Godfrey Saxe's version of the Indian legend about the blind men and the elephant, each finding language to be like what we see in our own research. Perhaps each of us is "partly in the right and all [of us are] in the wrong." Perhaps some approaches are proceeding down blind alleys and will eventually have to be abandoned. Or perhaps someone will discover a brilliant unifying insight, like general and special relativity and quantum field theory in physics, which shows how all these approaches are related. Our understanding of the world and of human behavior has advanced in all these ways. We introduce our readers to the study of language in its current breadth and diversity of approaches with their tensions and unresolved issues.

How to approach this book

We close this introduction with a suggestion for how to approach this book. While we would like to think that all of it will be useful to everyone, language is a broad and diverse phenomenon and there are many ways to explore such a vast terrain. Different aspects of language – and linguistics – are more relevant to various other academic disciplines, and so we propose several "routes" through different territories of language.

- **General linguistics:** The sounds of language; Words and their parts; The structure of sentences; Meaning; Language change; Writing; Dialect variation; Discourse
- **Language acquisition, learning, and teaching:** The sounds of language; Words and their parts; The structure of sentences; Meaning; Child language acquisition; Second language acquisition; Discourse; Language in culture; The politics of language
- **Sociolinguistics:** The sounds of language; Words and their parts; Meaning; Language change; Writing; Dialect variation; Discourse; Language in culture; The politics of language
- **Language processing:** The sounds of language; Words and their parts; The structure of sentences; Meaning; The brain and language; Computational linguistics; Child language acquisition
- **Linguistics and other fields:** Child language acquisition; Language in culture; Computational linguistics; The politics of language; The brain and language; Writing; Dialect variation; Discourse; Second language acquisition
- The companion website for this textbook can be found at www.cambridge.org/fasold. There you will find sound files, further reading suggestions, additional exercises, links, and electronic versions of key figures.

CHAPTER

1 The sounds of language

ELIZABETH ZSIGA

CHAPTER PREVIEW

KEY TERMS

acoustic phonetics
active and passive articulators
allophone
alternation
articulatory phonetics
complementary distribution
derivation
distinctive features
fundamental frequency
formant
intonation
manner of articulation
minimal pair
natural class
obstruent
phoneme
phonology
phonotactic constraint
pitch track
place of articulation
sonorant
sonority
source-filter theory
spectrogram
stress
suprasegmentals
syllable structure
tone
vocal tract
voicing
waveform

This chapter is about the sounds of speech. Without sound, communication can still take place – with a nod or a wave, a photograph or a drawing. There can even be language without sound: those who cannot hear use languages based on manual signs instead. Yet for most of us most of the time, getting our message across involves encoding it in sounds. Even when we write, we use symbols that are based on speech (though sometimes not very directly).

The study of the sounds of speech can be divided into the disciplines of **phonetics** and **phonology**. Phonetics studies speech sounds as physical objects. Phoneticians ask questions such as:

- How are speech sounds made?
- How many different sounds do languages use?
- How does sound travel through the air?
- How is it registered by the ears?
- How can we measure speech?

Phonology studies how languages organize sounds into different patterns. Phonologists ask questions such as:

- How do languages organize sounds to distinguish different words?
- How do languages restrict, or constrain, sequences of sounds?
- What sorts of changes (alternations) do sounds undergo if sequences arise that don't obey the restrictions?
- How are sounds organized into larger constituents (syllables, words, phrases)?

We begin with phonetics, the study of how speech sounds are made and perceived, and then discuss phonology, the study of how a language organizes those speech sounds into a meaningful system.

GOALS The goals of this chapter are to:

- describe the basic anatomy of the vocal tract
- explain how the structures in the vocal tract are controlled to make speech sounds
- show how to transcribe English words using IPA transcription
- describe the basic properties of suprasegmental aspects of speech, and how languages differ in their use of them
- describe some of the physical properties of sound waves
- interpret some basic aspects of waveforms, pitch tracks, and spectrograms
- explain phonemic and allophonic distributions
- describe some of the most common phonological alternations
- introduce some of the major goals of phonological theories

Articulatory phonetics

The tools of phonetics

One of the biggest obstacles phoneticians face is that they can't see the objects they are studying. You can't see the tongue as it's moving around inside someone's mouth; you can't see the sound waves traveling through the air; you can't see the vibration of the fluid in the inner ear. Since ancient times, however, phoneticians have made the best of the information they had access to, employing careful listening, measuring, modeling, and notation. In addition, more sophisticated devices have been developed within the past decades – devices such as Magnetic Resonance Imaging (MRI), sonography, and digital acoustic analysis. Figure 1.1 shows some pictures of the vocal tract as seen by these devices.

With these aids, what have we learned about how humans make and hear speech sounds?

The vocal tract

Basically, sound is vibrating air. Speaking means using your **vocal tract** (lungs, trachea, larynx, mouth, and nose) to get air moving and vibrating, and then shaping that movement in different ways. Figure 1.2 shows a diagram of the upper parts of the vocal tract.

Most speech sounds are made with air exiting the lungs; therefore, speech begins with breath. To begin to speak, you pull down your **diaphragm**, the big muscle that separates your chest cavity from your stomach. This enlarges the lungs, which draws air in. Then the diaphragm relaxes and the muscles around the ribs contract, slowly squeezing the lungs and forcing the air out and up the windpipe, or **trachea**.

At the top of the trachea is a little box of cartilage, called the **larynx** (the "adam's apple"). Inside the larynx, two folds of soft tissue, called the **vocal folds** (sometimes called 'vocal cords'), lie across the top of the trachea. If the vocal folds are held in the correct position with the correct tension,

A

B

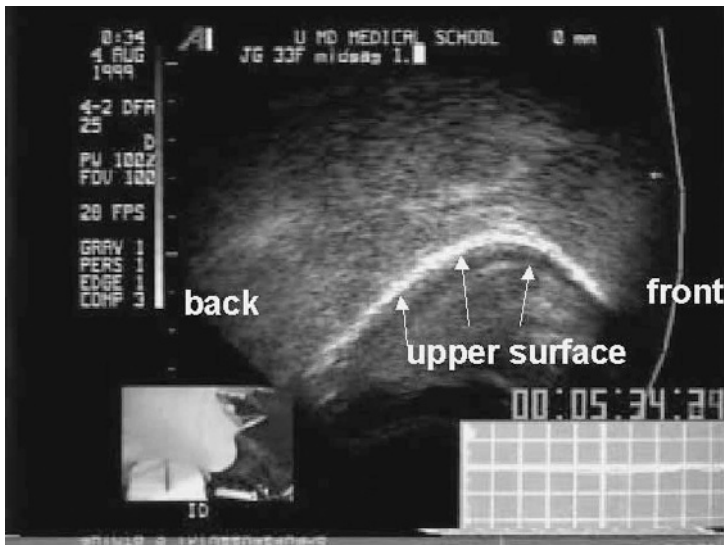

C

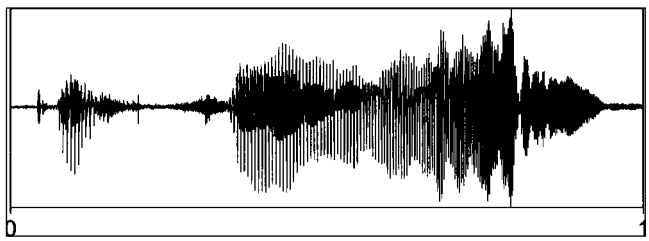

FIGURE 1.1
Views of the vocal tract. A. A magnetic resonance image of a mid-sagittal section (sideways slice) of the vocal tract B. A sonograph image of the surface of the tongue. C. A digital waveform showing sound pressure variations during one second of speech. (Images A and B courtesy of Dr. Maureen Stone, Vocal Tract Visualization Laboratory, University of Maryland, Baltimore)

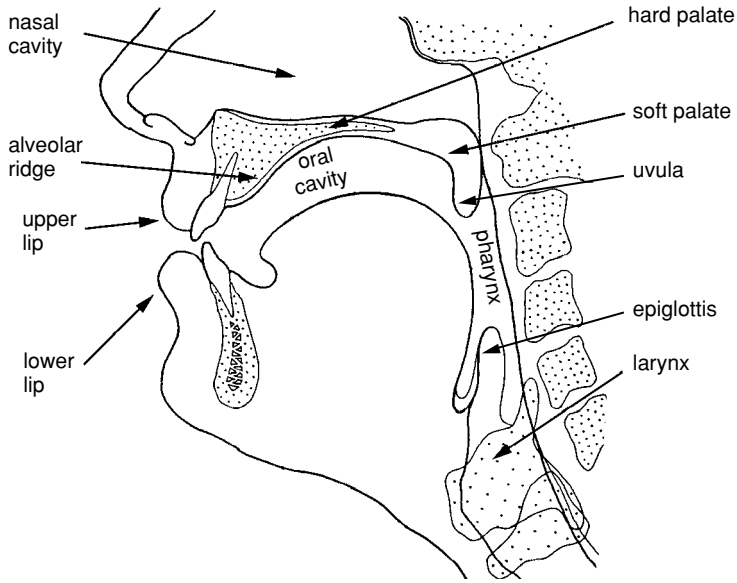

FIGURE 1.2
Parts of the vocal tract

the air flowing out of the trachea causes them to flap open and closed very quickly (around 200 times per second). You can feel this opening and closing motion as vibration in your throat. Find your larynx (you should be able to feel the bump of the adam's apple at the front of your throat), and then hum a tune. Muscles attached to the cartilages of the larynx allow you to adjust the tension of the folds, thus adjusting the rate of vibration and raising or lowering the pitch. The faster the vibration, the higher the pitch of the voice. Other muscles also allow you to draw the folds apart so that no vibration occurs.

Just above the larynx, at the base of the tongue, is the **epiglottis**. The epiglottis is a muscular structure that folds down over the larynx when you swallow to prevent food from going down into the lungs before it enters the passage to the stomach. The payoff for the risk of a larynx located low in the throat is an open area at the back of the mouth, the **pharynx**. The pharynx allows the tongue freedom for front and back movement. Other mammals, including nonhuman primates, have the larynx high up at the back of the mouth, connected to the nasal passages. Because they have no pharynx, chimps could never learn to talk. (This is why scientists who try to teach primates to communicate with language use gesture-based languages instead.)

Inside the mouth itself, there are many different structures – **active articulators** and **passive articulators** – that we use to shape speech sounds as the air passes through the vocal tract. The active articulators move toward the passive articulators in order to constrict and shape the air that is moving out from the lungs. Active articulators include the lips, which can be opened or closed, pursed or spread, and the tongue. What we usually see of the tongue is the small, pink tip, but it is actually a large mass of interconnected muscles that fills the floor of the mouth.

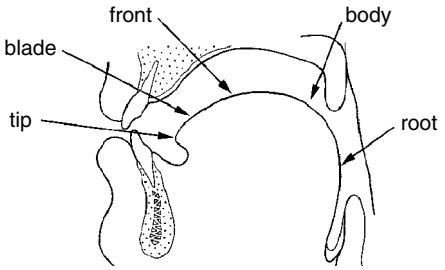

FIGURE 1.3
Areas of the tongue

Although the tongue has no bones or cartilage, different parts of the tongue can move fairly independently. The **tongue front** (including the **tip** and the **blade**, which extends a few centimeters back from the tip), the **tongue body** (the main mass of the tongue, also known as the dorsum), and the **tongue root** (the lowest part of the tongue, back in the pharynx), are considered separate active articulators. (See Figure 1.3.)

The passive articulators lie along the top of the vocal tract. Run your tongue along the top of your mouth beginning behind your upper teeth. You will first encounter the **alveolar ridge**, the bony rise just behind your teeth. The postalveolar region arches from the alveolar ridge toward the **hard palate**, the roof of the mouth. If you curl your tongue very far back in your mouth, you can feel that the bony structure of the hard palate gives way to softer tissue, which is known as the **soft palate**, or **velum**. The velum is a muscular structure that regulates the velar port, the opening in the back of the mouth that connects the mouth and nose. When the velum is lowered, as it is for breathing and for some sounds such as [m] and [n], the port is open and air flows freely between the nose and lungs. (It's a phonetic convention to write the symbols for sounds within square brackets. See the following sections for more on phonetic writing.) When the velum is raised, as it is for most speech sounds, the opening to the nose is closed off and all the airstream is directed through the mouth. At the very end of the velum is the **uvula**, the little pink pendulum you can see hanging down in the back of your mouth when you open wide and say "ah."

Articulation

Speaking involves using the structures of the vocal tract in different ways to control and shape moving air. We can think of the speaker producing the right combinations by making "choices" about which active and passive articulators to use and about how different constrictions will be made. These choices are not conscious; they are automated from long practice, just like the muscular routines of walking or reaching. When you reach for a cup of coffee, you don't say to yourself, "OK, now contract the tricep, relax the bicep," etc. Instead, without much conscious thought you select a goal, such as "fingers to mug," and your long-practiced routines of motion execute the goal-directed movement. Speech works the same way. The movements of speech are goal-directed gestures. Each sound is comprised of a set of articulatory goals that will get the vocal

tract in the right positions to make the sound you wish to make. An overall goal like "make an [m]" can be broken down into a set of subcomponents: "Close the lips," "open the velum," "make the vocal folds vibrate." These subroutines can be recombined in different ways to make different sounds, like a set of Lego blocks that can build a castle or a boat depending on the way they're put together.

The first choice is that of airstream mechanism: how will the speaker get the air moving in the first place? The usual choice is **pulmonic egressive** – that is, air moving out from the lungs. Most sounds used by most of the world's languages are pulmonic egressive. However, it is also possible to get air moving in other ways, such as by moving the larynx up or down, or by popping little pockets of air made with the tongue against the roof of the mouth, as in clicks. (Clicks include the expression we write in English as *tsk tsk* or *tut tut*, but in some languages of southern Africa sounds such as these are incorporated into the stream of speech as regular consonants like [p] or [t] in English.) The rest of this chapter discusses only sounds that are pulmonic egressive.

The second choice is what to do with the vocal folds. Sounds produced with vocal fold vibration are **voiced**; sounds produced without vocal fold vibration are **voiceless**. If you place your finger on your larynx and produce a sustained [z], you should be able to feel the vibration, or voicing. If you switch to [s], a voiceless sound, the vibration ceases. For some sounds, as in the initial [p] in *pop*, the vocal folds are held apart far enough and long enough to allow an extra "puff of air" to exit the mouth at the end of the [p]. This is called **aspiration**. You can feel the extra release of air if you hold your fingertips an inch or so in front of your lips as you say *pop* or *pill*. Aspiration can be indicated by a superscripted *h*: [p^h].

Besides deciding what to do with the larynx, the speaker must decide whether the velum will be open or not. If the velum is open, so that air flows into the nose, the sound is **nasal** (like [m]). If the velum is closed, the sound is **oral**.

Finally, the speaker must decide which active articulator will be used to make a constriction (lips, tongue front, tongue body, tongue root), where the constriction will be made (the **place of articulation**), and what sort of constriction will be made (the **manner of articulation**). The various places of articulation are discussed in following sections; we turn first to the various manners of articulation.

Manners of articulation

The manners of articulation include: stop, fricative, affricate, approximant, and vowel.

If the active and passive articulators are brought together to make a complete closure, so that airflow out of the mouth is completely cut off, the manner of articulation is a **stop**. The sounds [p], [t], and [k] in English are stops. Say the word *poppa* very slowly, and note that there is complete silence, with no air exiting the mouth, while the lips are closed for [p] in the middle of the word. You may even feel pressure building up behind the lips, as air continues flowing from the lungs and has nowhere to go.

This pressure is released with a slight pop, or burst, when the lips are opened. The sound [m] is a nasal stop. Even though the velum is open and air flows freely out of the nose, so that you can hum a tune while producing an [m] sound, the manner of articulation is still a stop, because the lips are completely closed, as they were for [p]. (Try pinching your nose closed for a moment while you're humming a tune, and see what happens.)

If the articulators are brought close together but not closed completely, so that the stream of air that is forced between them becomes turbulent and noisy, the manner of articulation is a **fricative**. The sounds [s], [z], [f], and [v] are fricatives. **Affricates** combine a sequence of stop plus fricative in a single sound. The sound usually written *ch* in English is an affricate. Try saying the word *achoo* as slowly as possible, paying attention to the movement of the tongue between the *a* and *oo* sounds. You first make a closure with the tongue front at or just behind the alveolar ridge, and then lower the tongue tip to let the air out through a narrow constriction slightly further back, between the tongue blade and postalveolar region.

If the active articulator moves to narrow the vocal tract, but not so much that fricative noise is created, the manner of articulation is an **approximant**. **Glides**, such as the sounds at the beginning of the words *yell* and *well*, are approximants, as are [l] and [r] in English. The *l*-sounds of the languages of the world are called **laterals**, because air flows out over the sides of the tongue. Try drawing out the initial sound in the word *lateral*. Now, without moving your tongue, take a deep breath in and out. You'll feel the air moving over the sides of the tongue. The *r*-sounds are called **rhotics**. The rhotic sounds of the languages of the world are quite varied, including quick taps of the tongue against the alveolar ridge, trills in which the tongue is set into vibration by air flowing over it, and the very odd shape of the American English [r], in which the body of the tongue is bunched up high and the tongue tip may be raised or curled backwards. (It is no surprise that non-native speakers of English have trouble with this sound.) **Vowels** are the most open manner of articulation. Different vowel sounds are made by moving the tongue body up or down, front or back, and by rounding or spreading the lips. During all vowel sounds, however, the vocal tract is relatively wide open, and air flows out freely. Oral stops, fricatives, and affricates together form a class of sounds called **obstruents**, because they make noise by obstructing the airflow in the vocal tract, causing a burst of sound as a closure is released or a hissing sound as the air passes through a narrow constriction. Nasal stops, approximants, and vowels (anything that's not an obstruent) form a class of sounds called **sonorants**. They make audible sounds not by obstructing the airflow, but by letting the air resonate. Sonorant sounds are almost always voiced. The vibration of the vocal folds causes the air inside the vocal tract to vibrate. If the vibration is strong enough, it produces an audible sound, like the ringing of a bell. Different vocal tract shapes (which we control by moving the active articulators) produce different patterns of vibration, which we hear as different sounds (more on this below). It is possible to produce voiceless sonorants, by moving a large volume of air through the open vocal tract. Languages like Hmong and

Burmese use voiceless nasals. Listen carefully, and you'll hear that the [l] in an English word like *play* is also voiceless.

Writing sounds: transcription

Before we discuss the different places of articulation used in English and other languages, we have to consider how to write down different sounds. Descriptive phrases like "the sound at the beginning of the word *yell*" or "in the middle of *achoo*" are cumbersome. We need a **phonetic alphabet**. Writing down sounds using a phonetic alphabet is called **phonetic transcription**.

In 1888, the International Phonetic Association (based in Paris) tackled the problem of how to precisely describe any sound the members might encounter in their efforts to describe all the languages of the world. They published symbols for a new alphabet, the *International Phonetic Alphabet* (IPA), based on two principles:

- The alphabet would be universal. There would be enough symbols so that every sound in every human language could be represented. (As new sounds have been discovered or old sounds reanalyzed, the IPA has been revised to incorporate the new findings. The latest revision took place in 1996.)
- The alphabet would be unambiguous. Every sound would have one symbol, and every symbol one sound.

Consider how English spelling falls short on these two principles. We certainly don't have letters for the clicks of southern Africa or the sounds made in the back of the throat in Arabic. And the one sound/one symbol correspondence is constantly violated. How many ways can you pronounce the letter *c*? (Consider *each vicious circle*.) The letter *x* stands for a sequence of two sounds, [k] followed by [s] (*box* rhymes with *locks*), while the sequence of letters *sh* stands for a single sound (a fricative made with the tongue blade against the postalveolar region).

Figure 1.4 shows the IPA symbols for all the sounds discussed in this chapter, and more.

In an IPA consonant chart, place of articulation is written across the top, and manner of articulation is written down the side, so that each cell indicates a combination of a specific manner and place. Cells that are shaded gray are physically impossible combinations. If there are two symbols in a cell, the one on the left is voiceless, the one on the right is voiced.

A good place to start learning IPA transcription is with the symbols for the consonants of English, which are given in Table 1.1, along with example words in which the sounds occur. Many, though not all, of the symbols will be familiar. As you learn the definitions of the different places and manners of articulation, you should be able to combine them to figure out the pronunciations of unfamiliar symbols in the IPA chart.

As you look through the example words in Table 1.1, you'll notice that not all of the cells are filled. Some of these gaps are *accidental*: English just doesn't happen to have a word *nen*, for example, though it has similar words like *nine* and *net*. Other gaps are *systematic*: no word in English ends

CONSONANTS (PULMONIC)

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap				ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Where symbols appear in pairs, the one to the right represents a voiced consonant. Shaded areas denote articulations judged impossible.

CONSONANTS (NON-PULMONIC)

Clicks		Voiced implosives		Ejectives	
ʘ	Bilabial	ɓ	Bilabial	’	Examples:
|	Dental	ɗ	Dental/alveolar	pʼ	Bilabial
!	(Post)alveolar	ʄ	Palatal	tʼ	Dental/alveolar
ǂ	Palatoalveolar	ɠ	Velar	kʼ	Velar
‖	Alveolar lateral	ʛ	Uvular	sʼ	Alveolar fricative

VOWELS

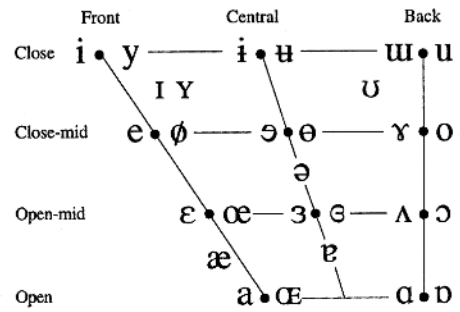

Where symbols appear in pairs, the one to the right represents a rounded vowel.

OTHER SYMBOLS

ʍ	Voiceless labial-velar fricative	ɕ ʑ	Alveolo-palatal fricatives
w	Voiced labial-velar approximant	ɺ	Alveolar lateral flap
ɥ	Voiced labial-palatal approximant	ʩ	Simultaneous ʃ and x
ʜ	Voiceless epiglottal fricative		
ʢ	Voiced epiglottal fricative		
ʡ	Epiglottal plosive		

Affricates and double articulations can be represented by two symbols joined by a tie bar if necessary.

k͡p t͡s

SUPRASEGMENTALS

ˈ	Primary stress
ˌ	Secondary stress ˌfoʊnəˈtɪʃən
ː	Long eː
ˑ	Half-long eˑ
˘	Extra-short ĕ
|	Minor (foot) group
‖	Major (intonation) group
.	Syllable break ɹi.ækt
‿	Linking (absence of a break)

DIACRITICS Diacritics may be placed above a symbol with a descender, e.g. ŋ̊

◌̥	Voiceless	n̥ d̥	◌̤	Breathy voiced	b̤ a̤	◌̪	Dental	t̪ d̪
◌̬	Voiced	s̬ t̬	◌̰	Creaky voiced	b̰ a̰	◌̺	Apical	t̺ d̺
◌ʰ	Aspirated	tʰ dʰ	◌̼	Linguolabial	t̼ d̼	◌̻	Laminal	t̻ d̻
◌̹	More rounded	ɔ̹	◌ʷ	Labialized	tʷ dʷ	◌̃	Nasalized	ẽ
◌̜	Less rounded	ɔ̜	◌ʲ	Palatalized	tʲ dʲ	◌ⁿ	Nasal release	dⁿ
◌̟	Advanced	u̟	◌ˠ	Velarized	tˠ dˠ	◌ˡ	Lateral release	dˡ
◌̠	Retracted	e̠	◌ˤ	Pharyngealized	tˤ dˤ	◌̚	No audible release	d̚
◌̈	Centralized	ë	◌̴	Velarized or pharyngealized	ɫ			
◌̽	Mid-centralized	e̽	◌̝	Raised	e̝	(ɹ̝ = voiced alveolar fricative)		
◌̩	Syllabic	n̩	◌̞	Lowered	e̞	(β̞ = voiced bilabial approximant)		
◌̯	Non-syllabic	e̯	◌̘	Advanced Tongue Root	e̘			
◌˞	Rhoticity	ɚ a˞	◌̙	Retracted Tongue Root	e̙			

TONES AND WORD ACCENTS

LEVEL			CONTOUR	
e̋ or ˥	Extra high		ě or ˧˥	Rising
é ˦	High		ê ˥˦	Falling
ē ˧	Mid		e᷄ ˦˥	High rising
è ˨	Low		e᷅ ˨˧	Low rising
ȅ ˩	Extra low		e᷈ ˩˨˧	Rising-falling
↓	Downstep		↗	Global rise
↑	Upstep		↘	Global fall

FIGURE 1.4
The International Phonetic Alphabet

Table 1.1 IPA symbols for the consonants of English

	Initial				Final	Medial
p	pat	pie	pen	pin	whip	upper
b	bat	buy	Ben	bin	bib	rubber
m	mat	my	men	minion	whim	summer
f	fat	fight	fen	fin	whiff	suffer
v	vat	vie	vendor	vintage	live	ever
θ		thigh		thin	with	Ethel
ð	that	thy	then		bathe	weather
t	tat	tie	ten	tin	wit	retool
d	data	dye	den	din	mid	redo
n	Nat	night		ninja	win	renew
s	sat	sigh	sensor	sin	miss	presser
z	zap		zen	zip	wiz	buzzer
l	lateral	lie	lentil	lip	will	filler
ɹ	rat	rye	rent	rip	where	terror
ʃ	shack	shy	shell	ship	wish	pressure
ʒ					beige	measure
tʃ	chat	chai	check	chip	witch	etcher
dʒ	jack	giant	gender	gin	edge	edger
k	cat	kite	Ken	kin	wick	wrecker
g	gap	guy			wig	mugger
ŋ					wing	singer
h	hat	high	hen	hip		ahead
w	whack	why	when	win		away
j	yak		yen	yip		

with the sound [h], for example. (Plenty of words end with the *letter h*, but we are concerned with sound, not spelling.) Systematic gaps will be discussed further below.

Consonants

We will organize our discussion of the consonants by active articulator and place of articulation. Generally, each articulator can move to more than one place of articulation, as shown in Table 1.2.

The lower lip can make constrictions at two different places. If the lower and upper lip come together, the sound is **bilabial**. The sounds [p], [b], and

Table 1.2 Active articulators, passive articulators, and place of articulation

Active articulator	Passive articulator	Place of articulation
lower lip	upper lip	bilabial
	upper teeth	labiodental
tongue tip or blade	upper teeth	dental
	alveolar ridge	alveolar
	postalveolar region	retroflex (tip)
	postalveolar region	alveopalatal (blade)
	hard palate	palatal (blade)
tongue body	hard palate	palatal
	soft palate	velar
	uvula	uvular
tongue root	pharyngeal wall	pharyngeal
larynx		laryngeal

[m] are bilabials. Note that [p] is voiceless and [b] and [m] are voiced. Alternatively, the lower lip can make contact with the upper teeth to produce a **labiodental** sound. [f] and [v] are labiodentals. Japanese has a bilabial fricative (IPA [ɸ] instead of [f]). This sound, as in the native Japanese pronunciation of *futon,* is made by blowing through pursed lips, the same motion as blowing out a candle.

The lower lip is rather limited in the places at which it can make a constriction. The tongue front is the most versatile of the active articulators, moving to at least four different places of articulation. The tongue tip moves forward to the upper teeth for the sounds at the beginning of *thin* and *then*. These **dental** fricatives are written [θ] (voiceless) and [ð] (voiced).

The English sounds [t], [d], [n], and [l] are made with the tongue tip at the alveolar ridge, the **alveolar** place of articulation. The fricatives [s] and [z] are also alveolar. For these fricatives the tongue forms a narrow groove under the alveolar ridge like a spout that shoots a stream of air against the teeth, producing a high-pitched hissing sound. Though the place of articulation for these fricatives is alveolar, the front teeth are necessary to create the proper high-pitched hiss, as all children discover when they lose their baby teeth.

The fricatives [ʃ] and [ʒ] (as in the middle of *pressure* and *measure*) are made further back, with the blade of the tongue making a constriction at the **palatoalveolar** place of articulation. (Interestingly, [ʒ] doesn't occur in initial position in English, except in obvious borrowings from French, such as *genre*.) The affricates in *church* and *judge* are also palatoalveolar. The IPA symbols for these sounds are [t͡ʃ] and [d͡ʒ]: two symbols are used for the combination of stop plus fricative, linked by a ligature. (In other transcription systems commonly used in linguistics books, [ʃ], [ʒ], [tʃ], and [dʒ]

are written with hatchecks: [š], [ž], [č], [ǰ].) English doesn't have any palatoalveolar stops or nasals, though other languages do (for example, French *agneau* [aɲo] 'lamb', Spanish *año* 'year').

Usually the blade of the tongue is used to make a palatoalveolar constriction. It is also possible, however, for the tip of the tongue to curl back to make a constriction in this area. If the tip of the tongue curls back, the sound is called **retroflex**. (The IPA symbols for these sounds have a little hook under the symbol, recalling the curling back of the tongue.) For some (but not all) American speakers, [ɹ] is a retroflex approximant. Can you determine whether your own tongue tip curls back in a word like *road*? There are no other retroflex sounds in English, though other languages, notably Hindi and other languages of India, have a full set of retroflex stops, fricatives, and nasals. A telltale sign of an Indian accent in English is substituting retroflex stops for English alveolars.

The sound at the beginning of the English words *you* and *yacht* is **palatal** (a palatal glide, to be exact). The whole middle section of the tongue, including blade and body, is pushed straight up to narrow the space between the tongue and hard palate. The IPA symbol for a palatal glide is [j]. (Think Scandinavian *ja*.) English doesn't have any other palatal sounds, but they're not hard to make. Start with [j] (as in *you*), and then make it voiceless. The result is a voiceless palatal fricative, used for example in German words such as *ich* [iç], meaning 'I'.

Moving further back in the vocal tract, the next place of articulation is **velar**, in which the tongue body moves up to make constriction against the velum, high in the back of the mouth. The English sounds [k] and [g] are velar stops. In English, the sequence of letters *ng*, as at the end of *song* or *ring*, usually indicate a velar nasal. In the word *song*, you don't make a sequence of alveolar nasal followed by velar stop (*n-g*), but a single nasal sound at the same place as [k] or [g]. (Feel how little your tongue moves when you say the word *king*.) The IPA symbol for a velar nasal stop is [ŋ]. As with [ʒ], English uses [ŋ] only at the end of words (*song*) or in the middle (*singer*), never at the beginning, although with practice you can learn to pronounce words like Thai [ŋaː] 'tusk' or Australian names like Ngaio. The native German sound at the end of the name *Bach* is a voiceless velar fricative, [x]. To make this sound, begin with a [k], then loosen the constriction slightly, letting a little turbulent airflow pass through.

The tongue body can also make constrictions further back, at the **uvular** place of articulation. To make a uvular stop, begin with a [k] or [g], then move the tongue a few centimeters back. Uvular stops are common in many Native American languages. Constrictions can also be made deep in the throat, with the tongue root moving back toward the pharyngeal wall. Voiced and voiceless **pharyngeal** fricatives are found in Arabic and Hebrew.

Finally, consonants can be made with the larynx as the only articulator. The sound [h] consists of the noise of air rushing through the open vocal folds, and may be considered a **laryngeal** fricative. It is also possible to close the vocal folds up tight, stopping the airflow at the larynx, a **glottal stop** (IPA [ʔ]). This is the sound in the middle of the English expression

uh-oh. If you pronounce this slowly, you can feel the constriction in the larynx. In other languages, like Hawai'ian, the glottal stop is used as a regular consonant. In the word *Hawai'i* the apostrophe stands for a glottal stop.

One English consonant remains to be discussed: the glide [w], as in *wear*. This sound combines a narrowing of the vocal tract at the velar place of articulation with rounding of the lips. It is thus a **double articulation**, a labiovelar glide. While double articulations at various places of articulation are not hard to make (given the independence of the active articulators), they can be hard for the ear to distinguish, so double articulations other than labiovelars are rare.

In summary, there are eleven common places of articulation – bilabial, labiodental, dental, alveolar, alveopalatal, retroflex, palatal, velar, uvular, pharyngeal, and laryngeal – though no single language makes consonants using all of the places of articulation.

Vowels

Vowels are harder to describe than consonants. By definition, vowels have an open vocal tract, so the tongue doesn't actually touch the upper surface of the vocal tract at any particular place and the term *place of articulation* isn't really appropriate. Instead, different vowels are described in terms of the ways in which the tongue body and lips move. Linguists classify vowels by the height of the tongue body, whether it is bunched toward the front or back of the mouth, and whether the lips are rounded.

If describing vowel systems in general is a difficult task, describing the vowels of English is even more so. One reason is because there are a lot of them. The most common number of vowels for a language to have is five. Though English writers use just five letters to encode their vowels (relics of an older system), the English language uses more than a dozen different vowel sounds. Another reason is because the exact number of vowels and exact vowel quality differ from dialect to dialect, much more so than for the consonants. For example, for most speakers on the east coast of the United States, the words *caught* and *cot* have two different vowel sounds; but for most speakers on the west coast, the two words are pronounced the same. Table 1.3 gives the IPA symbols, with example words, for the vowels of "General American" English – that is, English as it is more or less spoken in the central United States. The exact vowel sounds you use may well be slightly different.

Figure 1.5 charts the positions of English vowels relative to each other, based roughly on the position of the highest point of the tongue during that vowel. The vowel space is larger at the top than at the bottom, because there is more room for tongue movement closer to the palate. (The astute reader will notice some differences in exactly where the vowel symbols are placed in Figures 1.4 and 1.5, particularly for low and central vowels. This is because the IPA provides many more symbols than are needed in any one language. If a linguist is concerned primarily with a general description of the set of sounds that are needed to distinguish the words of a language (as we are here), he/she will usually choose the more familiar symbols from among those the IPA makes available.)

Table 1.3 IPA symbols for the vowels of English				
i	bead	key	he	
ɪ	bid	kit		
e	bade	kate	hey	
ɛ	bed	ketchup		
æ	bad	cat		
u	booed	coot	who	
ʊ	book	cook		
o	bode	coat	hoe	
ɔ	baud	caught	haw	
a	body	cot	ha	
ʌ	bud		cut	
aʊ	bowed	count	how	
ɔɪ	boy	coy	ahoy	
aɪ	bide	kite	high	
ə	about			rosa's [roz´z]
ɨ				roses [roz^z]

The terms we use to classify different vowels refer to the highest point of the tongue during the vowel. The tongue body moves up for the **high** vowels [i, ɪ, ɨ, u, ʊ], down for the **low** vowels [æ, a], and stays in the middle for the **mid** vowels [e, ɛ, o, ɔ, ʌ, ə]. The tongue moves forward in the mouth for the **front** vowels [i, ɪ, e, ɛ, æ] and backward for the **back** vowels [u, ʊ, o, ɔ, a]. The vowels [ɨ, ʌ, ə] are **central**. Vowels also differ with respect to lip

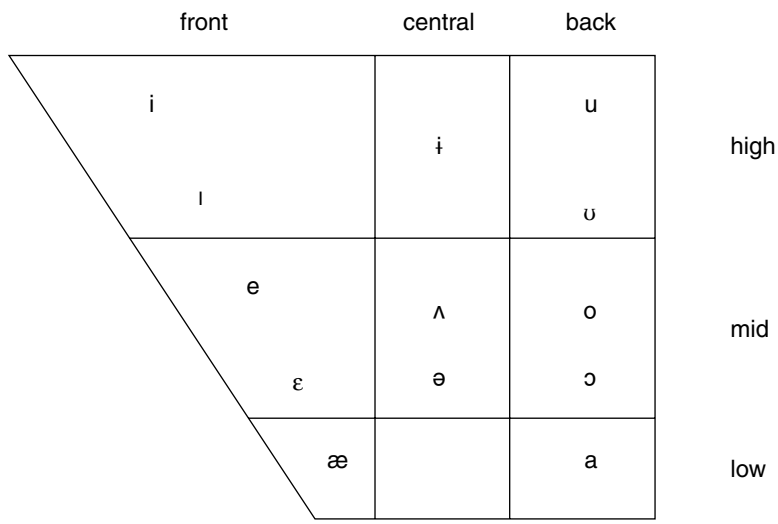

FIGURE 1.5
English vowels

rounding. In General American English, the back vowels [u, ʊ, o, ɔ] are **round**, all other vowels are **unround**. Finally, English divides its vowels into two sets, **tense** and **lax**. The tense vowels [i, e, o, u] are longer, slightly higher, and produced with greater stiffening of the tongue root than their lax counterparts [ɪ, ɛ, ɔ, ʊ]. The tense/lax distinction doesn't really apply to low vowels. These descriptive terms can be combined to pick out a specific vowel: [ɪ] is high, front, lax, unround; [o] is mid, back, tense, round.

Another sort of vowel is a **diphthong**, which combines two different positions in sequence. The diphthong [aɪ], as in General American *high*, moves from a low central position to high front. The diphthong [aʊ], as in General American *how,* moves from low central to high back. And [ɔɪ] moves from mid back to high front.

Two other English vowels are used only in short, unstressed syllables. (See the discussion of stress below.) The mid-central vowel [ə], called **schwa**, is heard in the first syllable of *about* and the second syllable of *rosa's*. The high, central [ɨ] occurs in the second syllable of *roses*.

As mentioned above, in English, all the nonlow back vowels are round, and all the low vowels and front vowels are unround. Combining lip and tongue position in this way makes it easier for the ear to distinguish the different vowel sounds. Any language that has at least three vowels will have front vowels that are unround and back vowels that are round. Some languages (such as French, Dutch, and German) also have front round vowels, and others (such as Japanese and Korean) also have back unround vowels. In linguistics, when a sound is unusual or difficult to hear or to say, we say that that sound is **marked**. The easier, more common sound is **unmarked**. Front round vowels and voiceless sonorants, for example, are marked. Front unround vowels and voiced sonorants are unmarked. If a language uses the marked version of a sound, it also uses the unmarked version.

Box 1.1 Summary of vocal tract choices

The terms and symbols of phonetics describe the choices that a speaker must make in order to produce a linguistic sound:

1. How should I get the air moving? Generally, this will be pulmonic egressive: air forced out of the lungs.
2. Which active articulator should I use: lips, tongue tip, tongue body, tongue root, or larynx?
3. What kind of constriction should I make: stop, fricative, affricate, approximant, vowel?
4. Where should I make the constriction? For consonants, the choices are bilabial, labiodental, dental, alveolar, alveopalatal, retroflex, palatal, velar, uvular, pharyngeal, and laryngeal; for vowels, the choices are high/mid/low, front/central/back, tense/lax, and round/unround.
5. Should the velum be open or closed?
6. What should I do with the larynx: voiced or voiceless, aspirated or unaspirated?

Introductory linguistics textbooks often simplify phonology by emphasizing the role of place and manner of articulation in making sounds, but you can see that there are many more choices involved. Speakers control the airstream mechanism, voicing, and nasality as well as the place and manner of articulation. Every sound is composed of smaller components that can be combined in different ways to make other sounds, and each of these components offers a (typically binary) opposition: voiced or voiceless, nasal or oral, open or closed, front or back, etc. Speakers of English are often biased by its alphabetic writing system; they automatically think of each sound as a letter – an autonomous atomic unit, equally related or unrelated to every other letter, so that /p/ and /b/ are no more closely or distantly related than say /g/ and /s/. But sounds are 'built' from lower-level vocal tract 'choices', and you change one sound into another by switching parameters for each choice (voiced to voiceless, stop to fricative, etc.). Think of 'voiced', for example, not just as an adjective that describes a sound but as one parameter (a choice, a building block, a specific vocal tract configuration) that, in combination with other parameters, creates the sound. The phonetic symbol representing a given sound isn't the sound itself, but a 'cover symbol' for the set of choices. This also means that the speech sounds of a language are related to each other in important ways; some sets of sounds differ only by changing a single parameter, while others differ in the settings of several parameters. As we will discover in the section on phonology, it is these parameters (the **distinctive features**) of a sound or group of sounds, not the individual sounds or symbols themselves, that are important in describing sound patterns within a linguistic system.

Suprasegmentals

Thus far, we have learned about individual sounds, but speaking involves stringing sounds together into larger units. Aspects of speech that influence stretches of sound larger than a single segment are called **suprasegmentals**. Suprasegmental aspects of speech include length, tone, intonation, syllable structure, and stress. Because suprasegmental aspects of speech involve the organization of sounds into larger units, the study of suprasegmentals straddles the domains of phonetics (the study of speech sounds as physical objects) and phonology (the study of how languages organize sounds into different patterns).

Length

Many factors influence how long it takes to articulate a given segment. Sometimes differences in vowel length are unintentional results of how different vowels are articulated. Low vowels, for which the mouth has to open wide, take longer to articulate than high vowels, for which little movement is necessary. In some languages, however, two segments may differ in length alone: the long segment and short counterpart are exactly the same, except that the former is (intentionally) held for a longer period of time, an extra "beat." These long segments may be written with a

double symbol ([aa], [pp]) or with a colon after the usual symbol ([a:], [p:]). (Sometimes a horizontal bar, or macron, is used to indicate a long segment.) For example, Japanese makes length distinctions in both vowels and consonants. In Tokyo, you want to be careful to order [bi:ru] 'a beer' rather than [biru] 'a building,' and to ask directions to a certain [tori] 'street,' rather than a [to:ri] 'gate' or [to:ri:] 'bird.' We will reserve the term "long vowel" to refer to distinctions such as those in Japanese, where everything about the long/short vowel pair is the same, except for length.

Long consonants are known as **geminates**. English can create long consonants when two words come together – compare *bookcase* [bʊkkes] to *book ace* [bʊkes] or *top part* [tappart] to *top art* [tapart] – but we don't distinguish long and short consonants within words. When double consonants are written, for example in *supper* vs. *super*, they actually tell us about the quality of the vowel, not the length of the consonant.

You probably learned in elementary school about long and short vowels in English spelling (like "long a" as in made *and "short a" as in* mad*), but this is more a distinction of vowel quality ([e] vs. [æ]) than of length. The long–short vowel terminology, while no longer linguistically accurate for English, is not completely random: 500 years ago, the difference between* made *and* mad *really was one of length, and the English vowel system was very similar to that of modern Japanese. Over the years, however, a series of sound changes affected the long and short vowels differently, pushing them out of alignment.*

Tone and intonation

The pitch of the voice carries a lot of information. It can tell you whether the speaker is a male or female, a large person or small, old or young. High pitch can tell you that a person is frightened; low pitch that he/she is angry. This sort of information isn't really linguistic, however, but physical or emotional. The terms tone and intonation refer to linguistic uses of pitch. **Tone** refers to the use of pitch to convey meaning at the word level; **intonation** refers to the use of pitch to convey meaning at the sentence or discourse level.

Intonation distinguishes different kinds of sentences or focuses attention on a particular word. For example, try reading the following sentences out loud (and dramatically):

"That's a cat?"

"Yup. That's a cat."

"A *cat*? I thought it was a mountain lion!"

The pitch of your voice moves in different directions on the word *cat*. On the first *cat*, pitch goes up, indicating a question. On the second, pitch falls, indicating a statement or confirmation. On the third *cat*, a more complicated fall–rise pattern indicates incredulity. (Typographically, we indicate these different "readings" with a question mark, period, and italics, respectively.) In each case, the sequence [kʰæt] refers to the same object, a feline. The pitch differences indicate only the role that the reference to the feline is playing in the current conversation: asking for information about the cat, providing it, or expressing disbelief regarding the information offered. All languages use intonation to some extent, though the patterns and meanings differ across languages.

In addition to intonation, most languages also use pitch to distinguish different words. In English, whether you say [kʰæt] with a rising pitch or falling pitch, the word still refers to a feline. In Thai, if you say [kʰa:] with rising pitch, it means 'leg'; but if you say it with falling pitch, it means 'value.' (There are actually five contrasting pitch patterns in Thai: high, low, mid, falling, and rising.) These words are as different as *cat* and *cut* to an English speaker. This use of pitch, to distinguish different words, is known as *tone*.

Although the idea of tones seems very strange to English speakers, the majority of the world's languages are tonal. The major European languages and their relatives are exceptional in *not* having tone.

Syllable structure

How many syllables are in the word *Appalachicola*? *Massachusetts*? *Antidisestablishmentarianism*? (6, 4, and 11, respectively.) English speakers have little trouble counting the number of syllables in a word, but linguists have a harder time defining what a syllable is.

One preliminary answer might be "a vowel and its surrounding consonants." Most of the syllables we encounter in English, in words like *pin*, *print*, or even *sprints*, fit this definition. However, it's perfectly possible to have a syllable without a vowel. We would all agree that *hidden* has two syllables, even if pronounced [hɪdn], with no vowel between the two consonants. Also, defining a syllable as "a vowel and the consonants around it" doesn't explain why some sequences of consonants are allowed and others are not. The sequence [prɪnt] is acceptable as a syllable in English, but the sequence [rpɪtn] is not acceptable as a single syllable in any language. Why is [prɪnt] a good syllable when [rpɪtn] is not?

The best answer (though not perfect) lies in the concept of **sonority**. Sonority can be defined as relative openness of the vocal tract, which corresponds directly to the relative loudness of a sound. The most sonorous sounds are the low vowels; the mouth is wide open, and the sound flows freely out. The least sonorous sounds are the voiceless stops; the mouth is completely shut, and no sound is made at all. Other sounds range between these two extremes.

The speech stream is organized into peaks and valleys of sonority. Languages generally do not choose long strings of consonants nor long strings of vowels. Rather, we alternate sounds that are more sonorous and less sonorous: each stands out better against the background of the other. A syllable, then, may be defined as a way of organizing sounds around a peak of sonority.

Take the simple syllable *pin*. The vowel [ɪ] is the most sonorous sound in the sequence, flanked by less sonorous consonants. Thus there is a single sonority peak, and a single syllable. The syllable *print* also follows the principle of sonority. Sonority rises from [p] (voiceless stop) to [r] (rhotic) to [ɪ] (vowel), then falls from vowel to [n] (nasal) to [t] (stop). A single peak, a single syllable. Meanwhile, the sequence [rpɪtn] has three peaks; higher sonority [r], [ɪ], and [n] are interrupted by lowest sonority [p] and [t]. Thus (if it is pronounceable at all) it has three syllables, not one.

The most sonorous element of a syllable, the peak itself, is called the **nucleus**. Lower sonority sounds preceding the nucleus are called the **onset**; those following the nucleus are called the **coda**. The nucleus and coda together form the **rhyme**. A syllable structure tree diagram for the word *print* is shown in Figure 1.6. (In linguistics, such tree diagrams are often used to show how constituent parts of a larger unit are related, so you'll see more tree diagrams in other chapters!)

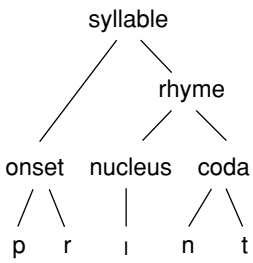

FIGURE 1.6
Syllable structure for the word *print*

Since vowels are the most sonorous sounds, they usually constitute syllable nuclei, but that's not always the case. Sounds other than vowels may form sonority peaks, so that *hidden* and *prism* have one vowel but two syllables.

Sonority thus seems to capture most of our intuitions about syllable structure and explains a lot about possible syllables in the languages of the world. But sonority doesn't account for everything. There are some English words that clearly violate the principle of sonority – *sprints* and *sixths*, for example. Linguists aren't sure exactly how to deal with words like this. It may be that endings like plural [s] and ordinal [θ] are not really part of the syllable at all; rather they're tacked on in an "appendix" to the end of the word.

Stress

Linguistic **stress** is a prominence relation between syllables: certain syllables are longer, louder, higher-pitched, or more clearly articulated than those around them. Just as we can generally count the syllables in a word, we can generally pick out the syllable that's most prominent: *phoNOlogy*, *phoNEtics*, *SYNtax*.

There are at least three different levels of stress in English. Consider the word *Alabama*. The third syllable [bæ] is the most prominent, and bears the *main or primary stress* of the word, but the other three syllables do not receive equal stress. The first syllable has a full vowel quality [æ], though it is not quite as long or loud as the third, but the second and fourth syllables are short and weak, with the tongue not moving far from its central position. We say that the first syllable has *secondary stress*, while the second and fourth are completely unstressed.

Some languages do not use stress at all. In Japanese, for example, all the syllables in a word are sometimes pronounced with equal prominence. This gives quite a bit of trouble to English speakers. In 1996, for example, when the Winter Olympics were held in Japan, English commentators had a great deal of trouble deciding how to say the name of the host city: *NAgano*? *NaGAno*? *NagaNO*? In fact, the word is pronounced with all syllables equally stressed (or equally unstressed), a difficult task for an English speaker. Conversely, Japanese speakers often have a great deal of trouble with reducing the unstressed vowels of English, sounding overly careful to English ears.

Because stress is a prominence relation, stressed and unstressed syllables tend to alternate across the word: *Apa**la**chi**CO**la*, for instance. In order for a syllable to be heard as prominent, it helps if it's surrounded by non-prominent syllables. In English, stress sometimes even moves around in order to accommodate an alternating pattern. In a word like *sixteen*, stress usually falls on the second syllable (*How old are you*? *Six**TEEN***.). But put the word next to one that begins with a strongly stressed syllable, and stress may shift back in order to maintain an alternating pattern (*How long have you worked here*? "***SIX**teen **YEARS***.) Linguists refer to a grouping of a stressed syllable and adjacent unstressed syllables as a **foot**. Choosing different kinds of feet (da-**DUM** vs. **DA**-dum) is important not only in poetry, but as part of defining a language's unfolding rhythm of speech.

Languages in which stress is completely predictable are called **fixed stress** systems. Fixed stress may be alternating, as in Pintupi, an Australian language where stress always falls on the first syllable and then on every other syllable after that: [**KU**ra**ɳu**lu**lim**pa**tju**ra] 'the first one (who is) our relation.' Other fixed stress systems may pick out just one syllable in the word as especially prominent. In Farsi, stress is always on the initial syllable, in Turkish always on the last syllable, in Polish always on the second-to-last syllable.

In other languages, stress is unpredictable: you just have to memorize which syllable gets stressed when you learn the word, in the same way you have to memorize whether it begins with [b] or [p]. This is called **lexical stress**. For example, Russian has a lexical stress system: if stress is placed on the first syllable, [duxi] means 'spirits'; stressed on the second syllable, it means 'perfume.'

The third type of system is **paradigmatic stress**, in which the stress patterns depend on what part of speech a word is – for example, a noun or a verb. The English system is mostly paradigmatic, with some unpredictable aspects. Generally, English verbs and adjectives follow one set of rules, nouns another. Thus English has pairs of words that differ only in stress, where one word is a noun and the other a verb: we re**JECT** the **RE**ject, re**CORD** the **RE**cord, con**VERT** the **CON**vert, in**SULT** with an **IN**sult, etc.

Acoustic phonetics

Thus far, we have talked about articulation, how speech sounds are made inside the mouth. But what happens inside the mouth is only part of the process. In order to understand how people use sound to communicate, we must also understand how the articulators turn air movements into sound, what happens to sound after it passes through the lips, how it travels through the air, and how it impacts on the ears and brain (and sometimes the microphones, recorders, and computers) of those who listen. These aspects of the linguistic study of sound fall under the heading of **acoustic phonetics**.

Sound waves

Speech sounds are moving air. Articulation is all about getting air to move in ways that can be heard – vibrating, popping, or swishing. In this short introduction, we will focus on the patterns of vibration in voiced sounds, particularly vowels. But speech sounds are very complex, so it helps to start by thinking about something much simpler – a tuning fork.

Strike a tuning fork against a table or other hard object, and the tines of the fork will vibrate. Depending on the exact shape and size of the fork, the tines will vibrate at a particular rate, or **frequency**. A fork tuned to "orchestral A," for example, will vibrate at a frequency of 440 cycles per second (cps), where a *cycle* equals one back and forth motion of the end of the tine. ("Cycles per second" are often called "Hertz" – abbreviated Hz – after a famous physicist, Heinrich Hertz.)

Different objects have different inherent frequencies of vibration, which determine the pitch of the sound. A small handbell will vibrate faster and make a higher pitched sound than a huge churchbell, which will vibrate slower and make a lower pitched sound. Human beings can hear frequencies as low as about 20 Hz, and as high as about 20,000 Hz. Humans will not perceive vibrations faster or slower than that as sound, although dogs can hear frequencies up to 45,000 Hz and bats and dolphins over 100,000 Hz.

How does the vibration of the tuning fork make it to our ears as sound? As the ends of the tuning fork vibrate, they set the air particles next to them vibrating as well, following the same back and forth motion. These moving air particles alternately push and pull on the particles next to them, and those on the particles next to them, and so on, so that the pattern of vibration moves outward from the tuning fork like ripples in a pond. These moving patterns of vibration are called **sound waves**. When the sound waves reach our ears (or a microphone), they set the eardrum (or the membrane in the microphone) vibrating according to the same pattern. Inside the ear, the vibrations set off nerve impulses, which are interpreted by our brains as sound.

Simple and complex sounds

The motion of a tuning fork is very simple (back and forth, like a pendulum), and thus the sound wave it creates is also very simple. The sound such a simple vibration makes is a pure tone of a single frequency (like 440 Hz). But the vibrations of the vocal tract – and the sounds it makes – are more complicated.

The vocal tract is like a clarinet. The vocal folds are the reed, the source of the vibration. The column of air in the mouth is the column of air in the instrument, the filter of the vibration. And the speaker, of course, is the musician, changing the shape of the column of air to modulate the sound produced. Thinking about speech sounds in this way is called the **source-filter theory** of speech production.

This is how it works. As air passes out of the trachea and over the vocal folds, if the speaker holds them in the correct position, the folds begin to vibrate. They flap open and closed at a frequency between about 100 times per second (for a large adult male) and 300 times per second (for a child). But on top of this basic flapping motion, there are many different subripples in the moving vocal folds. Think of sheets flapping in the breeze on a laundry line; they don't flap stiffly up and down, but ripple in complicated patterns. Each of these little ripples contributes its own pattern of vibration to the sound, creating "overtones," or **harmonics**, in addition to the basic pitch of the speaker's voice. The basic rate of vibration, the **fundamental frequency**, determines the pitch, but the overtones create the different qualities of different sounds.

As the vocal folds vibrate, they start the air inside the vocal tract vibrating in the same complex way. The air in the vocal tract then filters the harmonic structure of the sound produced at the vocal folds. Certain ripples (that is, certain harmonics) are amplified, and other ripples (other

harmonics) are damped out. *Which* harmonics are amplified depends on the shape of column(s) of air inside the vocal tract. Recall that different objects have different characteristic patterns of vibration, depending on their size and shape (handbells vs. churchbells, for instance). What is true of bells is also true of enclosed bodies of air. Differently shaped bodies of air will tend to vibrate at different frequencies. Harmonics that are "in tune" with the characteristic frequencies of a particular vocal tract shape will be amplified, those that are not in tune will be reduced. The speaker controls the filter by moving the tongue and lips to different positions, amplifying some harmonics and blocking out others. The most strongly amplified frequencies are called **formants**. Different vowel sounds have different formant structures.

So, depending on the shape of the tongue and lips, each vowel sound has a characteristic, complex pattern of vibration. The vibration moves out past the lips, and propagates into the world. The sound waves travel through the air at the rate of about 340 meters per second, until they impinge on a membrane tuned to receive them, such as the eardrum.

Hearing

The ear has three parts – the outer, middle, and inner ear. The outer ear consists of the visible shell of the ear (the **pinna**) and the ear canal leading down to the eardrum. The pinna helps to capture sounds better, and to locate sounds in space. The ear canal protects the eardrum, and also helps to amplify the sounds that are relevant for speech. Sound waves travel down the air in the ear canal until they reach the eardrum, which begins to vibrate.

Behind the eardrum, in the middle ear, are the three smallest bones in the body (the **ossicles**). The patterns of vibration are transferred from the eardrum through the bones of the middle ear to the inner ear, the **cochlea**. This additional stage of transfer, through the middle ear, helps to amplify very soft sounds, and tone down very loud sounds. It is in the cochlea that hearing actually takes place. The cochlea looks like a curled-up snail shell, smaller than a marble. The cochlea is divided into upper

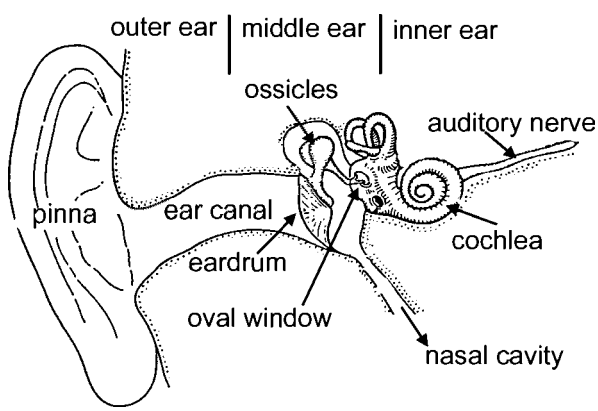

FIGURE 1.7
The hearing mechanism

and lower chambers, which are filled with fluid. Crucially, the membrane separating the two chambers, the cochlear membrane, is not uniform along its length, but is tapered. It is 3 millimeters thick and bony at one end and thin (just 1 millimeter) and flexible at the other. Embedded all along the membrane are tiny hair cells, the **cilia**, each attached to a nerve ending, and waving gently in the cochlear fluid.

As the eardrum vibrates, the bones of the middle ear vibrate, which in turn cause the membrane of the inner ear (the **oval window**) to vibrate, starting up waves of vibration in the fluid of the inner ear. The patterns of vibration in the cochlear fluid mirror the complex pattern of sound created in the vocal tract of the speaker – the fundamental frequency of the vibrating vocal folds as well as the harmonic characteristics of the different vowels.

Again recall that objects of different size and shape tend to vibrate at different frequencies. Because the cochlear membrane varies in shape along its length (thick at one end, thin at the other) different places along the membrane respond to different frequencies of vibration. The thick end vibrates in tune to low-pitched sounds, the thin end in tune to higher pitched sounds, and parts of the membrane in between vibrate in tune to mid-range sounds. In response to a given pattern of vibration, cilia at different places along the cochlear membrane are activated, sending signals to the brain about the frequencies present in the incoming wave. The brain recombines the frequency information it perceives into the sounds of language.

Measuring speech

Until quite recently, the phonetician's ears were the only instruments available for the analysis of sound, and they are probably still the most finely tuned and reliable. Beginning in the 1900s, however, phoneticians began to discover ways to make sound visible. Early devices, such as oscilloscopes and sound spectrographs, used a microphone to transfer patterns of vibration in the air into patterns of variation in electrical current. These variations could then be displayed – on paper or a screen – or passed through banks of capacitors and resistors for further measurement and study.

In the first part of the twenty-first century, speech analysis is done by computer. Microphones still convert the vibration of the membrane into variations in electrical current. But then computers convert a continuously varying sound wave into a series of numbers; this is called an analog-to-digital (A-to-D) conversion. (Working in the opposite direction, getting a computer, digital video disk, or digital audio player to *make* sound, is D-to-A conversion.)

Once represented and stored in a digital format, sound files can be mathematically analyzed to separate out the different frequencies, and the results of the analysis displayed onscreen in various formats. Three common and useful displays are shown in Figure 1.8. All three show the utterance "A phoneme?" (These figures were created on a laptop computer, using software that was downloaded free from the internet.)

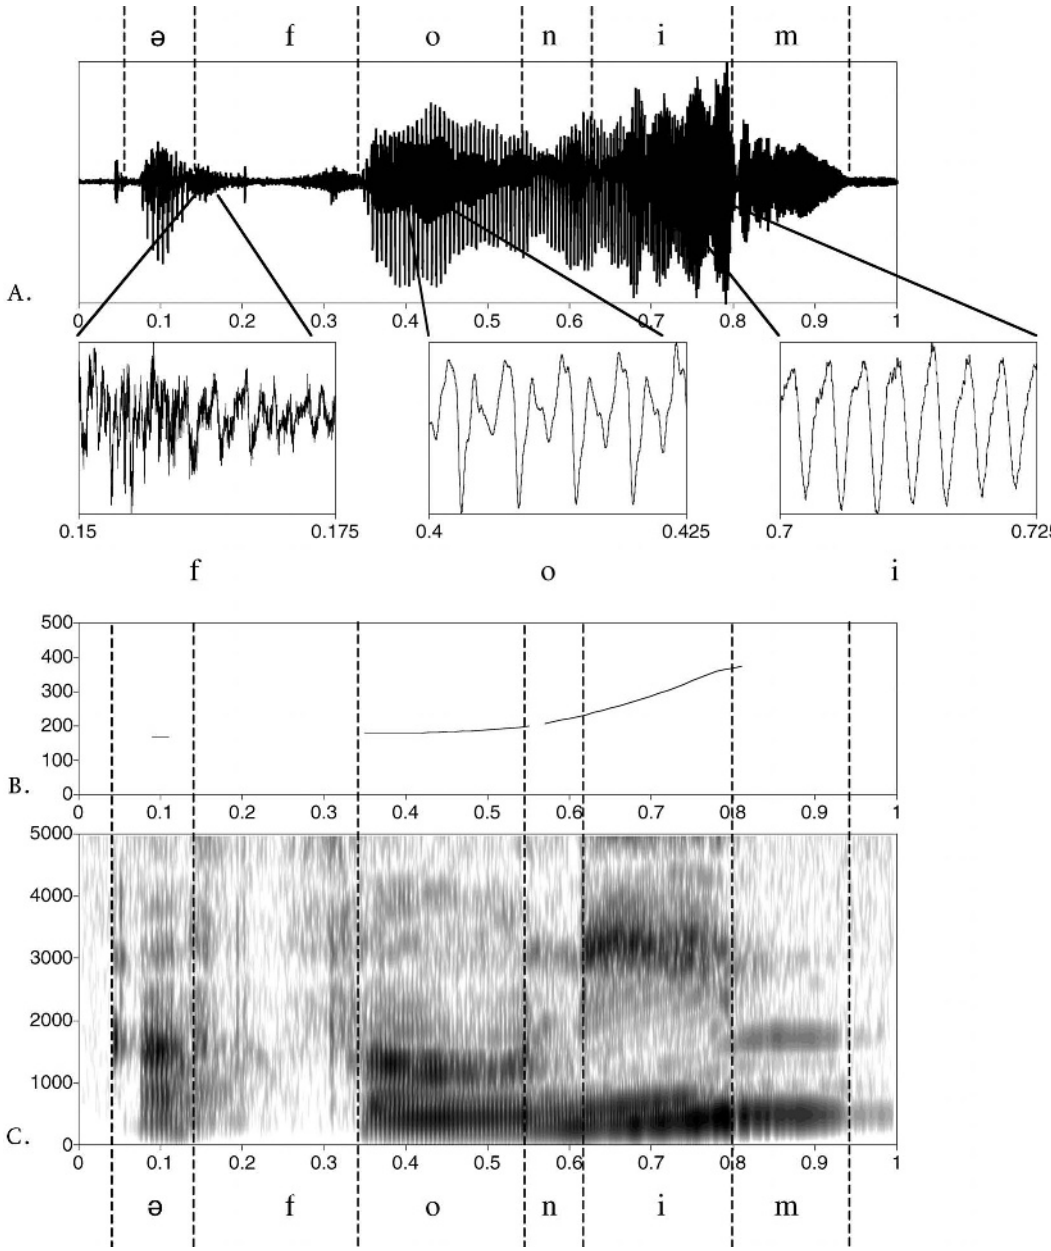

FIGURE 1.8
Waveforms for the utterance *A phoneme*?

Figure 1.8A shows a **waveform**; the varying amplitude (corresponding to loudness) of the vibrations (on the y-axis) is plotted over time (on the x-axis). The duration of the utterance is one second. Above the waveform are the segments (roughly) corresponding to each part of the waveform. Note that the vowels [o] and [i] have the greatest amplitude, the obstruent [f] the

least, and the nasals and unstressed [ə] have an intermediate amplitude. The figure also shows a close-up view (magnified 12 times) of portions of the waveform. At this magnification, you can see the random noisiness of the fricative, and the complex, repeating pattern of vibration that is characteristic of each vowel. The patterns for [o] and [i] are different, and thus the vowel qualities are different. The waveform is a useful representation for seeing the overall structure of an utterance, identifying the different types of sounds (stop, vowel, nasal consonant), and measuring the durations of different aspects of the speech signal.

Because of all the communicative information carried by pitch changes, linguists often want to measure fundamental frequency over the course of an utterance. Figure 1.8B shows the result of such a computation, a **pitch track**. In this representation, the y-axis shows frequency, and time is on the x-axis. Notice that no frequency is calculated during the [f]. Since there is no vocal fold vibration during this segment, there is no frequency to measure. This pitch track shows that this speaker's voice has a baseline fundamental frequency of about 200 Hz, but that this frequency rises toward the end of the utterance, in a contour typical of a question.

Neither of these figures, however, has told us much about the quality of the vowel sounds. We can see from the waveforms that the vowel qualities in [o] and [i] are different, because the patterns look different, but the complex squiggles don't tell us much about the component overtones and thus about the vocal tract shape that made them. The computer can further analyze the sound wave to tease apart its component frequencies. The result of this analysis is a spectrogram, shown in Figure 1.8C.

As with a pitch track, the x-axis represents time, and the y-axis shows frequency. But instead of a single line graph, we see a complicated pattern of the many frequencies present in each sound. We want our analysis to tell us which frequencies are the loudest, because once we know which frequencies have been amplified, we can use known formulas to figure out the (approximate) shape of the vocal tract that made them. This crucial third dimension, amplitude, is represented by the darkness of the lines. A dark bar at a certain frequency means that that frequency is strongly represented in the sound. Each vowel has a pattern of two or three most prominent frequencies, which are called **formants**, above the fundamental frequency of the speaker's vocal folds. For example, during the [o] vowel, we see a formant at about 1200 Hz. The [i] vowel has a different formant structure, with the second formant at about 3000 Hz.

Because the basic relationships between vocal tract shape and formant structure are known, we can infer the position of the tongue, which we can't see, from the spectrogram, which technology has made visible for us. Across speakers, the general relationships between speech sounds remain the same: for example, the second formant is always higher in [i] than in [o]. But because every person's vocal tract size and shape is unique, every person's formant structure is unique too. We recognize familiar voices, regardless of what they're saying, and in the hands of an expert, a spectrographic voice print is as unique as a fingerprint.

Phonology

What else could there be to say about the sounds of speech? Quite a lot. Language isn't just in the mouth or ears, but also in the brain. When we turn from analyzing the physical aspects of speech sounds to studying their cognitive organization, we move from phonetics to phonology. Phonology can never be completely divorced from phonetics, since sound patterns can never be completely separated from how they are produced and heard, and production and perception are always influenced by the overarching linguistic organization.

All human beings have basically the same structures in their vocal tracts and in their ears. So why are languages so different? To some extent, it is because they use different sounds from the repertoire of possible human vocal tract noises. Arabic uses pharyngeal and uvular fricatives, while English does not. French selects front round vowels (such as [y] and [œ]), English selects lax high vowels ([ɪ] and [ʊ]), and Spanish sticks to basic [i, e, a, o, u]. Thus, to some extent, learning to speak a new language is about learning to make new sounds.

But there's more to it than that. Languages differ not only in the sounds they use, but in how they organize those sounds into patterns. Consider, for example, the voiced obstruents of English and Spanish.

Phonemes and allophones

Both English and Spanish have the sounds [d] and [ð]. For example, English has *den* [dɛn] and *then* [ðɛn], while Spanish has [dama] 'lady' and [laðo] 'side'. Both sounds occur in both languages, but they differ in their *distributions* – that is, where the sounds appear – and in the information the difference between them conveys.

In General American English, there are several pairs of words that differ only in that one has [d] where the other has [ð]: *den* [dɛn] and *then* [ðɛn], *eider* [aɪdər] and *either* [aɪðər], *bade* [bed] and *bathe* [beð]. Pairs of words that differ in only a single sound in the same position within the word are called **minimal pairs**. (If you pronounce *either* as [iðər], then you have a near-minimal pair – two words that are almost exactly alike.) The existence of minimal pairs means that the difference between the two sounds is **contrastive**; change one sound into another and you've created a contrast in meaning (that is, you've made a different word). If two sounds are contrastive, their distribution is *unpredictable*. If I tell you I'm thinking of an English word that rhymes with *when*, and starts with either [d] or [ð], you cannot predict which sound (or word) I have in mind. Look back at Tables 1.1 and 1.3, which give many examples of minimal and near-minimal pairs for different sounds in English.

Now compare this to the situation in Spanish:

dama	'lady'	laðo	'side'
demonias	'demons'	universiðað	'university'
dulse	'sweet'	imbaliðo	'invalid'
disfras	'disguise'	grenaða	'Grenada'

Notice that [d] and [ð] have a different distribution in Spanish. There are no minimal pairs – no [ðama] contrasting with [dama], no [lado] contrasting with [laðo]. The difference between [d] and [ð] is not contrastive; the [d] versus [ð] difference is never used in Spanish to signal a difference in meaning. Instead, the two sounds have different distributions: only [d] is found in initial position, and only [ð] is found between vowels. The distribution is *predictable*. If you know the context (the position in the word, or the surrounding sounds), you can predict whether [d] or [ð] will be used.

When the occurrence of two different sounds is predictable based on the context, we say that the two sounds are in **complementary distribution**. (Not "com*pli*mentary" in the sense of "nice," but "com*ple*mentary" in the mathematical sense that one half of a circle is the complement of the other half.) The sounds [d] and [ð] are in complementary distribution in Spanish. (The situation is slightly more complicated when contexts other than word-initial and between vowels are considered, but the principle of complementary distribution holds. There is only [d] after [n], for instance, and only [ð] after [r].)

When two sounds in a language form minimal pairs (that is, if their distribution is unpredictable and contrastive), those two sounds represent different **phonemes**. When two sounds in a language are in complementary distribution (that is, their distribution is predictable and noncontrastive), the two sounds are **allophones** of the same phoneme. In English, [d] and [ð] represent different phonemes. In Spanish [d] and [ð] are allophones of the same phoneme.

Another way to say this is that a phoneme is a label for a group of sounds that are perceived by the speaker to be the "same" sound, and the allophones are the different ways of pronouncing that sound depending upon the context in which it is produced. To the Spanish speaker, if he/she pays any conscious attention at all, [ð] is just "a way of saying *d*." A speaker of the language knows when the [d] allophone is called for and when the [ð] allophone is appropriate. If you're learning Spanish as a second language, you may have been taught this distribution as a rule, something like "the sound [d] is pronounced as [ð] between vowels." There is no such rule relating [d] and [ð] in English. They are separate phonemes.

We may diagram the situation as follows:

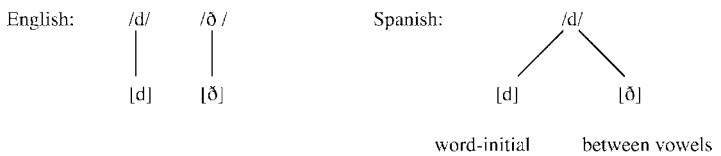

Phonemes are indicated by slashes, while allophones are indicated by brackets. At the allophonic level, English and Spanish have the same sounds. At the phonemic level, English has a contrast where Spanish has none.

Differences in phonemic and allophonic distribution pose significant problems for speakers of one language who are learning to speak another. A native speaker of Spanish learning English will have trouble with the distinction between *den* and *then*. To him/her, [d] and [ð] count as the same

sound, so he/she will tend to hear them that way, and to pronounce them according to the principles of his/her own language. A Spanish speaker may tend to say [dɛn] for *then* (using the word-initial allophone), and [æðer] for *adder* (using the intervocalic allophone). These are not random errors, but a result of imposing the phonological organization of the first language onto the words of the second.

The first thing a phonologist working to describe a previously unknown language wants to figure out is its inventory of sounds: what sounds does the language use? But the second thing he/she wants to figure out is which sound differences the language uses to encode differences between words: what are the phonemes? Next, a phonologist tries to figure out the different allophones of each phoneme by identifying predictable patterns of complementary distribution. Answering questions about contrast and predictability of sounds in a language is the main work of phonology.

Box 1.2 Another look at phonemes and allophones

The example above shows a case where a distinction that is contrastive in English is not contrastive in another language. The reverse occurs too, of course. Consider the case of voiceless stops in Thai and English.

A voiceless stop ([p], [t], or [k]) may be produced with or without an extra puff of air, called *aspiration*. In English, you can feel the aspiration if you hold your fingertips an inch or so in front of your lips as you say *pop* or *pill* . But you won't feel any aspiration for a [p] that occurs after [s] as in *spot* or *spill*. In English, stops produced with an extra puff of air are aspirated, those without the extra puff are unaspirated. Speakers of both Thai and English produce a full set of aspirated voiceless stops [p^h, t^h, k^h] and unaspirated voiceless stops [p, t, k]. Though the inventory of voiceless stops is the same, the languages use the inventory in different ways. Some Thai and English words using these sounds are shown below.

Thai:	Aspirated:		Unaspirated:	
	p^hàt	'to stir fry'	pàt	'to wipe'
	t^hun	'a fund'	tun	'to hoard'
	k^hâ:w	'step'	kâ:w	'rice'
English:	Aspirated:		Unaspirated:	
	p^hɪl	pill	spɪl	spill
	t^hɪl	till	stɪl	still
	k^hɪl	kill	skɪl	skill

In Thai, aspiration is contrastive. The difference between [p] and [p^h] makes a difference in meaning, and thus minimal pairs with aspirated and unaspirated stops (such as [p^hàt] and [pàt]) are easy to find. [p] and [p^h] are two different phonemes, as are [t] and [t^h], and [k] and [k^h].

There are, however, no minimal pairs for aspiration in English. Aspiration is never the *only* difference between two words. Instead, we can *predict* whether a stop in an English word will be aspirated or unaspirated depending on its context. At the beginning of a word, a voiceless stop will be aspirated; after /s/, it will be unaspirated. In English, the single phoneme /p/ has two allophones: [p] and [pʰ].

The allophonic nature of aspiration in English is reflected in the orthography. English doesn't have two different written symbols for these two sounds (for [p] and [pʰ], for example), whereas Thai does. An English speaker needs a lot of convincing to believe that the *p* sounds in *sport* and *port* are not phonetically identical. To him/her, they count as the same, because they are allophones of a single phoneme. The Thai speaker needs no such convincing. To him/her, [p] and [pʰ] are as different as [t] and [k].

Phonotactics

Languages do not allow random sequences of sounds; rather, the sound sequences a language allows are a systematic and predictable part of its structure. Languages have **phonotactic constraints** – restrictions on the types of sounds that are allowed to occur next to each other or in particular positions in the word. In English, for example, no word begins with the sequence [tl]. There are words like *train* and *plane*, but no *tlane*. This sequence is perfectly pronounceable and occurs in other languages (e.g. Navajo [tlee] 'night'). The sequence can occur in English if the two sounds are in different syllables, as in *At.lan.tic.* But English has a phonotactic constraint: **tl* ("*t-l* is illicit") at the beginning of a syllable – and therefore at the beginning of a word, as well. (The asterisk is used to indicate any sequence that is not possible in a given language.)

Another phonotactic constraint of English limits the sounds that are allowed to follow the diphthong [aʊ]. We have many words like *out, crowd*, *town*, *mouse*, *couch*, *south*, and *rouse*. But there are no words **aup*, **auk*, **aub*, **awm*, **aug*, etc. If you make a list of all the [aʊ] words you can think of, you will discover that the only sounds allowed to follow [aʊ] are those made with the tongue tip and blade: [t], [d], [s], [z], [n], [θ], [ð], [tʃ], [dʒ] (and possibly [l] and [r], if you say *owl* and *hour* as one syllable). So one way of stating this English phonotactic constraint would be "The diphthong [aʊ] can only be followed by [t], [d], [s], [z], [n], [θ], [ð], [tʃ], [dʒ]." But we can do better. We want to capture the *generalization* that these sounds are not a random list, but all share some property. A better way to state this constraint is "The diphthong [aʊ] can only be followed by a consonant made with the tongue front." Phonological constraints seldom (if ever) target random collections of sounds. Rather, they almost always target groups of sounds that have one or more phonetic properties in common. Such groups are called **natural classes**.

One of the most common phonotactic constraints across languages is "Nasals must agree in place of articulation with a following stop." Think of English words that have a sequence of nasal + oral stop: *camp*, *hamper*, *bombard*, *paint*, *intelligent*, *wind*, *window*. Though there are some exceptions, the general pattern is for the bilabial stops [p] and [b] to be preceded by the bilabial nasal [m] and for the alveolar stops [t] and [d] to be preceded by the alveolar nasal [n]. What about words with [k] and [g]? We spell these words *nk* and *ng*, but if you listen carefully to your pronunciation of words like *think* and *linguistics*, you'll realize that these sequences are really [ŋk] and [ŋg] (the velar nasal [ŋ] preceding the velar stops [k] and [g]), and thus also follow the constraint. (Many words whose spelling ends in *ng*, like *sing* and *rang*, actually end in just the nasal [ŋ], with no [g] stop at all.)

There are many different phonetic dimensions that can define natural classes. Voicing is often important in phonotactic constraints. A common constraint is that sequences of obstruents must be either all voiced or all voiceless. In Russian, there are words like [**vzbutʃki**] 'scolding' and [**fsp**leska] 'splash,' but not *[**fzputʃgi**]. In English, we have [siksθs], but not *[siksθ**z**]. In native Japanese words (borrowings from Chinese or other languages are exempt), there can be only one voiced obstruent per word. So there are words like [saki] 'wine,' [kaze] 'wind,' and [zaseki] 'seat' (each with one or more voiceless obstruents), but there is no *[gaze] or *[guda].

Phonotactic constraints are often related to syllable structure. Many languages put constraints on the types of syllables they will allow. All languages allow syllables of the CV (Consonant–Vowel) type – that is, with an onset and a nucleus, but without a coda. This is the least marked and most common type of syllable. Some languages, like the West African language Senufo, allow only CV syllables. The preference for CV syllables results from two phonotactic constraints: "Syllables must have onsets" (a C before the V), and "Syllables must not have codas" (no C after the V).

Some languages have only one of these constraints. Hawai'ian, for example, does not allow codas, but does allow syllables that don't have onsets. Therefore, it allows syllables of the type V and CV, but never CVC. Hawai'an is also famous for having one of the smallest segmental inventories around – only five vowels [i, e, a, o, u] and eight consonants [p, k, ʔ, l, w, n, m, h]. That gives a grand total of 45 possible syllables (5 V and 40 CV combinations) to encode the thousands of words in the language. How do you use 45 syllables to make thousands of words? Hawai'ian words can be very, very long, with lots of repetition! The state fish, for example, is called the [hu.mu.hu.mu.nu.ku.nu.ku.a.pu.a.ʔa]. Note that every syllable is either CV or V, and that [ʔ] functions as a regular consonant.

Other languages allow CV and CVC but not VC. Every syllable in these languages must have an onset. Arabic and German are languages of this type. Words that are spelled with initial vowels in writing are actually pronounced with an initial glottal stop – for example, [ʔapfɛl] 'apple' in German, and [ʔal] 'the' in Arabic.

Many languages that allow syllables to have codas (CVC and VC) limit the sounds that can occur there. These restrictions are called **coda constraints**. In Thai, [p], [pʰ], and [b] may occur in onsets, but only [p] may

appear in the coda. In Japanese, the only coda consonants allowed are nasals or the first half of geminates: [nip.pon] 'Japan,' [ʃim.bun] 'newspaper,' [gak.ko] 'school.' Constraints on onsets exist (e.g., no [ŋ] in onsets in English), but are much rarer. Finally, many languages allow at most one consonant in the coda and one in the onset. English is rather rare in allowing CCVCC, or even CCCVCCC words, like *shrink* or *sprints*. Russian words like [**fsp**leska] (four consonants in a row) are rarer still.

Box 1.3 Phonotactics and borrowing

Phonotactic constraints often give rise to interesting forms when a more constrained language borrows a word from a less constrained language. Japan, for example, has borrowed the game of baseball from America, and many of the terms that go with it. But English baseball terms must be adapted to the syllable constraints of the Japanese language. A common strategy is to insert extra vowels to break up consonant clusters and turn codas into onsets. So *bat* becomes [bat.to], *strike* becomes [su.to.rai.ku], and *baseball* becomes [be.su.bu.ro]. Hawai'ian has the double challenge of a restricted inventory of sounds and limited syllable structure possibilities. The phrase *Merry Christmas* becomes [meli kalikimaki] in Hawai'ian. It looks unrecognizable, but you can easily work it out: change [r] into [l], [s] into [k], and add vowels as needed. Note that even English speakers don't pronounce the [t].

Of course, English does the same with words it borrows from other languages. We pronounce *futon* as [futon] not [ɸuton], and the country of Qatar (pronounced, in Arabic, just like it's spelled [qatar]) as *cutter*. In those rare cases where we borrow a word from a language that allows even more syllable types than we do, the English strategy tends to be to drop the offending consonants rather than insert extra vowels. Greek, for example, allows words to begin with [pn] and [ps] as in [pneu] 'breath' and [psaike] 'spirit.' We preserve the Greek origin of words like *pneumonia* and *psychology* in the spelling, but not the pronunciation.

Alternation and allomorphs

Phonology is about contrast and predictability. We want to know what sounds in a language are contrastive phonemes, and we want to be able to predict the contexts in which the different allophones of a phoneme appear. Sometimes we do this by studying lists of words in order to discover different kinds of sound distributions. Sometimes we look at what happens to borrowed words. Another way is by studying **alternations**.

An alternation is seen when the same **morpheme** is pronounced different ways depending on the context. A morpheme is any entry in your mental dictionary; it can be a word or a part of a word, like a prefix (added to the beginning of a word) or a suffix (added to the end of a word). For example, in English, *cat* is a morpheme, and so is the plural suffix *-s*. As we

put different morphemes together to create larger words, we create different contexts. Depending on the context, different forms of the morpheme, different **allomorphs**, are pronounced. The same phonotactic constraints that disallow certain sequences of sounds *within* morphemes also give rise to alternations when the combination of morphemes creates new environments that violate phonotactic constraints.

We saw that, within morphemes in English, nasals agree with the place of articulation of a following stop: remember *camp*, *wind*, and *think*. Because these nasals are within a single morpheme, we never see alternations in how they are pronounced. But some English prefixes that end with a nasal are pronounced differently when the prefix is added to words beginning with different sounds. Consider the negative prefix *in-*. Before a vowel or an alveolar, *in-* is pronounced [ɪn]: *inedible*, *inaudible*, *indelible*, *intolerable*, *insufferable*. But before the bilabials [m], [b], or [p], it becomes [ɪm]: *immobile*, *imbalance*, *impossible*. And before [k] or [g], it's pronounced [ɪŋ]. We don't change the spelling, because English doesn't have a separate letter for the velar nasal, but listen to your pronunciation of a word like *incorrect* and you'll realize the nasal is velar rather than alveolar. The phonotactic constraint "Nasals must agree in place of articulation with a following stop" applies not just within morphemes (like *camp* and *wind*) but also when morphemes are combined.

Natural classes help us to define the set of sounds that is targeted for alternation, the change itself, and the environments where the change takes place. This can be clearly seen in the case of the Spanish voiced stops and fricatives described above. We noted there that /d/ has two allophones in Spanish: [d] in initial position, and [ð] between vowels. But the change is not limited to /d/ alone. It also affects /b/ and /g/. Whenever a voiced stop [b, d, g] is preceded by a vowel, it changes to a voiced fricative [β, ð, ɣ].

barselona	'Barcelona'	a βarselona	'to Barcelona'
bola	'ball'	la βola	'the ball'
data	'date'	la ðata	'the date'
gata	'female cat'	la ɣata	'the female cat'

The set of sounds that undergoes change is a natural class: voiced stops. They all undergo the same change; they become voiced fricatives. And the environment in which the change occurs can be defined by a single phonetic parameter: after sounds that have an open vocal tract and continuous flow of air, like vowels, [s], and [r].

Types of phonological alternations

There are certain types of phonological processes that we find over and over again across languages. (See Chapter 8 for more examples of the following alternations in a diachronic (across time) perspective.)

Assimilation

The most common type of alternation is **assimilation**: two sounds that are different become more alike. Assimilation may be either *local*, when the

two sounds are next to each other, or *long-distance*, where two sounds seem to affect each other even through other intervening segments.

Voicing assimilation is one very common alternation across languages. In English, the plural suffix agrees in voicing with a previous stop: *two cats* [kæ**ts**] and *two dogs* [da**gz**]. Another example is seen in Russian. The Russian word for 'from' is either [ot] or [od]: [o**d vzb**útʃki] 'from a scolding,' [o**t fspl**éska] 'from a splash.'

Local assimilation may also affect place of articulation. The English negative prefix *in-* is an example of a nasal assimilating in place of articulation to a following stop (*indecent* vs. *impossible*). Another example comes from Twi, a language of West Africa.

me-pɛ	'I like'	me-m-pɛ	'I do not like'
me-tɔ	'I buy'	me-n-tɔ	'I do not buy'
me-ka	'I say'	me-ŋ-ka	'I do not say'

The morpheme that means 'not' in Twi has three different allomorphs. It is always produced as a nasal consonant, but the nasal takes on the place of articulation of whatever consonant it's next to – bilabial next to bilabial, velar next to velar, etc.

The most extreme type of local assimilation is complete assimilation: two sounds that are next to each other become identical. Complete assimilation of adjacent vowels is found in many West African languages. In Yoruba, for example [ow**o**] 'money' plus [**e**po] 'oil' becomes [ow**ee**po] 'oil money.' In Igbo, [nwok**e**] 'man' plus [**a**] 'determiner' becomes [nwok**aa**] 'that man.' The English prefix /ɪn/ can also undergo complete assimilation in certain words. Before most consonants, the nasal assimilates just to the place of articulation, as in *impossible*. But before [l] or [r], the assimilation of the nasal is total: *in* + *regular* becomes *irregular*, *in* + *legal* becomes *illegal*.

Assimilation can also take place long-distance. Two segments that are not immediately adjacent may come to share some phonetic property. **Vowel harmony** is the prime example of long-distance assimilation. For example, in Turkish, the vowels in a word must be all front or all back, and the high vowels must be all round or all unround. The suffix meaning (roughly) 'of' therefore has different allomorphs, depending on the quality of the vowel in the preceding syllable: [ip-**in**] 'of the rope,' [pul-**un**] 'of the stamp.' In harmony systems, the vowels assimilate to each other, even though consonants intervene. Consonant harmonies also exist, though they are rarer.

Dissimilation

The opposite of assimilation is **dissimilation**. Two sounds that are similar become different. One impetus for dissimilation may be ease of articulation. Two sounds that are similar but not exactly the same seem to be particularly difficult to pronounce correctly right next to each other. (This principle forms the basis of many tongue twisters, such as the *sixth sheik's sixth sheep*.) One solution to the problem would be to make the similar sounds identical; another solution is to make them more different. An example of dissimilation is found in the history of Greek. In Ancient

Greek, the word for *school* was [sxolio], with two adjacent voiceless fricatives. In Modern Greek, this has become [skolio], with a fricative–stop sequence instead. Dissimilation may also help hearers to realize that two segments are present, not just one. In Setswana (spoken in Botswana), voiced stops become voiceless after a (voiced) nasal consonant, perhaps to make the presence of the stop more obvious: [bona] 'see' becomes [mpona] 'see me.'

Insertion

Another common type of alternation is insertion (which phonologists call **epenthesis**). Insertion is usually related to syllable structure, as we saw above, when vowels are inserted to break up strings of consonants. This can happen in loan words (as in Japanese and Hawai'ian) or when morphemes come together. In English, when we want to add the plural suffix [z] to a word that already ends in [s] or [z], we insert a vowel [ɨ] to break up the two high-pitched fricatives: one dress [drɛs], two dresses [drɛsɨz].

Deletion

The opposite of insertion, of course, is **deletion**. Instead of breaking up a sequence of consonants with a vowel, a language may choose to delete one of the consonants (as in the loss of the initial [p] in *pneumonia*). *Grand* ends with [nd], *mother* starts with [m], but when the two words are put together, the medial [d] is usually deleted: [grænmʌðər]. (The [n] may then assimilate to the [m], resulting in [græmmʌðər].) Lardil (a language spoken in Australia) deletes final vowels from words of three or more syllables: [yalulu] 'flame' becomes [yalul]. But sometimes vowel deletion leaves sequences that violate the phonotactic constraints of the language, and further surgery is called for. Lardil words may end in at most one consonant, and that consonant must be produced with the tongue tip. So [tʃumputʃumpu] (dragonfly) becomes not *[tʃumputʃump], with final [mp], but [tʃumputʃu], with the final three segments deleted.

Lenition and fortition

Another type of change is **lenition** – sounds become softer or weaker. Stops change to fricatives, fricatives change to approximants. The Spanish alternation [b, d, g] becomes [β, ð, ɣ] is an example of lenition. The opposite of lenition is **fortition**. Here, fricatives change into stops, as in Kikuyu (East Africa) "postnasal hardening": [ɣora] 'buy,' [ŋgoreetɛ] 'I have bought.'

Metathesis and reduplication

Other less common alternations include **metathesis** and **reduplication**. Metathesis means switching the order of sounds. For example, the English word *horse* used to be [hros], before the [ro] sequence was metathesized. Reduplication means copying. In English, we sometimes copy parts of words to convey a pejorative, diminutive sense: *teeny-tiny, itsy-bitsy*, or the more dismissive *syntax-schmintax*. But in other languages it's a regular part of the sound system. In Manam (Austronesian), [gara] means 'scrape' and

[gara-gara-si] means 'scraping it.' You may also recall Hawai'ian [humuhumunukunukuapua?a] and Lardil [tʃumputʃumpu].

Phonological theory

The preceding sections have discussed some contrasts and alternations that are common in human languages. But phonologists want to know more than this. They want a generalization, not just a list. Phonologists don't want to know just "What is the inventory of sounds in Polish?" but "What is a possible inventory in any language?" They want to know not just "What alternations are common across languages?" but "What alternations are possible in any language?" They want to know not just "How are Russian and Ukrainian different?" but "How different can languages be?

We have seen that languages choose different phonetic dimensions (different vocal tract gestures) to encode their contrasts. Voicing, aspiration, manner, and place of articulation can all be used contrastively. Similarly, we have seen these same dimensions used to define sets of sounds that are relevant for phonological alternations. **Distinctive feature theory** aims to encode all the phonetic dimensions that languages have available to encode contrasts and natural classes.

The linguist Roman Jakobson (1896–1982) proposed that each relevant dimension could be thought of as a plus or minus contrast. The speaker chooses whether a sound will be [+voice] (with vocal fold vibration) or [−voice] (without vocal fold vibration); [+nasal] (open velum) or [−nasal] (closed velum); [+sonorant] (a sonorant sound, without airstream obstruction) or [−sonorant] (an obstruent sound, with airstream obstruction); [−continuant] (a stop) or [+continuant] (not a stop), etc. Jakobson proposed both an acoustic and an articulatory definition for each of his features.

Every phoneme could be defined in terms of a set of distinctive features: [m], for example, would be [+labial, +nasal]. Features could also be used to define natural classes, and the changes that affect them. The class of voiced stops in Spanish would be [−sonorant, −continuant, +voice], and the change from stop to fricative would be [−continuant] becomes [+continuant].

By proposing a fixed, finite set of universal features, Jakobson attempted to define all the phonetic dimensions that could be phonologically relevant – that is, that could be used for contrasts and alternations. A possible human language would use these, and only these, features. Many of the features Jakobson proposed are still in use today; others have been replaced or refined.

Phonologists are also concerned with describing the relationship between phonemes (the **underlying representation**, or UR – the way words are stored in the brain) and allophones (the **surface representation**, or SR – the way words are actually pronounced). Noam Chomsky and Morris Halle in their influential 1968 book, *The Sound Pattern of English*,

proposed that allophones are derived from phonemes by the application of **phonological rules**. These rules took the form

$X \rightarrow Y / A \_\_ B$

where X is the class of sounds affected, Y is the change, and A__B is the context, all written in terms of distinctive features. The rule above can be read "X becomes Y in the context of between A and B" or "AXB becomes AYB." Other special characters were used for word and morpheme boundaries: #___ means "word-initial position," and ___# means "word-final position." Can you read the following two phonological rules, which describe alternations discussed above?

Rule 1: [−continuant, −voice] → [+aspirated] / # ___
Rule 2: [−continuant, −sonorant, +voice] → [+continuant] / [+continuant] ___

The first rule states that a voiceless stop becomes aspirated when it is produced at the beginning of a word. This describes word-initial aspiration of [p], [t], and [k] in English. The second rule states that a voiced stop becomes continuant (that is, a fricative) when it follows another nonstop sound. This describes how a voiced stop in Spanish [b, d, g] changes to voiced fricatives [β, ð, ɣ] whenever it is preceded by a vowel.

In a **derivation**, rules change the UR into the SR. Recall the discussion, above, of the pronunciation of the word *grandmother*. The basic, underlying, representation has a sequence of three consonants [ndm] in the middle. First, this sequence is simplified, by deletion of the [d]. Then, the [n] assimilates in place of articulation to the following [m], to which it is now adjacent. The result is that the medial consonants are in fact pronounced [mm]. A derivation shows this series of changes. Thus, the derivation of [græmmʌðər] looks like this:

Underlying representation:	grændmʌðər
Rule 1, deletion:	grænmʌðər
Rule 2, assimilation:	græmmʌðər
Surface representation:	græmmʌðər

Phonologists have proposed different ways of writing rules. For example, it might make more sense to think of assimilation as feature *sharing* rather than feature changing. One might write a rule of voicing assimilation by diagramming a single feature being shared from one consonant to another:

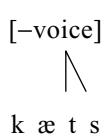

This way of writing rules is called **autosegmental representation**, because the features are treated as autonomous from the string of segments.

More recently, phonologists have questioned whether phonological rules are the best way to think of the relationship between underlying and surface representations. Rules can describe alternations, but they don't capture the fact that alternations tend to take place for a reason – to bring words into conformity with phonotactic constraints. Recent work in phonology has emphasized the importance of constraints in mediating between underlying and surface representations. Such work proposes that all languages share the same set of constraints, but rank them differently. A given form may violate a lower-ranked constraint in order to meet a more important one. For instance, the pronunciation of /grændmʌðər/ as [græmmʌðər] indicates that English places a higher priority on the constraint that says "Don't have three consonants in a row" than on the constraint that says "every segment in the UR must appear in the SR."

Chapter summary

The goal of this chapter has been to describe the sounds of speech, from the point of view of a phonetician and of a phonologist. Phoneticians study the physical aspects of linguistic sounds: Movements of the structures of the vocal tract, place and manner of articulation, the propagation of sound waves through the air, hearing and speech perception, computer measurement of fundamental frequency and formant structure. Phonologists study the more abstract organization of sound patterns: Syllable structure, phonotactic constraints, alternations, the relationship between underlying and surface representations.

We began this chapter by posing questions that phonologists and phoneticians ask, and have attempted to survey some of the preliminary answers that have been proposed. Phonologists and phoneticians have learned a lot about how speech sounds are made, how they are perceived, and how they are organized. But many questions remain.

- What new sounds and sound patterns remain to be discovered?
- How can we best (and most quickly) describe and preserve the sounds and sound patterns of the diverse languages that are dying out?
- As we gain more and more knowledge of how our mouths and ears work (from more and more sophisticated measuring devices), how can we incorporate this knowledge into our acoustic and articulatory models?
- As we gain more and more knowledge of how our brains work, how can we incorporate this knowledge into our phonological models?
- How do cognitive patterns and articulatory events influence and constrain each other?
- What is the right set of distinctive features? Are phonological features and articulatory gestures one and the same?

- Is the relationship between underlying and surface representation rule-based or constraint-based?
- How can we account for language change and language variation?
- How can we better understand how children acquire the phonology and phonetics of their native language?
- How can we better understand the process of learning a non-native language, and help those who are struggling to do so?

The studies that will answer these questions remain to be written.

Exercises

Exercises 1.1

Draw a mid-sagittal diagram of the vocal tract, using Figure 1.2 as a model. In your diagram:

– shade the larynx yellow
– shade the pharynx green
– shade the oral cavity red
– shade the nasal cavity blue

Then, label each of the active and passive articulators:

– trace the active articulators in red
– trace the passive articulators in blue

Exercise 1.2

Fill in the blanks to describe these consonants of English:

	voicing	nasality	place of articulation	manner of articulation	sonorant/ obstruent
m	voiced	nasal	bilabial	stop	sonorant
s	voiceless	oral	alveolar	fricative	obstruent
g					
ʃ					
n					
d					
f					
ð					

Exercise 1.3

Fill in the blanks to describe the vowel symbols:

	height	front/back	rounding	tense/lax
ɪ	high	front	unround	lax
i				
ɛ				
ɔ				
æ				
o				
ʊ				

Exercise 1.4

Transcribe the following words in IPA:

transcribe
yellow
computing
movie
phonology
housetop
beans
climb
flimsy
books
thoughtful
annoyed
joyous
than
choice
baseball
caught
code
spring break
sunrise

Exercise 1.5

Write out the following passage (from Sapir 1933) in English orthography. Note that the symbol [ɾ] stands for an alveolar "tap," an allophone of American English *t* and *d* found in *pretty* ['pɹɪɾɪ] and *later* ['leɪɾəɹ]. The symbol ['] precedes the stressed syllable.

ðə 'kansɛpt əv ðə 'fonim, ə 'fʌnkʃənəli sɪg'nɪfikənt 'junɪt ɪn ðə 'ɹɪdʒɪdli də'faɪnd 'paɹərn ɔr kʌnfɪgjə'ɹeʃən əv saʊndz pə'kjulər tu ə 'læŋgwɪdʒ, æz dɪ'stɪnkt frəm ðæt əv ðə saʊnd ɔr fə'nɛɾɪk 'ɛləmɛnt æz sʌtʃ, æn əb'dʒɛktɪvli də'faɪnəbəl 'ɛntɪɾi ɪn ði aɹ'tɪkjuleɪɨd ænd pɛɹ'sivd to'tælɪɾi əv spitʃ, ɪz bi'kʌmɪŋ mɔr ænd mɔr fə'mɪljər tu 'lɪŋgwists. ðə 'dɪfɪkəlti ðæt 'mɛni stɪl sim tə fil ɪn dɪs'tɪŋgwɪʃɪŋ bə'twin ðə tu mʌst ə'vɛntʃuəli dɪsə'pir æz ðə rilə'zeʃən groz ðæt no 'ɛntɪɾi ɪn 'hjumən ɛk'spiɹiəns kæn bi 'æɹəkwɪtli də'faɪnd æz ðe mə'kænɪkəl sʌm ɔr 'pradəkt əv ɪts 'fɪzɪkəl 'prapəɹtiz.

Exercise 1.6

Transcribe the following words, and draw a syllable tree for each. Explain how sonority is relevant in assigning sequences of consonants to the onset or coda.

arthritis
handbag
complained
linguist

Exercise1.7

Make a list of American English words that contain the sequence [ju], such as *beautiful* ['bjuɾɪfəl] and *music* ['mjuzɪk]. What natural class of consonants is prohibited from preceding [ju]?

Exercise 1.8

Consider the distribution of voiced and voiceless vowels in Japanese (data from Tsuchida 1997). Are [i] and [i̥] representatives of two distinct phonemes, or are they allophones of a single phoneme? What about [u] and [u̥]? Argue for your answer, either by citing (near-)minimal pairs from the data, or by describing the distributions of two sounds.

all voiced vowels:

kokoro	'heart'
sensei	'teacher'
kakko	'parenthesis'
suberu	'slip'
buka	'subordinate'
kidoo	'orbit'
omoʃiroi	'interesting'
banira	'vanilla'
jume	'dream'

some voiceless vowels:

su̥peru	'spell'
ɸu̥ka	'incubation'
ki̥too	'prayer'
kosu̥to	'cost'
çi̥too	'hit'
ʃu̥to	'capital'
ki̥ppari	'clearly'
ɸu̥kaɸu̥ka	'soft'
su̥tasu̥ta	'quickly'
koku̥sai	'international'
su̥ʃimai	'sushi rice'

Exercise 1.9

Consider the distribution of [l] and [lʲ] (a palatalized version of [l], with a raised and fronted tongue body position) in Russian. Do they represent two different phonemes, or are they allophones of a single phoneme? Argue for your answer, either by citing (near-)minimal pairs from the data, or by describing the distributions of two sounds (data courtesy of Maria Gouskova).

words with [l]:

lat	'agreement'
gala	'gala'
polka	'shelf'
mel	'chalk'
pol	'heat'

words with [lʲ]:

milʲ	'of miles'
nebolʲ	'imaginary tale'
lʲat	'demon'
lʲot	'ice'
molʲ	'moth'
polʲka	'polka'

Exercise 1.10

Consider the distribution of [k] (voiceless velar stop) and [x] (voiceless velar fricative) in Florentine Italian (data from Villafana 2005). Do they represent two different phonemes, or allophones of the same phoneme? Argue for your answer, either by citing (near-)minimal pairs from the data, or by describing the distributions of two sounds.

laxasa	'the house'
poxo	'little'
bixa	'stack'
amixo	'friend'
fixi	'figs'
kwuoxo	'cook'
kwando	'when'
kapella	'chapel'
blaŋko	'white'
makkina	'machine'
kabina	'booth'

Exercise 1.11

Consider the following alternations that occur in English. Describe each as assimilation, dissimilation, lenition, fortition, epenthesis, or deletion. For each, list three more words that undergo the alternation.

a. Vowels are nasalized before a nasal consonant, in words such as *camper* [kæ̃mpər], *wrong* [rɔ̃ŋ], and *tone* [tõn].
b. In some dialects, [ð] and [θ] are pronounced [d] and [t], in words such as *then* [dɛn], *thin* [tɪn], and *mouth* [mɔt].
c. In many dialects, [r] is not pronounced when in follows a vowel, as in *car* [ka], *park* [pak], and *sure* [ʃuə].
d. Words with a sequence of nasal plus fricative are sometimes pronounced with a stop between nasal and fricative: *tense* [tɛnts], *something* [sʌmpθɪŋ], *sense* [sɛnts].
e. The prefix *con-*, meaning 'with,' has three different allomorphs: *conduct* [kəndʌkt], *complain* [kəmplen], *congress* [kaŋgrɛs].

Suggestions for further reading

Johnson, Keith 2003, *Acoustic and auditory phonetics*, 2nd edition, Oxford and Malden, MA: Blackwell. A very readable introduction to the physics of sound, speech perception, and computer speech processing.

Kenstowicz, Michael 1994, *Generative phonology*, Oxford and Malden, MA: Blackwell. A more advanced and thorough introduction to phonology, with extensive data sets.

Ladd, D. R. 2000, *Intonational phonology*, Cambridge University Press. Introduction to the study of intonation from a linguistic perspective.

Ladefoged, Peter 2001, *A course in phonetics*, 4th edition, Oxford and Malden, MA: Blackwell. A practical guide to phonetics, emphasizing practice in articulation and transcription.

Ladefoged, Peter and Maddieson, Ian 1996, *The sounds of the world's languages,* Oxford and Malden, MA: Blackwell. A survey and detailed description of the articulation and acoustics of all the sounds in the IPA chart.

Spencer, Andrew 1996, *Phonology*, Oxford and Malden, MA: Blackwell. A very accessible introductory text.

Yip, Moira 2002, *Tone*, Cambridge University Press. A comprehensive overview of tonal phonology.

CHAPTER

2 Words and their parts

DONNA LARDIERE

CHAPTER PREVIEW

KEY TERMS

ablaut
affix
agreement
allomorph
aspect
base
case
compound
derivation
feature
gender
infix
inflection
lexeme
lexicon
mood
morpheme
morphology
number
paradigm
person
prefix
reduplication
root
stem
suffix
suppletion
tense
word
zero derivation

This chapter introduces the subject of morphology, the study of the internal structure of words and their meaningful parts. Morphological processes fulfill two basic purposes: (1) to create new words in a language and (2) to modify existing words. We may associate a word with a certain basic idea, image or event, but modifying the exact form of a word can also contribute important information, such as who is participating in an event, when or how it occurred, or something about the speaker's attitude toward it. The more complex the word, the more information of this sort it is likely to convey. By manipulating various parts of a word, we can shade, intensify, or even negate its basic meaning, or change its grammatical role within a sentence. Different languages, of course, have different ways of doing this.

GOALS The goals of this chapter are to:

- introduce key concepts in the study of complex word analysis
- provide a concise description of some of the varied morphological phenomena found among the world's languages
- illustrate methods used to derive and support linguistic generalizations about word structure in particular languages
- touch briefly on how knowledge of complex word forms comes to be acquired

What is a word?

Imagine you were in an environment where everyone around you was speaking a language you'd never heard before, and you couldn't understand a single word of what they were saying. That typical phrase – "couldn't understand a single word" – underscores our intuition that words are the fundamental building blocks of language. The foremost task of any language learner, including young children acquiring their native language, is to figure out how to segment and analyze the wall of talking-noise around them into meaningful units – namely, words and their meaningful parts.

But what is a word, exactly? *Webster's Unabridged Dictionary* (1989) defines a *word* as the smallest independent unit of language, or one that can be separated from other such units in an utterance. In the following conversational exchange, (1b) demonstrates the independence of the word *tea*.

(1) a. Which do you like better – coffee or *tea*?
b. *Tea*.

Words can enter into grammatical constructions, such as phrases and sentences. For example, the word *tea* can be used in different positions in a sentence according to its grammatical role:

(2) a. *Tea* is good for you.
b. She doesn't drink *tea*.
c. There are beneficial antioxidants in *tea*.

Tea is the subject of the sentence in (2a), the direct object in (2b), and the object of a preposition in (2c).

Our definition from *Webster's* continues: words are "usually separated by spaces in writing and distinguished phonologically, as by accent" (p. 1643). But this is only partially accurate. Although spaces are placed between words in the written form of many languages (like English), orthography (the written form of a language) cannot be a crucial component of wordhood. There are languages like Chinese which don't insert spaces between words in writing, but speakers of these languages still know what a word is in their language. Similarly, people who can't read

and speakers of languages without writing systems know what words are in their languages, too.

On the other hand, phonology does play an important role across languages in identifying the boundaries between words. For example, consider the string /grinhaʊs/. (Recall from Chapter 1 that symbols between forward slashes represent how something is pronounced, using the phonemes of a language.) Phonological stress disambiguates the meaning of the utterances in (3a) and (3b), indicating that /grinhaʊs/ is a single (compound) word in (3a) but two distinct words in (3b):

(3) a. They walked past a *GREENhouse.*
    b. They walked past a *green HOUSE.*

Phonology can help us identify words, but we need other information as well. Consider the following:

(4) a. *Tea's* good for you.
    b. That shop sells *teas* from around the world.
    c. I asked him not to *tease* the cat.

Is *tea's* in (4a) one word or two? The sound form of *tea's* is phonetically identical to that of *teas* in (4b) and even *tease* in (4c); all are pronounced /tiz/. But your intuition is probably that the wordhood status of *tea's* is somehow grammatically different from that of *teas* or *tease*. There is an additional element in *tea's* which, although phonologically dependent on *tea* (as a contracted form of the word *is*), is nonetheless a distinct *grammatical* word.

Webster's also states that words are "typically thought of as representing an indivisible concept, action, or feeling, or as having a single referent." Clearly, the word *tease* in (4c) has a different referent than *teas* in (4b). But the word *teas* also means something a bit different than the simple word *tea* – something like 'more than one kind of tea.' This difference in meaning is conveyed by the ending *-s* (pronounced [z]) on the word *tea*. But this *-s* ending is not an independent word; rather, it must be attached directly to an independent word whose basic meaning it is modifying – in this case, to indicate *plural* meaning. We can conclude that even though *teas* is just one word, the *-s* ending is a distinct subpart that contributes some piece of additional information to the word's overall meaning.

It appears that we require a fairly complex definition of *word*, defining it in relation to meaning, grammar, and phonology. For now, let us more simply define a **word** (a surprisingly difficult term in linguistics) as an abstract sign that is the smallest grammatically independent unit of language.

All languages have words, but the particular sign a language uses to express a particular meaning is *arbitrary*. For example, there's nothing inherent in the sound form of the word *water* that actually carries the meaning of 'water.' French speakers refer to the very same stuff as *eau*, Japanese speakers call it *mizu*, and Italians *acqua*. The fingerspelled form of 'water' that was signed onto Helen Keller's outstretched palm as water flowed over the other enabled her to "break into" the system of words as

abstract signs. Later, she would also learn how the meaning of 'water' was represented in another kind of abstract form – the system of raised dots known as Braille:

(5)

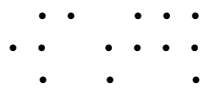

'water'

The human impulse to discover and create words, it seems, transcends even profound differences in physical capabilities.

The words of one's language make up its **lexicon**. One might think of the lexicon as a kind of mental dictionary where words are stored. Our knowledge of each word, like the lexical entries in a dictionary, includes several kinds of information. Consider what you know, for example, about the word *sleep*:

- how it is pronounced: /slip/
- what it means – informally, something like to repose or rest in the body's natural periodic unconscious state. Your knowledge of the meaning of *sleep* also includes the information that only animate objects – like babies, cats, and students (but not trees or ideas) – can get sleepy.
- the grammatical contexts in which the word can be used. *Sleep* is an intransitive verb (it doesn't take a direct object), as in the sentence *Sally sleeps late on weekends*. But it can also be a noun as in *John talks in his sleep*. It can be found in compound words such as *sleepwalking* and *sleep-deprived* and in idioms such as *to let sleeping dogs lie*.
- that it is an **irregular** verb for past-tense marking in English, requiring that we memorize its past form *slept* /slɛpt/ instead of simply adding the regular past marker to produce **sleeped* /slipt/.

When you stop to consider for a moment all the (tens of thousands of) words that are in your lexicon and everything you already know about each of them, you can begin to appreciate the magnitude of the accomplishment of this impressive feat. Moreover, new items are continually being added, just as dictionaries are continually revised and updated (e.g. *beer goggles*, *DVD-player*). The meanings of the listed words might also change over time, or acquire (or lose) different shades of meaning (e.g. *dude*, *gay*).

However, the contemporary study of word formation is not as much about the study of existing, listed dictionary words as it is the study of *possible* words in one's language and the mental rules for constructing and understanding them. Not all of the words you can produce and interpret are listed in the lexicon, because the number of possible words is infinite. For example, a recent quick look through a single magazine turned up the following words:

(6) outgeneraled
scrounginess
on-messagism
unanswerability
extraterritorialization
hyperparenting
transhumanists
balconied

In that same issue there were also many freely coined compound-word expressions, including the following:

(7) thwack-time interval
poultry-litter composting
receipt-management strategy
cringe-making
floppy-haired
cultural studies semiotics junkies
snowy-headed
puzzled-chimp expression

It's possible that one of these newly created words will "stick" in your lexicon – perhaps popping up again someplace else as more people adopt it or maybe because you just like it. Most of these words, however, are destined to be immediately forgotten, but even though they are ephemeral, they demonstrate the human capacity to mentally represent the complex structure of words in one's language.

To further illustrate what you know about words, let's consider a word you're not likely to know (because I've made it up): *frimp*. If you heard it in the context of an English sentence such as *John likes to frimp on weekends*, then you would deduce that it's a verb that can be used intransitively (that is, without a direct object). And once you knew that, then even before learning its exact meaning (which would depend on the context and your knowledge that it's an "action"), you would already know how to construct several other word forms based on this verb. You'd know how to use its past form (*he frimped all day yesterday*) and progressive form (*he was in the kitchen frimping when I called*). You'd also know how to turn it into an adjective (*I wish he'd mend his frimping ways*). You would know to look up (or list) *frimp*, not *frimped* or *frimping*, as the "dictionary" form, because you'd assume that *frimp* is a **regular** verb. Since the *-ed* and *-ing* endings can attach to *all* regular verbs, the forms *frimped* and *frimping* don't really need to be listed in the lexical entry for *frimp*. You'd also know that John was a *frimper*. As you can see, you already know quite a lot about this hypothetical word!

Each language has its own rules and processes for creating new words, and these words are interpretable in their contexts even if they are never recorded in a dictionary. The forms of words may be simple or extremely complex; our knowledge of the mental rules and categories that enable us to produce and interpret them makes up the subject of **morphology**.

Morphology: the study of word structure

The branch of linguistics that is concerned with the relation between meaning and form, within words and between words, is known as **morphology**. Morphology literally means 'the study of form' – in particular, the forms of words. Although "form" in this context usually refers to the spoken sound

or *phonological form* that is associated with a particular meaning, it doesn't necessarily have to – signed languages also have word forms. Instead of the articulators of the vocal tract, signed languages make use of the shape and movement of the hands. All languages, whether spoken or signed, have word forms.

Morphologists describe the constituent parts of words, what they mean, and how they may (and may not) be combined in the world's languages. The pairing of a meaning with a form applies to whole words, like *sleep*, as well as to parts of words like the 'past' meaning associated with the ending *-ed* as in *frimped*.

Morphology applies within words, as in the addition of a plural ending to *cat* /kæt/ to change its form to *cats* /kæts/ and its meaning to 'more than one cat.' It also applies across words, as when we alter the form of one word so that some part of it matches, or *agrees* with, some feature of another word, as shown in (8):

(8) a. That cat sleep*s* all day.
b. Those cat*s* sleep all day.

In the sentence in (8a), the word *cat* is a **third-person singular** (3SG) subject, which in most varieties of English requires that we add an *-s* to another word – the verb – when they occur together in a sentence. This verbal suffix "means" something like 'my subject is third person and singular.' In (8b), however, the word *cats* is plural, which in English doesn't require the verb to add any special agreeing form. (English is highly unusual among the world's languages in this regard!) In the examples above, notice that the words *that* and *those* also crossreference the singular vs. plural meaning distinction between *cat* and *cats*. This kind of morphological **agreement** between matching parts of words is widely observed among the world's languages.

Languages vary widely in their amount and functions of morphology (often as a result of historical development – see Chapter 8). For example, all languages need a way to signal grammatical roles such as subject and direct object (or, who did what to whom). English depends quite strictly on the order of words in a sentence to do this. The meaning of (9a) is very different from that of (9b):

(9) a. Brutus killed Caesar.
b. Caesar killed Brutus.

Latin, however, marks grammatical roles *morphologically*, and word order is consequently much freer; all the Latin sentences below mean 'Brutus killed Caesar':

(10) *Latin* (Bauer 2003: 63)
Brūtus Caesar*em* occīdit.
Caesar*em* occīdit Brūtus.
Occīdit Caesar*em* Brūtus.
'Brutus killed Caesar.'

In Latin, the addition of *-em* to the noun *Caesar* indicates that Caesar is the direct object, or the one who got killed. (The subject, Brūtus, in this case is unmarked.) Here the morphological form of a noun, rather than its position in the sentence, signals its grammatical function. One of the most important functions of morphology is to distinguish the roles played by the various participants in an event; we could not interpret language without this information.

In the remainder of this chapter, we will examine some other functions of morphology, and also take a closer look at some key terms and concepts that linguists use to describe the processes of morphology.

Morphemes

We said earlier that *tea* and *teas* are both words with slightly different meanings, and that this difference is due to the *-s* ending on *teas*. But since *-s* is not itself a word, how can it have its own meaning? In fact, it is not words, but rather **morphemes**, that are the smallest units of language that combine both a form (the way they sound) and a meaning (what they mean). Words are made up of morphemes. Simple words consist of a single morpheme. Complex words consist of more than one morpheme. For example, *cat* is a simple word compared with *cats*, which contains two morphemes – the noun *cat* plus a plural marker *-s*. Similarly, in the word *unfriendly*, there are three morphemes: *un-*, *friend*, and *-ly*, each of which contributes some meaning to the overall word. Some words in morphologically rich languages can contain so many morphemes that we need an entire complex sentence in English to translate them. Consider the following complex word from Turkish, which contains the lexical root *Avrupa-* 'Europe' plus eleven additional morphemes (don't worry for now about the function of each morpheme as glossed below the word):

(11) *Turkish* (Beard 1995: 56)

Avrupalılaştırılamıyacaklardansınız
Avrupa-lı-laş-tır-ıl-a-mı-yacak-lar-dan-sın-ız
Europe-an-ize-CAUSE-PASSIVE-POTENTIAL-NEG-FUT.PART-PL-ABL-2ND-PL
'You (all) are among those who will not be able to be caused to become like Europeans.'

However, even in English, we can come up with a reasonably complex word that can more compactly express the approximate meaning of the latter part of the translation 'not able to be caused to become like Europeans':

(12) You are among those who will be *unEuropeanizable*.

The English word *unEuropeanizable* consists of the root *Europe* plus the morphemes *un-*, *-(i)an*, *-ize*, and *-able*. Each of these morphemes contributes to the overall meaning of the entire word.

In building words (and phrases and sentences), two basic kinds of morphemes are used. Morphemes with richer lexical "vocabulary" meaning

When a language has a morpheme with a grammatical meaning, we gloss it in small capital letters in the interlinear translation, in the same order in which the morphemes appear. So, in (11), tır means 'causative,' ıl means 'passive,' a means 'potential,' and so on. When a gloss appears with a dot within it, it means the morpheme includes both meanings, hence, yacak means 'future participle' in Turkish. Example (11) has two morphemes glossed PL*, for "plural" – the first, lar, is the plural for the future participle, the second, ız, is the plural for second person.*

(referring to things and qualities and actions in the world) are called lexical morphemes or **lexemes**. Lexemes typically belong to the "major" part-of-speech categories of nouns (N), verbs (V), or adjectives (A); simple lexemes may serve as the **root** of more complex words. On the other hand, morphemes that contribute mainly grammatical information or indicate relationships between the lexemes are called **grammatical morphemes**. In the sentence in (13) below, the words *maniacal*, *little*, *dog*, *attempt*, *bite*, and *mailman* are all lexemes. The grammatical morphemes, which have been underlined, are *their*, *-al*, *-ed*, *to*, and *the*.

(13) Their maniacal little dog attempted to bite the mailman.

Grammatical morphemes are the glue that holds the lexemes in a sentence together, shows their relations to each other, and also helps identify referents within a particular conversational context. In the sentence in (13), the pronoun *their* consists of grammatical features (third person, plural) that partially identify via agreement some previously mentioned referents (say, *John and Mary*), and simultaneously signals a possessor relation between them and the lexeme *dog*. Like the morpheme *the* (in the phrase *the mailman*), *their* also has definite reference, indicating that the speaker assumes the hearer knows who is being referred to. The past tense marker *-ed* tells us that an event (the 'biting-attempt') already happened. The morpheme *to* is a formal device (called an **infinitive** marker) for marking an untensed verb (*to bite*). Finally, although the word *maniacal* is a lexeme (because it is an adjective that refers to an attribute of the dog, just like *little*), it is a complex word derived by adding the morpheme *-al* to the root noun *maniac*. Thus, *-al* has the grammatical function of turning a noun into an adjective meaning 'having the qualities of' that noun.

Both lexemes and grammatical morphemes can be either **free** or **bound**. Bound morphemes must be attached either to a root or another morpheme, but free morphemes can stand alone. Most lexemes in English, such as *dog* and *bite*, are free morphemes. Suffixes, like *-ed* and *-al*, are bound. In many other languages, however, lexical roots are not free morphemes; they must be bound with other morphemes to yield a grammatical word. In Italian, for example, the root of the verb *lavor-* 'work' *must* be bound with grammatical morphemes such as tense and agreement markers:

(14) *Italian*

Lavor-ano	a	casa.
work-PRES.3PL	at	home

'They work at home.'

A morpheme performing a particular grammatical function may be free in one language and bound in another. For example, the English infinitive marker *to* (as in the verb phrase *to win the election)* is a free morpheme. It can be separated from its verb by one or more intervening words (despite what prescriptive grammar books say!), as in *to very narrowly win*

the election. In French, however, the verb *gagner* 'to win' consists of the root *gagn-* and the infinitive morpheme *-er*, which are tightly bound together in a single word and cannot be split up. Conversely, the regular past-tense marker *-ed* in English must be tightly bound to a verb, as in *attempt-ed*, *walk-ed*, or *call-ed*, but in Koranko (a West African language), the morpheme used to encode past tense is a free morpheme, an independent word:

(15) *Koranko* (Kastenholz 1987: 109, cited in Julien 2002: 112)

à	*yá*	kɔ̀lɔmabolɔ	kári
3SG	PAST	tree-branch	break

'he/she broke a branch'

We can see that the past morpheme *yá* in Koranko is free because it can be separated from its verb (*kári* 'break') by an intervening direct object (*kɔ̀lɔmabolɔ* 'tree branch').

Box 2.1 Identifying morphemes

Can you identify all the morphemes in the following English sentence?

The musicians reconsidered their director's unusual proposal.

Let's go word by word:

- *the*: a grammatical morpheme indicating that the referent of the following noun is definite (not just any musicians) and known to both the hearer and the speaker
- *musicians:*
  - the root lexeme *music*
  - the morpheme *-ian* indicates a person who works in some capacity connected to the meaning of the root
  - the plural marker *-s*, meaning 'more than one'
- *reconsidered:*
  - the root lexeme *consider*
  - the morpheme *re-*, meaning 'again'
  - the past-tense marker *-ed*
- *their*: a grammatical morpheme indicating possession of the following noun by some plural third persons
- *director's:*
  - the root lexeme *direct*
  - the morpheme *-or*, denoting someone who performs the action of the verb
  - the morpheme *-s*, indicating possession of something by the noun to which it is attached
- *unusual:*
  - the root lexeme *usual*
  - the morpheme *un-*, meaning 'not'
- *proposal:*
  - the root lexeme *propose*
  - the morpheme *-al*, turning the root verb into a noun

The forms of morphemes

We observed earlier that morphemes combine both a form and a meaning. However, sometimes the exact form of a morpheme systematically varies under certain conditions, much like the way in which phonemes can be pronounced as different allophones depending upon the context in which they are produced, as observed in Chapter 1. And in fact, one of the most common factors influencing the forms morphemes take is phonology, or more precisely, some aspect of the local phonological environment. For example, in English orthography, regular nouns are marked for plural by adding *s* (or in some cases *es*) but the actual *sound* of the plural morpheme *s* varies between [s], [z], and [ɨz]. Consider the following pluralized English words:

(16)

[z]		[s]		[ɨz]	
peas	[pʰi:z]	puffs	[pʰʌfs]	peaches	[pʰitʃɨz]
charms	[tʃarmz]	charts	[tʃarts]	charges	[tʃardʒɨz]
mills	[mɪlz]	myths	[mɪθs]	misses	[mɪsɨz]
caves	[kʰevz]	cakes	[kʰeks]	cases	[kʰesɨz]
flags	[flægz]	flaps	[flæps]	flashes	[flæʃɨz]
plays	[pʰlez]	plates	[pʰlets]	phrases	[frezɨz]

The examples in (16) show that there are three possible forms of the plural suffix for regular nouns in English. These regular plural morpheme variants are in complimentary distribution and are called **allomorphs**. In other words, the particular regular plural form is completely predictable depending on, or phonologically conditioned by, the final sound of the base. The [-ɨz] form follows a class of fricatives called **sibilants** (/s, z, ʃ, ʒ, tʃ, dʒ/). The [-s] form follows all the other voiceless consonants (/p, f, θ, t, k/). The [-z] form follows all the voiced sounds (all vowels and all the voiced consonants). Each of these three plural forms makes it easier to hear the plural marking following different root word endings.

Although phonological factors are often responsible for allomorphic variation, allomorphy may also be conditioned by factors other than phonology. Many languages, for example, have different verb classes (called **conjugations**), which condition the form of affixes such as agreement markers. In Italian, the verbs *lavorare* 'to work,' *scrivere* 'to write,' and *dormire* 'to sleep' belong to three different conjugation classes (sometimes known as the *-are*, *-ere*, and *-ire* classes respectively, based on their infinitive ('to') forms, which are also allomorphs). The suffix *-o* attaches to a verb root in any class and means 'agreement with a first-person, singular subject' (1SG), as shown below:

(17) *Italian*

lavor-*o*	scriv-*o*	dorm-*o*
work-1SG	write-1SG	sleep-1SG
'I work'	'I write'	'I sleep'

For some other agreement categories, however, such as 'second person, plural' (2PL), verbs in these conjugation classes require different forms of the agreement marker:

lavor-*ate*	scriv-*ete*	dorm-*ite*
work-2PL	write-2PL	sleep-2PL
'You (pl.) work'	'You (pl.) write'	'You (pl.) sleep'

This is an example of *morphologically* conditioned allomorphy, since conjugation classes are formal morphological categories.

Finally, semantic factors may also play a role in determining how morphemes can be realized. The English prefix *un-* (meaning 'not') can readily attach to adjectives in the first column in (18) but not the second (Katamba 1993: 79). Why not?

(18)		
	unwell	*unill
	unloved	*unhated
	unhappy	*unsad
	unwise	*unfoolish
	unclean	*undirty

If we wanted to express the negated form of *ill*, *hated*, *sad*, *foolish*, and *dirty*, we'd simply have to use the free-morpheme variant *not* instead of *un-*. In a pair of words representing opposite poles of a semantic contrast (such as *happy* and *sad*), the positive value (*happy*) is usually the **unmarked** (or more neutral or normal) quality, from which the more **marked** negative value can be derived by adding the affix *un-*. Lexemes already containing the negative value (*sad*) often cannot take a negative affix (**unsad*).

Box 2.2 Practice with conditioned allomorphy

The Turkish data below exhibit allomorphy. Can you identify the allomorphs and determine what is conditioning them?

Turkish

adamlar	'men'	günler	'days'
anneler	'mothers'	ipler	'threads'
atlar	'horses'	jıllar	'years'
aylar	'months'	kalemler	'pencils'
bankalar	'banks'	kediler	'cats'
başlar	'heads'	kitaplar	'books'
camiler	'mosques'	kızlar	'girls'
çocuklar	'children'	masalar	'tables'
dersler	'lessons'	mevsimler	'seasons'
dişçiler	'dentists'	oteller	'hotels'
eller	'hands'	sonlar	'ends'
elmalar	'apples'	umutlar	'hopes'
gözler	'eyes'	üzümler	'grapes'

From the English translations, you can see that every Turkish word above has a plural meaning. Can you identify a likely plural morpheme for each word? (In a true acquisition situation, you would hear – and be able to contrast – each of these words in both plural and singular

contexts, such as *elma* 'apple' vs. *elmalar* 'apples,' and *kedi* 'cat' vs. *kediler* 'cats.') In some cases, the plural form is *-lar* whereas in the others it's *-ler*. Is this difference predictable or is it random? In other words, if you were learning Turkish and came across a new singular noun that you hadn't heard before, would you know how to pluralize it?

First, it doesn't look as if the meaning of the word conditions the morpheme's form. For example, looking at body parts, animals, and buildings, we find *baş-lar* 'heads' but *el-ler* 'hands'; *at-lar* 'horses' but *kedi-ler* 'cats'; *banka-lar* 'banks' but *otel-ler* 'hotels', and so on. Gender doesn't look promising either; *adam-lar* 'men' is male and *anne-ler* 'mothers' is female, but *kız-lar* 'girls' is also female. (In fact, Turkish does not distinguish grammatical gender.) The final sound of the lexical base might determine which form to use (as it does in English plurals), but compare, for example, *ip-ler* vs. *kitap-lar*, both of which end in /p/, or *gün-ler* vs. *son-lar*, both of which end in /n/. Bases ending in a vowel also contrast, as in *cami-ler* vs. *masa-lar*.

Now look at the last vowel in every root. Making a list of which ones co-occur with each plural form, we get the following distribution:

a, u, ɪ (= /ɨ/), o	→	*-lar*
e, i, ö, ü	→	*-ler*

The roots whose final vowels have the feature [+front] are followed by the plural suffix with a [+front] vowel (*-ler*), whereas roots with [−front] final vowels are followed by the plural suffix with a [−front] vowel (*-lar*). This process is called **vowel harmony**, and since it involves (phonological) feature matching, you might think of it as a kind of phonological "agreement." The plural forms *-lar* and *-ler* are allomorphs in complimentary distribution, and the choice of one or the other allomorph is determined in this case by vowel harmony, a kind of phonological conditioning.

Some morphological operations of the world's languages

In this section we'll examine some of the morphological processes languages use to modify the form and meaning of lexemes:

- affixation: the addition of a discrete morpheme either before, after, inside of, or around a root or another affix
- reduplication: the copying of some part of a root
- root change: the change or replacement of some part of a root
- suprasegmental change: a shift in tone or stress to signal a grammatical function

Affixation

The most common morphological process for modifying a root is by adding something to it – the process of affixation. Most of the world's

languages use some kind of affixing to indicate grammatical information about a word or its relation to other words. An **affix** is a grammatical morpheme which (by definition) must be bound to a root or to another affix. Any form an affix attaches to, whether simple or complex, is called a **base** (or a **stem**). Affixes which attach to the right, or end, of a base are called **suffixes**. Affixes which attach to the left, or front, of a base are called **prefixes**. The complex English word *uninterpretability*, for example, consists of the root lexeme *interpret*, a prefix *un-* and the suffixes *-able*, and *-ity*.

(19) *English*
un- interpret -able -ity → uninterpretability

Some languages, such as Turkish, are primarily suffixing. This is shown in the Turkish example we observed earlier, repeated below as (20). In this word, there are eleven suffixes following the root *Avrupa-* 'Europe':

(20) *Turkish* (Beard 1995: 56)
Avrupa-lı-laş-tır-ıl-a-mı-yacak-lar-dan-sın-ız.
Europe-an-ize-CAUSE-PASSIVE-POTENTIAL-NEG-FUT.PART-PL-ABL-2ND-PL
'You (all) are among those who will not be able to be caused to become like Europeans.'

Other languages, like Chichewa (the national language of Malawi), are mostly prefixing; the example below shows six prefixes preceding the root *phwany* 'smash':

(21) *Chichewa* (Mchombo 1998: 503)
Mkângo s-ú-na-ká-ngo-wá-phwányá maûngu.
3lion NEG-3SUBJ-PAST-go-just-6OBJ-smash 6pumpkins
'The lion did not just go smash them, the pumpkins.'

Chichewa, like other African Bantu languages, has a complex system for classifying nouns, which we'll briefly discuss in Box 2.7. The numbers '3' and '6' indicate that the noun mkângo *'lion' belongs to noun class 3 and* maûngu *'pumpkins' to class 6. The numbers glossed in the verb prefixes indicate subject agreement (with 'lion') and object agreement (with 'pumpkins') respectively.*

The simplest and most common way to build word structure is to cumulatively add suffixes or prefixes to derive a more complex word, as in the English, Turkish, and Chichewa examples above. This kind of affixation is sometimes referred to as concatenative morphology, since discrete morphemes appear linked together (or concatenated) like beads on a string.

Other types of affixation

In addition to prefixes and suffixes, some languages make use of **infixes**, a kind of affix that is inserted inside a lexical root. Infixing is less common than suffixing or prefixing across the world's languages. An example of verb root infixing from Tagalog (the national language of the Philippines) is shown in (22):

(22) *Tagalog* (Himmelmann, to appear)

tulong	'help'	t-*um*-ulong	'helped
bili	'buy'	b-*um*-ili	'bought'
hanap	'search'	h-*um*-anap	'searched'

Another unusual kind of affixation is **circumfixing**, in which a two-part or discontinuous morpheme surrounds a root. Many past participles in German are formed this way, as shown in (23):

(23) *German*

ge-kann-*t*	'known'
ge-läute-*t*	'rung'
ge-schüttel-*t*	'shaken'
ge-zeig-*t*	'shown'

Although infixing and circumfixing involve either splitting up the lexical root or splitting up the grammatical morpheme (and thus are not truly concatenative processes), there is still a one-to-one correspondence between a particular morpheme and a particular grammatical function. In Tagalog, the infix *-um-* indicates past-tense. In German the past participle is formed by circumfixing the discontinuous morpheme *ge-t* around the verb.

An even more interesting kind of infixation is found in Semitic languages such as Arabic and Hebrew. In these languages certain sequences of vowels are interspersed throughout an abstract lexical root consisting only of consonants. For example, the root *ktb* in Arabic is associated with the meaning 'write' but it is a bound root; the sequence of sounds *ktb* is never uttered by itself but must be filled in with specific vowels in a specific pattern in order to derive an actual word. The particular choice and position of vowels determines the overall meaning of the word, as shown in (24):

(24) *Arabic*

katab	'to write'
kataba	'he wrote'
kutib	'has been written'
aktub	'be writing'
kitaab	'book'
kutub	'books'
kaatib	'clerk'
maktaba	'library, bookstore'

In this kind of root-and-vowel-pattern morphology, certain grammatical categories like 'past-tense' cannot be associated with an individual affix. Instead, they're associated with an entire pattern, or template, of vowels superimposed upon a lexical root consisting of only consonants. Thus, a past-marked verb in Arabic consists of a consonant-and-vowel template that looks like *CaCaCa* (e.g. *kataba*), and the present progressive form is the *aCCuC* template (e.g. *aktub*).

Box 2.3 Arabic word formation

Try Arabic word formation using the lexical root *drs*, meaning 'study.' Based on the examples in (24), how do you say 'he studied,' 'has been studied,' 'be studying,' and 'school' in Arabic?

- The template for the past-tense form 'he studied' is *CaCaCa* → *darasa*.
- The template for the perfect passive form meaning 'has been studied' is *CuCiC* → *duris*.
- The template for the present progressive form 'be studying' is *aCCuC* → *adrus*.
- 'school' is made up of the prefix *ma-* (which means roughly 'a place') combined with the template *CCaCa* → *madrasa* 'a place for studying.'

Reduplication

Suppose you were learning Tagalog and heard the future-tense forms for the verbs *tawag* 'call' and *takbo* 'run' shown in (25):

(25) *Tagalog* (Schachter and Otanes 1972; Himmelmann, to appear)

ROOT		FUTURE	
tawag	'call'	ta-tawag	'will call'
takbo	'run'	ta-takbo	'will run'

You might guess that the future-tense morpheme in Tagalog is the prefix *ta-* 'will.' This would be a good guess on the basis of such a limited sample of data. But suppose you then heard the future-tense forms of *bisita* 'visit' and *bili* 'buy':

bisita	'visit'	bi-bisita	'will visit'
bili	'buy'	bi-bili	'will buy'

Now you have at least two possible hypotheses: either there are two (or more) distinct morphemes (*ta-* and *bi-)* for expressing the future tense, or something else a little more interesting is going on. More data should help you decide:

pasok	'enter'	pa-pasok	'will enter'
alis	'leave'	a-alis	'will leave'
dalo	'attend'	da-dalo	'will attend'
lakad	'walk'	la-lakad	'will walk'
gawa	'make'	ga-gawa	'will make'
kain	'eat'	ka-kain	'will eat'
sunod	'obey'	su-sunod	'will obey'

For each of the future forms, the first syllable of the root is copied and prefixed to the root to form the future tense. Although it looks like a case of simple prefixation in some respects, there is no single pre-specified morpheme or even a set of morphemes that we could point to as having the meaning 'future.' Rather, the future tense is derived by a morphological process that copies a subset of the phonemes of each individual root. This copying process is called **reduplication**.

In reduplication, sometimes the entire word is copied, or sometimes just part of the word (as in the Tagalog examples above), and sometimes

part of the root is copied along with some fixed or prespecified morpheme. In the following example from Ilokano (another language spoken in the Philippines), reduplication is used to derive a verb with a "pretentative" meaning from an abstract noun:

(26) *Ilokano* (Rubino 2002)

singpet	'virtue'	agin-si-singpet	'pretend to be virtuous'
baknang	'wealth'	agim-ba-baknang	'pretend to be rich'

In the 'pretend' verbal forms there is a specific prefix *agiN-* that attaches to the first CV (consonant and vowel) copied from the root. The capital *N* in the prefix indicates that the last sound in the prefix is a nasal which assimilates to the place of articulation of the first (copied) consonant of the root. So in the case of *singpet* 'virtue,' the initial *si-* of the root is copied, which in turn determines (or conditions) the final nasal consonant of *agin-* (since /s/ and /n/ share the same place of articulation), to yield the 'pretend' verbal form *aginsisingpet* 'pretend to be virtuous.' Similarly, for *baknang*, the initial *ba-* of the root is copied and conditions the form of the prefix *agim-* (since /b/ and /m/ share the same place of articulation). The Ilokano pretentative construction offers another example of how phonological processes like nasal assimilation often condition the ultimate form of morphemes.

Box 2.4 English Pig Latin

Try pronouncing the following phonetically transcribed utterance, which involves reduplication in which, for every derived word, part of an English root is copied and combined with some pre-specified material. Can you reconstruct the English roots and decipher the utterance?

Then see if you can analyze the derived words in terms of (partial) reduplication.

u-dej u-jej o-nej aʊ-hej u-tej ik-spej ɪg-pej ætɨn-lej ?

If you learned the English-based word game Pig Latin as a child, the pattern in the words above should sound familiar. How might a linguist who knew English but didn't know Pig Latin analyze this utterance? Observe that every word ends with a morpheme containing at least one consonant and the fixed sequence *-ej*, yielding the suffix *-C(C)ej*. The consonant in this suffix is different for every word, which should give us a clue that the consonant is what's being copied from the root. But the base to which the suffix is affixed in each of the first four words above doesn't even contain a consonant. In fact *every* word in the utterance begins with a vowel sound. That's a clue that the consonant copied from the root has been deleted from the initial position of the root in the derived word. Following up on our hypothesis, we get the following

English roots (in phonetic transcription) and the corresponding derived Pig Latin words:

English		*Pig Latin*	
du	'do'	u-dej	'do'
ju	'you'	u-jej	'you'
no	'know'	o-nej	'know'
haʊ	'how'	aʊ-hej	'how'
tu	'to'	u-tej	'to'
spik	'speak'	ik-spej	'speak'
pɪg	'Pig'	ɪg-pej	'Pig'
lætɨn	'Latin'	ætɨn-lej	'Latin'

Ablaut and suppletion

Another kind of morphological operation, called **ablaut**, signals a grammatical change by substituting one vowel for another in a lexical root. Consider past-tense marking in English verbs, which may be regular or irregular.

(27) *English*

Regular		Irregular	
Root	Past	Root	Past
call	called	fall	fell
glide	glided	slide	slid
like	liked	strike	struck
bake	baked	take	took
live	lived	give	gave
share	shared	swear	swore
confide	confided	ride	rode

The regular verbs are marked for past-tense by simply concatenating the root with the grammatical morpheme *-ed*. But forming the past-tense for irregular verbs is a nonconcatenative process; the past-tense is not marked by adding a prefix or suffix but by a vowel change. Moreover, for these verbs there is no clearly identifiable, prespecified morpheme that we can associate with the grammatical meaning 'past.' The particular vowel change will depend rather idiosyncratically on the particular (irregular) root lexeme, and must be memorized. In these cases, we can't say that the past-tense morpheme is just the changed vowel; rather, the whole word form (for example, *rode*) resulting from the morphological process of ablaut is the past-tense form. In these verbs, the lexical meaning and the grammatical function 'past' are more synthetically fused than in regular past-tense marking.

More radically, sometimes grammatically related forms bear very little resemblance to each other. Consider the following pairs:

(28) *English*

ROOT	PAST
catch	caught

buy	bought
think	thought
teach	taught
seek	sought

In these cases, all of the root after the initial consonant has been deleted. All the past-tense forms in (28) have the same rhyme (/-ɔt/) despite having very different lexeme roots. This is a case of partial **suppletion**, in which nearly the entire root appears to have been replaced by a completely different form, leaving only the original root onsets. The English pair *go–went* is a case of total suppletion – *went* shares nothing at all with *go*.

Tone and stress

Another nonconcatenative morphological process often used to signal a contrast in grammatical meaning is the use of tone. In Somali, one way in which some nouns can be pluralized is by shifting the high tone on the penultimate (next-to-last) syllable in the singular form onto the final syllable in the plural:

(29) *Somali* (Lecarme 2002)

SINGULAR		PLURAL	
árday	'student'	ardáy	'students'
díbi	'bull'	dibí	'bulls'
mádax	'head'	madáx	'heads'
túug	'thief'	tuúg	'thieves'

Some languages use changes in syllable stress to indicate grammatical information. For example, some English nouns have been derived from verbs simply by shifting the stress from the second to the first syllable. Compare the nouns derived by stress shift (in the second column) with the nouns derived from the same verbs via affixation (in the third column).

(30)

VERB	NOUN	NOUN
convíct	cónvict	conviction
permít	pérmit	permission
progréss	prógress	progression
rebél	rébel	rebellion
recórd	récord	recording

In this section we have looked at some of the ways in which the world's languages build and modify the structure of words, including affixation, reduplication, internal root change, and shifts in tone and stress. What purposes do all these operations serve? We turn next to that question.

Two purposes of morphology: derivation and inflection

Morphological processes are traditionally classified into two broad types, each with a rather different function. Among the major lexical

categories – nouns (N), verbs (V), and adjectives (A), **derivational** morphology creates new lexemes from existing ones, often with a change in meaning. In example (30) above, we saw that it is possible to derive two different nouns from the verb *convict* via two different operations. The first noun, *cónvict*, was derived by stress shift and denotes a person who has been convicted. The second noun, *conviction*, was derived by affixation and denotes the outcome or result of being convicted. In this case, the morphological operations of stress shift and affixation were both used for derivational purposes, since two lexemes (nouns) were created from another lexeme (a verb).

Inflectional morphology, on the other hand, adds grammatical information to a lexeme, in accordance with the particular syntactic requirements of a language. Consider the following English sentence:

(31) He plans to contact her in a few weeks.

The particular (suppletive) forms of the pronouns *he* and *her* are required by the syntactic roles they play in the sentence as subject and object, respectively; furthermore, the verb *plan* must be affixed with *-s* to agree with its third-person, singular (3SG) subject, and the noun *week* must also be affixed with plural *-s* as required by the quantifier phrase *a few*. Thus, in this example, the morphological mechanisms of suppletion and affixation were both used for inflectional purposes – to convey grammatical information. Consider how ungrammatical the result would be if these particular syntactic requirements were not met; that is, if the wrong suppletive forms of the pronouns were used or if the required affixes were not added:

(32) *Him plan to contact she in a few week.

Both derivation and inflection often co-occur within the same word, although in English there is typically only one inflectional operation per word. (There may be several derivational ones.) Consider the complex English word *dehumidifiers*. Creating this word requires three derivational operations and one inflectional operation, each subsequent step building on the base of the previous one:

(33)

humid	– an adjective, the lexical root
humidify	– step 1: a transitive verb is derived by suffixing *-ify*, meaning to 'cause something to become humid'
dehumidify	– step 2: a transitive verb is derived from the base *humidify* by prefixing *de-*, meaning 'to remove, reverse or perform the opposite action'
dehumidifier	– step 3: a noun is derived from the base *dehumidify* by adding the suffix *-er*, meaning 'something which performs the action of'
dehumidifiers	– step 4: the noun is made grammatically plural (inflected) by adding the regular plural suffix *-s* (in its allomorphic form [-z])

Typically, if a morphological operation causes a word-category change (such as from adjective to verb, as in *humid* → *humidify*), the process is considered derivational, since a new major-category lexeme has been created from an existing one. However, not all derivational operations cause a category change; for example, the prefix *de-* above attaches to a verb and derives another verb with a different meaning. Inflection, on the other hand, does not usually produce a category change. Adding a plural affix to a noun, for example, only grammatically augments that noun; it does not change its category. Similarly, adding an agreement marker to a verb results in the same verb, but one with an added formal feature. We will take a closer look at derivation in the next section, and at inflection in the following section.

Derivation

As mentioned above, derivation creates or *derives* new lexemes from existing ones. It allows new words to enter a language, even if sometimes only fleetingly in a particular conversation or magazine text or e-mail message. Derivation is also extremely useful for expressing phrases more compactly. It is much more efficient, for instance, to refer to someone working in the field of science, politics, or banking as a scient*ist*, politic*ian*, or bank*er* than to have to repeatedly use more cumbersome phrases such as "someone who works in the field of …". Derivation is a kind of shorthand system that allows us to economize – by packing more information into shorter utterances.

Although all the morphological processes outlined above are employed among the world's languages to derive new words, in this section we will focus on two: affixation (the most common process), and another kind of morphological process especially relevant to derivation – **compounding**.

Derivational affixes

Derivational affixation is the most common way among the world's languages to derive one lexeme from another. As mentioned above, derivation often changes the lexical category of a word, or its meaning, or both. We can observe some examples of this in various languages below:

(34) *Mandarin Chinese* (Li and Thompson 1981: 41–42)

gōngyè-*huà*	dòngwù-*xué*	kēxué-*jiā*
industry-V	animal-'ology'	science-ist
'industrialize'	'zoology'	'scientist'

German

Zerstör-*ung*	Einsam-*keit*	erb-*lich*
destroy-N	lonely-N	inherit-A
'destruction'	'loneliness'	'hereditary'

French

faibl-*esse*	chant-*eur*	rapide-*ment*
weak-N	sing-er	rapid-Adv
'weakness'	'singer'	'rapidly'

In each of the examples above, a suffix has applied to a particular kind of lexeme to derive another. In many cases, there is a category change; for example, the suffix *-ung* in German applies to verbs to derive a noun indicating a result of the verb (*zerstör-* 'destroy' → *Zerstörung* 'destruction'). (Nouns are conventionally capitalized in German orthography.) The French suffix *-esse* attaches to adjectives to derive nouns meaning something like 'the state or quality of being A' (*faible* 'weak' → *faiblesse* 'weakness'). The Chinese suffix *-jiā* derives a noun from another noun; here there is no category change but an agentive meaning is added – that of someone who practices in the field of the base noun (*kēxué* 'science' → *kēxuéjiā* 'scientist').

Let's turn now to English, a language that is quite rich in derivational morphology, with several different affixes sometimes sharing a similar function. Consider the following data:

(35)

sing-er	appli-cant
violin-ist	prank-ster
magic-ian	cook

In each word in (35), a noun has been derived that bears an obvious agentive relation to the root: a singer sings, a violinist plays the violin, a magician performs magic, an applicant applies for something, a prankster commits pranks, and a cook cooks. The agentive meaning in these examples is expressed by five different suffixes or, as in the case of *cook*, by nothing at all (the latter process is called **zero derivation**).

However, not all the affixes in (35) above can attach freely to any root. The suffix *-er*, for example, can only attach to verbs (*singer*, *smoker*), while the suffix *-ist* attaches only to nouns or adjectives (*violinist*, *cartoonist*), and *-ian* attaches only to nouns, especially those of Greek origin (*mathematician*, *politician*).

Because derivational affixes are selective in what they can modify, they generally apply in a particular order within a complex word. Consider again our earlier example of *dehumidifier.* Although the order of each step was spelled out in (33) above, we can more formally notate the derivational order by using various methods, such as a tree diagram as shown in (36a), bracketing (36b), or simple numbering (36c):

(36) a.

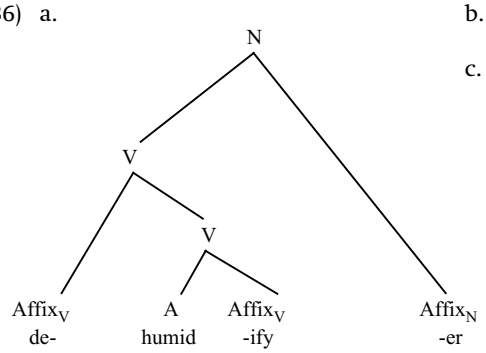

b. [[de$_V$[[humid$_A$] ify$_V$]] er$_N$]

c. de$_V$-humid$_A$-ify$_V$-er$_N$
2 ROOT 1 3

Some of the many derivational affixes of English are shown below.

(37)

Prefixes	*Category selected*	*Category derived*	*Examples*
de-	V	V	demagnetize, decompress
dis-	V	V	disentangle, dislocate
mis-	V	V	mismatch, mismanage
pre-	V	V	preview, predigest
re-	V	V	reappear, repossess
un-	A	A	unhappy, unproductive
un-	V	V	unwrap, unzip
Suffixes			
-able	V	A	bearable, washable
-al	V	N	approval, rebuttal
-ant	V	N	applicant, inhabitant
-ate	A	V	activate, validate
-en	A	V	redden, shorten
-er	A	A	singer, gambler
-ful	N	A	plentiful, beautiful
-ian	N	N	magician, musician
-ify	A/N	V	purify, beautify
-ion	V	N	detection, discussion
-ist	N/A	N	artist, activist
-ity	A	N	sensitivity, portability
-ive	V	A	oppressive, instructive
-ize	N	V	vaporize, magnetize
-ment	V	N	management, settlement
-ness	A	N	happiness, fullness
-y	N	A	watery, snowy

Some derivational affixes are very **productive**; that is, they can apply almost without exception to a certain kind of base. For example, the affix *-able* freely attaches to transitive verbs, deriving a new adjective with the meaning 'able to be V-ed' (as in *washable*, *faxable*, *analyzable*). On the other hand, some derivational affixes occur in only a small number of words and aren't productive, such as *-dom* (*kingdom*, *wisdom*, *boredom*) and *-th* (*warmth*, *truth*, *width*). Derivational affixes that are very productive at some point in the history of a language may become less so over time. The feminizing suffix *-ess* used to be more productive than it is today. Although there are still some words in common usage such as *actress*, *princess*, and *goddess*, the words in (38) were also once more widely used in English:

(38)

authoress	janitress
ambassadress	manageress
editress	mayoress
governess	poetess
huntress	proprietoress

This decline in the productivity of *-ess* affixation has likely been fueled by social factors that favor terms that de-emphasize or are completely neutral with regard to sex, such as *flight attendant* (instead of *stewardess*) and *server* (instead of *waitress*).

Compounding

Compounding is the concatenation of two (or more) lexemes to form a single new lexeme. Because compounding always results in the creation of a new lexeme, it is a good example of a derivational process. In English and many other languages, compounding is highly productive and a primary source of new vocabulary. Some English examples are shown below:

(39) *English*

greenhouse	soy sauce	man-made
moonlight	coast guard	brown-eyed
download	long shot	long-range

As shown above, English compounds are sometimes written as a single word, or as words separated by a space or a hyphen; however, they are all considered by linguists to be compounds. The characteristic pronunciation for English compounds is for the stress to fall on the first lexeme in a two-lexeme compound (although there are exceptions). Recall the difference in pronunciation between the compounded word *gréenhouse* (a place to grow plants) vs. the phrase *green hóuse* (a house that is green).

In some languages, such as English, German, and Dutch, compounding can be highly **recursive**, meaning that a derived compound can serve as the base for further compounding:

(40) toe + nail
toenail + clipper
toenail clipper + accident
toenail clipper accident + insurance
toenail clipper accident insurance + company
toenail clipper accident insurance company + employee
toenail clipper accident insurance company employee + benefits
etc.

Box 2.5 Compound word or syntactic phrase?

We might ask whether compounds like *toenail clipper accident insurance company . . .* are really words rather than syntactic phrases. Everyone seems to agree that *toenail* is a word and even *toenail clipper* still seems pretty wordlike. Is there some definable point at which complex words lose their wordhood status and become syntactic phrases? If you recall the highly complex derived Turkish word shown earlier in example (11) that must be translated as an entire complex sentence in other languages, you'll understand why linguists have been debating

this question for years! And yet, even the most complex form in (40) above is still more wordlike than phrasal. Recall that morphological derivation renders utterances more compact; now imagine the complexity of the sort of syntactic phrase we'd need to capture a similar meaning: 'the benefits of the employees of a company that sells insurance in case of accidents that involve the use of clippers that are used to clip the nails on one's toes.' (And even this horrid phrase contains 'shortcuts' – the derived words *employees*, *insurance,* and *clippers*.)

In each successive derivation in (40), it is the rightmost element that identifies what the compounded word is; in other words, a *toenail clipper* is a particular kind of clipper (as opposed to say, a hedge clipper or nose-hair clipper); a *toenail clipper accident* is an accident that somehow involves a toenail clipper, etc. This identifying element is called the **head**; its meaning and part-of-speech category determine that of the entire compound overall. English compounds are typically right-headed. Thus, *chocolate milk* is a kind of milk but *milk chocolate* is a kind of chocolate. Both compounds are nouns, because both *milk* and *chocolate* are nouns. The examples below demonstrate right-headedness in compounds involving various word categories in addition to the type of noun–noun compounding we have looked at so far.

(41)

Noun compounds	*Adjective compounds*	*Verb compounds*
N + N chocolate milk, toenail	N + A headstrong, skin-deep	N + V handpick, fingerspell
A + N softball, shortcake	A + A bittersweet, aquamarine	A + V blacklist, soft-pedal
V + N drawstring, driveway	V + A slaphappy, punch-drunk	V + V blowdry, shrinkwrap
P + N instep, oversight	P + A underripe, overgrown	P + V undertake, oversleep

However, there are many compounds in English in which the overall category and meaning are not determined by the rightmost element. A few examples are given in (42):

(42)

$over_{P}$ + $weight_{N}$	→	$overweight_{A}$
$make_{V}$ + $shift_{V}$	→	$makeshift_{A}$
$lack_{V}$ + $luster_{N}$	→	$lackluster_{A}$
$make_{V}$ + $believe_{V}$	→	$make\text{-}believe_{A/N}$
$speak_{V}$ + $easy_{A}$	→	$speakeasy_{N}$
up_{Adv} + $keep_{V}$	→	$upkeep_{N}$

Moreover, there is a quite productive class of compounding in English in which certain types of verb phrases (verbs plus adverbs, prepositions or verb particles) are compounded into nouns. In this case, the right-head generalization does not hold at all:

(43)

kickback	screw-up	payoff	buyout
sing-along	breakdown	pullover	sit-up
breakthrough	giveaway	workout	getaway
drawback	get-together	drive-in	heads-up

In these cases, the morphological process of zero-derivation has applied in addition to compounding, turning verbal phrases (like *break dówn*) into compound nouns (*bréakdown*). Since neither lexeme in the compound determines its overall grammatical category or meaning, these compounds are generally considered unheaded.

Box 2.6 Compounding in various languages

Some examples of compounds in other languages are provided below. Can you identify whether they are right-headed or left-headed?

Hebrew (Berman 1997: 323)

orex-din	tapúax-adama
conductor-law	apple-earth
'lawyer'	'potato'

Japanese (Shibatani and Kageyama 1988: 454)

hai-zara	kosi-kakeru
ash-plate	waist-hang
'ashtray'	'to sit down'

Mandarin Chinese (Li and Thompson 1981: 47–49)

fēi-jî	hé-mā
fly-machine	river-horse
'airplane'	'hippopotamus'

Jacaltec (a Mayan language) (Spencer 1991: 349)

potx'-om txitam	'il-om 'anma
kill-er pig	watch-er people
'pig-killer'	'people-watcher'

To help determine the head of each compound, we can apply the simple 'is a' test employed above; for example, *chocolate milk* 'is a' kind of milk, not a kind of chocolate, and thus *milk* is the head and the compound is right-headed. Let's analyze a few of the above compounds.

In Hebrew, the compound *orex-din* 'lawyer' is a person who practices ('conducts') law, rather than a kind of law; therefore, we deduce that the word *orex* 'conductor' is the head of the compound and the compound is left-headed. Now look at the other Hebrew compound, *tapúax-adama* 'potato.' Is a potato a kind of apple or a kind of earth? Well, metaphorically, it's much closer to being a kind of apple than a kind of earth. (For those who know French, the resemblance to the French word for 'potato' will be obvious: *pomme de terre* 'earth apple.') Thus we conclude that *tapúax* 'apple' is the head and the compound is

left-headed. From our very limited sample of data it appears that Hebrew compounds are left-headed.

Next let's turn to the Japanese compound *kosi-kakeru* 'to sit down.' The first word *kosi* 'waist' is a noun and the second *kakeru* 'hang' is a verb; the resulting compound is a verb. Which element is the head? Since sitting down is a kind of "hanging" of one's waist rather than a kind of waist and since the resulting compound is a verb overall, it appears that *kakeru* 'hang' is the head and the compound is right-headed. The other Japanese compound *hai-zara* 'ashtray' is also clearly right-headed: an ashtray is a kind of plate, not a kind of ash. It appears that compounding in Japanese is right-headed (but of course we'd need more data to confirm this).

Similar analyses indicate that the compounds in Chinese above are right-headed, whereas those in Jacaltec are left-headed.

Inflection

We observed earlier that morphological **inflection** adds grammatical information to a lexeme, depending on the particular syntactic requirements of a language. The kind of information added indicates a property or a **feature** within a set of grammatical contrasts, such as singular vs. plural, first person vs. second, masculine vs. feminine, past vs. nonpast, and many others. By "syntactic requirements of a language," we mean contexts in which a particular language requires us to make such a contrastive distinction. For example, the presence of the quantifier *two* in the following English sentence creates a context in which the grammatical feature [+plural] must be realized on any noun that could fill in the blank in the following sentence:

(44) Olivia bought *two* _____ online yesterday.

"Plural" is a feature associated with nouns in English (and many other languages), whether it's a simple noun, such as *book*, or a derivationally complex one, such as *dehumidifier*. It's also important to note that it is merely the grammatical information [+plural] that is required in this context by the rules of English syntax, rather than any particular morphological form. For example, we could still produce a grammatically well-formed sentence by filling in the blank with the irregularly pluralized nouns *mice* or *theses*. In other words, as long as the noun is plural, the syntax "doesn't care" which particular form is chosen. It's the job of the language-specific rules of morphology to determine how an abstract feature should be realized phonologically and to select the right form.

The following sections provide a mini-catalogue of some of the more common grammatical contrasts marked by inflectional morphology across the world's languages.

Person

Person is a grammatical feature that distinguishes entities referred to in an utterance. First person refers to the speaker (in English, *I*/*me*), and

second person refers to the addressee (*you*). Third person is a default category that refers to everything else (for example, *she/her*, *he/him*, *it*, *the dog*, *John*, *the fact that it might rain today*). Person is often combined with **number** (see the next section) and thus we often speak of person–number combinations such as 'third person singular' (3SG) or 'first person plural' (1PL), etc.

When we speak of the inflectional categories of person (and/or number), we're usually referring to a grammatical agreement relation, most often subject–verb agreement. Languages which distinguish grammatical persons typically require that a verb agree with its subject's person feature, and occasionally with that of its object as well. Subject–verb agreement helps to indicate which noun in a sentence is "doing" which verb. This is particularly valuable in languages with free word order, in which subjects can come before or after their verbs. English (which has relatively fixed word order) has only one inflectional agreement marker for its regular verbs (and only on present-tense verbs): 3SG *-s*, as in *he/she/it runs*. Below is a partial **paradigm** (an orderly display of related forms, for example, as here, contrasting persons and numbers) for present tense, subject–verb person/number agreement in Polish, a language with much richer verbal agreement contrasts, for the verb *kochać* 'to love.'

(45) *Polish present tense*

Person	*Singular*		*Plural*	
1	kocham	'I love'	kochamy	'we love'
2	kochasz	'you (sing.) love'	kochacie	'you (pl.) love'
3	kocha	'he, she loves'	kochają	'they love'

This table shows that Polish speakers use a different inflectional form of the verb depending on the person/number features of the subject.

Number

Number is a grammatical property of nouns (and as we've already seen, often marked via agreement on verbs and other elements such as determiners and adjectives). The most fundamental contrast is between **singular** ('one') and **plural** ('more than one'), although many languages also mark a distinct **dual** form ('two'). Slovenian marks this three-way number contrast, as shown in (46). (The noun *mest-* 'city' is in nominative case, to be discussed shortly.)

(46) *Slovenian*

mest*o*	(sing.)	'city'
mest*i*	(dual)	'two cities'
mest*a*	(plur.)	'cities'

Even in languages that have grammatical number distinctions there are some nouns that cannot be counted and therefore cannot be pluralized. For example, in English, abstract nouns such as *companionship*, *carelessness*, and

peace have no plural form. Nouns that denote non-individuated material like *rice*, *lettuce*, or *toilet paper* are also typically not pluralizable (unless we mean something like 'kinds of' rice, lettuce, or toilet paper). This type of noun is called a **mass** or **noncount** noun. A noun that is a mass noun in one language may be countable in another. In English, for example, *furniture* is a mass noun (**furnitures*), whereas in French, it is countable (*meuble*, *meubles*). In this situation, English must resort to using an additional word that can individuate parts of the mass, such as a *piece* of furniture, a *grain* of rice, a *head* of lettuce, or a *roll* of toilet paper.

In some languages which appear to lack regular grammatical number distinctions, such as Chinese, most nouns are individuated in this way. Consequently, Chinese has a highly developed system of **classifiers**, a kind of grammatical morpheme that affixes to quantifiers (like 'one,' 'some,' 'many,' 'three,' etc.) or demonstratives ('this' and 'that'):

(47) *Mandarin Chinese*

wǔ-*ge* rén	zhèi-*běn* shū	nèi-*zhāng* zhǐ
five-CL person	this-CL book	that-CL paper
'five people'	'this book'	'that sheet of paper'

There are dozens of classifiers in Mandarin Chinese. The choice of classifier depends on the particular noun and often must simply be memorized. Some patterns exist, such as the use of the classifier *-tiáo* for elongated objects, such as 'snake,' 'rope,' 'river,' 'tail,' and also for most four-legged mammals; however, there are several exceptions. For many native Chinese speakers today, however, the most frequently used classifier *-ge* has begun to take on a more general all-purpose character and is gradually replacing many of the more specialized ones (Li and Thompson, 1981: 112).

Gender

In many languages, nouns are sorted into different classes that other words – such as adjectives, determiners, pronouns, or verbs (or some combination of these) – must agree with. These noun classes are often referred to as **gender**, etymologically descended from Latin *genus* and French *genre*, meaning 'kind' or 'sort.' Gender agreement helps to indicate which adjectives, determiners, etc. are associated with a particular noun. In languages that mark grammatical gender, every noun is assigned to a class. You might be familiar with terms like "masculine" and "feminine" to describe these classes in some languages with a two-way distinction (like French and Spanish), or "masculine," "feminine," and "neuter" for some with a three-way distinction (like German and Russian).

Sometimes gender is indicated on the noun itself. In the Spanish examples in (48), the masculine noun *amigo* ('friend, masc.') ends in *-o* whereas its feminine counterpart *amiga* ('friend, fem.') ends in *-a*. Observe that the forms of the indefinite article *un/una* and the adjective *americano/a* agree with the gender of the noun:

(48) *Spanish*

un	amig*o*	american*o*	(masc.)	'an American friend (male)'
un*a*	amig*a*	american*a*	(fem.)	'an American friend (female)'

However, many nouns in Spanish are not so obviously marked. The gender of Spanish nouns that end either in consonants or vowels other than *-o* and *-a* must be learned:

(49)	el balcón	(masc.)	'the balcony'	el coche	(masc.)	'the car'
	la razón	(fem.)	'the reason'	la noche	(fem.)	'the night'

Even worse, there are "misleading" nouns ending in *-a* that are masculine and some ending in *-o* that are feminine:

(50)	el clima	(masc.)	'the climate'	la mano	(fem.)	'the hand'
	el dia	(masc.)	'the day'	la foto	(fem.)	'the photo'

The way to determine a noun's gender is by observing the type of determiners, adjectives, etc. that agree with it. The gender of a noun is conclusively confirmed by agreement.

Languages which mark gender differ in terms of the types of words that must agree with the noun. In the Spanish examples above, the articles (such as *un*/*una* and *el*/*la*) and adjectives (such as *americano*/*a*) agree. In the following German examples, only the articles agree:

(51) *German*

der junge Mann	(masc.)	'the young man'
die junge Frau	(fem.)	'the young woman'
das junge Mädchen	(neut.)	'the young girl'

In the Swahili sentence below, the subject noun *ki-kapu* 'basket,' which belongs to Swahili noun class 7, requires agreement on the adjective (*ki-kubwa* 'large'), quantifier (*ki-moja* 'one'), and verb (*ki-lianguka* 'fell') – all indicated by the 'noun class 7' prefix *ki-*:

(52) *Swahili* (Corbett 1991: 117)

ki-kapu	*ki*-kubwa	*ki*-moja	*ki*-lianguka
7-basket	7-large	7-one	7-fell

'one large basket fell'

Perhaps you're surprised at a language having a "noun class 7." Swahili (also called Kiswahili) belongs to the family of Bantu languages spoken throughout much of the southern half of Africa. Like most Bantu languages, it has a very elaborate gender system. There are between ten and twenty noun classes in Bantu languages (compared with the two to four found in most Indo-European languages), and Swahili has about fifteen (depending on how these are analyzed). Among the world's languages, the basis for assigning nouns to gender classes is often complex and, as you can see from the examples above, goes well beyond semantic-based concepts of "masculine" or "feminine" (see Box 2.7).

Box 2.7 How do nouns get a gender?

Although gender in many languages at least partially correlates with sex-based distinctions, you can see from the examples in the text that morphological gender is also assigned in most cases independently of semantic "maleness" or "femaleness." (For linguists the word *gender* means 'kind' rather than 'sex.') So how is a noun assigned to a particular class? There are two main ways – either on the basis of its meaning or its form. Most systems are mixed; that is, some nouns are assigned on the basis of meaning and others according to their form.

In French, for example, nouns denoting male referents are masculine and those denoting females are feminine – a meaning-based distinction for gender class assignment. What about all the other nouns – the overwhelming majority? Two studies (Mel'čuk 1958; Tucker, Lambert, and Rigault 1977) found surprisingly accurate predictors of the gender of non-semantic-based nouns in French derivational morphology and phonology. For example, noun compounds derived from verb phrases in French are masculine, even if the final (object) noun in the compound is feminine, as in *porte-monnaie* = carries-money (fem.) = 'purse' (masc.). Similarly, 99.8% of nouns ending in *-tion* /sjɔ̃/ are feminine (e.g. *action* /aksjɔ̃/ 'action'); 97.2% of nouns ending in /o/ are masculine (e.g. *mot* /mo/ 'word'), as are 94.2% of nouns ending in /ʒ/ (e.g. *potage* /pɔtaʒ/ 'soup'), and so on. Children acquiring French are able to keep track of these regularities, and probably learn the exceptions in much the same way that children acquiring English learn the past-tense forms of irregular verbs. In fact, the French gender-assignment system is so heavily phonology-based that deaf children who learn to speak French do not acquire it, because they cannot hear the language (Tucker et al. 1977: 59, cited in Corbett 1991: 61).

In languages with highly complex noun class systems, such as Swahili, the gender system is also mixed, with nouns assigned to particular classes either on the basis of their meaning or their form. Some of the semantic features used to classify nouns in these languages are humanness, sex, animacy (with distinctions for some kinds of plants and animals), body parts, size, and shape, although there are exceptions in many classes. Membership in other classes is phonologically or morphologically determined. In Bantu languages, for instance, there are distinct classes for the plural referents of nouns of another given class. Recall the Swahili example above in (52) meaning 'one large basket fell.' If we were to change 'one' to a plural number, e.g. 'three,' the noun class of *ki-kapu* 'basket' would change from 'class 7' to 'class 8' *vi-kapu* 'baskets,' and all the agreement markers on the other sentential elements would change as well from *ki-* to *vi-*:

vi-kapu	*vi*-kubwa	*vi*-tatu	*vi*-lianguka	(Corbett 1991: 44)
8-baskets	8-large	8-three	8-fell	
'three large baskets fell'				

Thus, the "meaning" of gender class 8 in Swahili is 'plural of class 7 nouns.'

What if a language (or more precisely, a speaker) can't decide which class to use? As we observed earlier with the Chinese classifier system, languages with complex noun class systems may often designate one class as a kind of general all-purpose default gender. This gender is typically overgeneralized by children and especially second-language learners; it gradually absorbs the nouns of smaller, more obscure classes, and may also serve as the primary class for new nouns that enter a language.

Case

We observed earlier that one of the most important functions of morphology is to distinguish the roles played by the various participants in an event. **Case** is a grammatical category that does this, by indicating a particular noun's relation to some other element in a clause or phrase. Typically, case marking indicates the relation of the noun to the verb (as its subject, direct object, or indirect object) or to another noun (as in a possessive or locational relation). Before we see how these relations are morphologically marked in different languages, let's briefly consider these relations as illustrated in the English sentence in (53).

(53) John gave Mary his sister's old bicycle.

In this sentence, the verb *gave* is related to three nouns, the giver (*John*), the gift (*bicycle*), and the recipient (*Mary*). The giver *John* is the subject of *gave*, the gift *bicycle* is the direct object, and the recipient *Mary* is the indirect object. In addition, there are two possessive relations – one between John and his sister (marked with the possessive pronoun *his*), and another between his sister and the bicycle (marked with the possessive affix *-s* on *sister*). In languages that mark case distinctions, many of these relations would be indicated by inflectional morphology. Consider the same sentence in Slovenian, for example. (In (54), case forms are labeled as follows: **nominative** for subjects, **accusative** for direct objects, and **dative** for indirect objects.)

(54) *Slovenian* (Rus, p.c.)

Janez	je	dal	Marij-i	sestr-in-o	star-o	kol-o.
John.NOM	is	given	Mary-DAT	sister-POSS-ACC	old-ACC	bicycle-ACC

'John gave Mary his sister's old bicycle.'

In Slovenian, nominative case is unmarked; the base form of the noun receives no affix. Other cases are marked – dative by the suffix *-i* and accusative by *-o*. Notice that adjectives such as *staro* 'old' and *sestrino* 'sister's' (which is considered a possessive adjective in Slovenian) also agree in case marking with the noun they are modifying.

In addition to indicating grammatical roles such as subject, object, and possession, many languages with extensive case-marking systems also use

inflectional morphology to mark the kinds of locational relations for which English uses prepositions (like *to*, *from*, *at*, etc.). These locational suffixes are called **locative** case inflections. Here are some examples of this type of case suffixing in Lezghian, a Caucasian language spoken in southern Russia:

(55) *Lezghian* (Beard 1995: 261-262)

sev	'bear'	sev-rek	'under the bear'
sev-rev	'at the bear('s)'	sev-rek-di	'to under the bear'
sev-rev-di	'toward the bear'	sev-rek-aj	'out from under the bear'
sev-rev-aj	'away from the bear'	sev-re	'in the bear'
sev-rex^{h}	'behind the bear'	sev-rej-aj	'out of the bear'
sev-rex^{h}-aj	'out from behind the bear'	sev-rel-aj	'off of the bear'

The locative case affixes of Lezghian provide another illustration of one language using bound morphemes to perform the same function that free morphemes (e.g. prepositions) perform in another.

Tense

All animal communication systems can convey information about a current or imminent situation – a warning signal, for example, or a mating call. Humans are exceptional in that we can talk about events that occurred in the past, no matter how distant, or speculate about situations that may or may not happen far into the future. All human languages have ways for locating situations in time – for example, through the use of lexical expressions like *yesterday*, *the day after tomorrow*, *when the sun sets*, or *in the rainy season*. In addition, many languages also use a morphological category called **tense** to locate an event or state in relation to a point in time. In simple tenses, such as past, present, and future, that reference point is "now," at the moment of speaking. The past-tense, for example, indicates that an event took place prior to the moment of speaking, as shown below for English and Japanese:

(56) Last night I *went* to a karaoke bar.

(57) *Japanese*

Yuube	karaoke	baa-ni	ikimasi-*ta*
last night	karaoke	bar-to	go-PAST

'Last night I went to a karaoke bar.'

The use of the past-tense here indicates that the event of my going to a karaoke bar happened sometime prior to the moment of my telling about it; that time is more precisely specified by the adverbial phrase *last night*.

Whereas past-tense marks an event as occurring prior to the moment of speech, the future tense situates an event sometime after the moment of speech. Note that while we often speak of a 'future tense' in English, it is not morphologically marked via inflection; rather we use various auxiliary forms or the simple present form:

(58) He's leaving next week.
He's going to leave next week.
He will leave next week.
He leaves next week.

For this reason, English is claimed to morphologically distinguish only two tenses: past and nonpast. Other languages, such as Italian, do have morphological future tense:

(59) *Italian*

Part-ir-à	la	settimana	prossima
leave-FUT-3SG	the	week	next

'He will leave next week.'

Not all languages morphologically mark tense. Dyirbal (an Australian language) and Burmese are claimed to lack tense (Comrie 1985). So is Chinese, as shown below:

(60) *Mandarin Chinese* (Li and Thompson 1981: 214)

zuótiān	yè-lǐ	wǒ	mèng-jiàn	wǒ	mǔqîn
yesterday	night-in	1SG	dream-perceive	1SG	mother

'Last night I dreamed about my mother.'

In (60), there is no past-tense marker on the verb *mèngjiàn* 'dream' (or anywhere else in the sentence). The past time of the event is simply lexically indicated by the expression *zuótiān* 'yesterday.' Although all languages are able to express semantic time reference lexically, this should not be confused with the formal grammatical category of tense.

Aspect

Whereas tense is concerned with locating an (entire) event in time relative to the moment of speaking, **aspect** is a grammatical category that encodes a different kind of temporal characteristic, such as whether an action is (or was) completed, ongoing, repeated (iterative), or habitual. For example, in English one can say:

(61) a. John is painting the kitchen.
b. John was painting the kitchen.
c. John painted the kitchen.

The difference between (61a) and (61b) is a tense distinction (i.e. present vs. past), whereas the difference between (61b) and (61c) is an aspectual distinction. Both (61b) and (61c) are in the past-tense; however, the form *was painting* in (61b), often called the **progressive** form in English, indicates that the action of painting was ongoing and it's quite possible that John never completed painting the kitchen. (Imagine if John was interrupted after he'd just begun painting.) The example in (61c), however, indicates that John did finish painting – that is, the painting of the kitchen was completed. This aspectual distinction is found in many languages; the "ongoing" meaning is sometimes referred to by linguists as the **imperfect(ive)** and the "completed"

meaning is called the **perfect(ive)**. An example of how imperfective aspect is used in French and Spanish is shown below:

(62) *French and Spanish* (Comrie 1976: 3)

Jean	*lisait*	quand	j'entrai
Juan	*leía*	cuando	entré
John	read-IMP	when	(I) enter-PAST

'John was reading when I entered.'

The verb forms *lisait* and *leía* (French and Spanish, respectively) indicate that John was in the process of reading when he was interrupted by my entering, and we cannot assume that he finished whatever he was reading.

The following example from Chinese illustrates perfective aspect:

(63) *Mandarin Chinese* (Li and Thompson 1981: 213)

wǒ	chī-*le*	fàn	zài	zǒu
I	eat-PERF	rice	then	go

'I'll go after I eat.'

In (63), the verb *chī* 'eat' bears the perfective marker *-le*, indicating that the action of going will occur after eating has been completed. (Note that in English we could also translate this sentence using an equivalent morphological construction – the present perfective: 'I'll go after I'*ve* eat*en*.')

Aspect should be distinguished from tense; in the Chinese sentence in (63), for example, the act of eating is marked as "complete" and yet is not in the past. In practice, however, these two categories are often highly interactive. Consider again sentence (61c) in English: *John painted the kitchen*. For events such as kitchen-painting, our English "past-tense" form also encodes perfective aspect. In other words, the entire painting event is situated prior to the moment of speaking (past tense), *and* the painting of the kitchen was completed (perfective aspect). Similarly, the nonpast (present) tense (for example, *is painting* in the sentence *John is painting the kitchen*) indicates that the event is situated in time at the moment of speaking (tense), and also indicates that the action is still ongoing (imperfective aspect).

Mood

Mood is a grammatical category that expresses the speaker's belief, opinion, or attitude about the content of an utterance. Although often morphologically marked on verbs, mood really applies to entire clauses, to indicate whether the speaker thinks a proposition is true, or likely, or doubtful, or is something he/she wonders about, or hopes or wishes for. Some common mood distinctions across languages are the **indicative**, used for making declarative assertions, the **interrogative**, used for asking questions, and the **imperative**, used for giving commands. We also find languages which use a special mood form known as the **subjunctive** to express desire, hope, or doubt, and the **conditional** to express what one

would or should do. Examples illustrating indicative (64a), subjunctive (64b), imperative (64c), and conditional (64d) moods of the Italian verb *venire* 'to come' are given below:

(64) *Italian*

a. *Viene* dalla biblioteca.
   come.3SG.PRES.*IND* from.the library
   'He/she is coming from the library.'

b. Spero che *venga* presto domani.
   hope.1SG.PRES. that come.3SG.PRES. early tomorrow
   IND *SUBJUNCT*
   'I hope that he/she comes early tomorrow.'

c. *Vieni* a casa presto.
   come.2SG(FAM.).*IMPER* to home soon
   'Come home soon.'

d. *Verrebbe* voluntieri, ma non serà in città.
   come.3SG. gladly but not be.3SG.FUT in town
   PRES.*COND*
   'He/she would gladly come, but he/she won't be in town.'

Another interesting modal function often expressed inflectionally in languages is **evidentiality**, in which the speaker indicates a degree of certainty or doubt about a proposition based on the kind of evidence available for it. Think for a moment about how important it is for humans to be able to evaluate the reliability of an information-source – our survival literally depends on it. For example, we're likely to believe what we've witnessed firsthand with our own eyes more than what we've heard from second-hand reports by others. In English, we convey our degree of confidence in the truth of an assertion by lexical means, as in phrases like *I know . . .* or *I heard . . .* or *I doubt . . .* or by the use of adverbs such as *apparently*. However, other languages directly inflect these meanings on the verb. In Quechua, for example, there is a three-way evidentiality distinction, as shown below:

(65) *Quechua* (Sánchez 2003: 21)

a. kaya-n-*mi* sanurya.
   be-3SG-*AFFIRM* carrots
   'There ARE carrots.'

b. huk hunaq-*shi* pukla-shun.
   one day-*REPORT* play-1PL.INCHOAT
   'We will play one day (they say).'

c. Pay rura-nqa-*tr*.
   he do-3-SG.FUT-*DUBITATIVE*
   '(I doubt) he will do it.'

In (65a) the verbal affix *-mi* indicates firsthand knowledge (affirmative), in (65b) the affix *-shi* relays information based on hearsay (reported), and in (65c) the affix *-tr* expresses doubt (dubitative). In Quechua these markers also serve to place focus on the word they are affixed to.

Grammatical mood may also express wonder, surprise or disappointment with a result that is contrary to expectation. The following examples illustrate the "five modes" for verbs in Menomini, an Algonquian Native American language:

(66) *Menomini* (Hockett 1958, cited in Palmer 1986: 2)

piw	'he comes/is coming/came'
piwen	'he is said to be coming / it is said that he came'
piʔ	'is he coming? / did he come?'
piasah	'so he is coming after all! (despite our expectation to the contrary)'
piapah	'but he was going to come! (and now it turns out that he is not)'

The semantically varied and complex category of mood shades the meaning of a sentence, by conveying some opinion, attitude, or emotion about its content as perceived by the speaker. All languages have a way to express mood distinctions, and the majority do so morphologically.

Acquiring inflectional contrasts

If you have formally studied a language that has rich subject–verb agreement or grammatical gender distinctions, such as Spanish, Arabic, Swahili, or German, then you've probably memorized tables of neatly-arranged verb forms or lists of vocabulary words in which information about the gender of nouns was provided. But what if you were a child acquiring Swahili or Spanish naturally? It's important to remember that in a child's natural acquisition environment, words do not come annotated with information such as "*coche* (masc.)" or "*noche* (fem.)". What the student in a language classroom learns via the memorizing of lists and paradigms, the preliterate child must acquire from unconscious attention to the contrasting distribution and co-occurrences of **morphophonological** distinctions in the language they hear. In languages with complex morphology, this is a remarkable feat.

Consider all the information that is artificially condensed in a morphology example or problem set (including the ones in this chapter), such as English glosses and translations. Children acquiring a language don't get this kind of help. Instead, they swim in a linguistic sea where the meaning of an utterance and distinctions such as case, number, person, gender, conjugation class, and so on, must be deduced in context and compared with minimally contrasting forms that have already been acquired. It is the gradual accumulation of such minimal contrasts that leads to the child's formulation and increasing refinement of grammatical categories.

The aggregate knowledge accumulated over countless hours of exposure to a language in context is quite literally "lost in translation" in the way data are carefully handpicked and presented in linguistics textbooks. For that reason, we encourage you to consider someday undertaking some real field work of your own; you've done it already – as a child!

Box 2.8 Instructed vs. naturalistic acquisition of morphological agreement: an example

If you've studied French in a foreign-language classroom setting, you probably learned that adjectives must agree with the nouns they modify in gender and number. Let's focus for a moment on gender agreement. It's likely that you were provided with a list of adjectives and the following rule: "To form the feminine of an adjective, add *e* to the masculine form" (e.g. Barson 1981: 101). (In some cases there might be some other minor spelling change as well.) The list below is quite typical (data from Bauer 2003: 38-39):

Orthography		Pronunciation		
Masc.	Fem.	Masc.	Fem.	
mauvais	mauvais*e*	/movɛ/	/movɛz/	'bad'
heureux	heureus*e*	/œʁø/	/œʁøz/	'happy'
grand	grand*e*	/gʁɑ̃/	/gʁɑ̃d/	'big'
long	longu*e*	/lɔ̃/	/lɔ̃g/	'long'
chaud	chaud*e*	/ʃo/	/ʃod/	'hot'
vert	vert*e*	/vɛʁ/	/vɛʁt/	'green'
froid	froid*e*	/fʁwa/	/fʁwad/	'cold'
petit	petit*e*	/pəti/	/pətit/	'little'
blanc	blanch*e*	/blɑ̃/	/blɑ̃ʃ/	'white'
faux	faus*se*	/fo/	/fos/	'false'

Note that this rule is completely based on written orthography; it suggests that final feminine *e* is an inflectional affix. This way of "looking at it" might help a second language learner (who already knows how to read) to learn the pattern. However, morphology is not based on orthography, but on morphemes which are sequences of *sounds*.

Now imagine a young child acquiring French who hasn't yet learned to read or spell. But children of preschool age already know and use correctly both the masculine and feminine forms of the adjectives listed above. So what "rule" does the preliterate child follow to acquire them? Let's look at how the words above are pronounced – the only kind of evidence available to the child.

The preliterate child cannot make use of a spelling rule that adds a letter (*e*). Note that the feminine form has an added consonant *sound* at the end of the word compared with its masculine counterpart. However, as Bauer (2003) points out, there is no way to predict from the masculine form which added consonant the feminine form will take. This is determined by each word. The only rule the child can establish on the basis of the phonological data is that the masculine adjective can be predictably derived by deleting the final consonant sound from the feminine adjective. This is the opposite order of derivation from the spelling-based orthographic rule presented in French language textbooks, and illustrates how different the mental representations of morphological rules can be between native vs. classroom-instructed language acquirers!

Chapter summary

A key part of knowing a language is the ability to construct and interpret the *words* of that language. The branch of linguistics that is concerned with the relation between form and meaning within words and between words is called *morphology*. The basic unit of language that combines both a form and a meaning is the *morpheme*. Simple words consist of just one morpheme, whereas complex words consist of more than one morpheme and may consist of many. There are different kinds of morphemes. Those bearing richer lexical meaning and belonging to the major lexical categories of nouns (N), verbs (V), and adjectives (A) are called *lexemes*, and may serve as the *root* for additional morphological operations. Those serving primarily to signal a grammatical function are called *grammatical morphemes*. The actual phonetic forms of morphemes can vary systematically depending on certain conditioning factors; these variant forms are known as *allomorphs*.

Languages make use of various morphological operations to modify the form and meaning of lexemes. The most common process is *affixation*, in which a morpheme is added to a base (either a root or another affix). Other processes include *reduplication* (the copying of all or part of a root), internal root changes such as *ablaut* and *suppletion (*involving the replacement of all or some part of a root's segments), suprasegmental change (in which a shift in word stress or tone is used to signal a morphological contrast), and *compounding* (in which two lexemes are combined to form a new lexeme). These morphological operations are used in the service of two major functions: *derivation* and *inflection*.

Derivational morphology creates new lexemes from existing ones, with a change in the word's lexical category or meaning, or both. *Inflectional* morphology adds grammatical information to a lexeme, as required by the particular grammatical rules of each language. Some common inflectional contrasts found among the world's languages are *person*, *number*, *gender*, *case*, *tense*, *aspect*, and *mood*.

Finally, we briefly considered how children manage to acquire some of these grammatical contrasts. Children appear to be able to keep track of distributional regularities found in the input – for example, correlations between gender distinctions and the phonological regularities of roots. The meanings of grammatical forms such as case and tense marking must be deduced from the use of these forms in informative and minimally contrastive language contexts.

Exercises

Exercise 2.1

Indonesian. Consider the Indonesian reciprocal forms below (from Sneddon 1996:104):

ROOT		RECIPROCAL	
tolong	'help'	tolong-menolong	'help each other'
pukul	'hit'	pukul-memukul	'hit each other'
kunjung	'visit'	kunjung-meŋunjung	'visit each other'
peluk	'embrace'	peluk-memeluk	'embrace each other'
telpon	'telephone'	telpon-menelpon	'telephone each other'

a. Given these data, can you derive the rule for forming the reciprocal in Indonesian? Pay attention to any phonological changes that may occur.
b. Given this rule, how would you form the reciprocal form of the roots *tikam* 'stab,' *pinjam* 'borrow,' and *tawar* 'bargain'?

tikam	'stab'	_____	'stab each other'
pinjam	'borrow'	_____	'borrow from each other'
tawar	'bargain'	_____	'bargain with each other'

Exercise 2.2

Finnish. The Finnish data below (from Leney 1993) are given in both the nominative and the partitive case. Partitive is used for quantified nouns, such as *monta kilometri-ä* 'many kilometer-PART,' *kolme kuningas-ta* 'three king-PART' or *lasi viini-ä* 'a glass of wine-PART'. Study the data and answer the questions that follow.

BASE (NOM)		PARTITIVE
kilometri	'kilometer'	kilometriä
kuningas	'king'	kuningasta
viini	'wine'	viiniä
runo	'poem'	runoa
hautausmaa	'cemetery'	hautausmaata
leipä	'bread'	leipää
tytär	'daughter'	tytärtä
tyttö	'girl'	tyttöä
katu	'street'	katua
kyljys	'pork chop'	kyljystä
sana	'word'	sanaa
olut	'beer'	olutta
tee	'tea'	teetä
mies	'man'	miestä
tie	'road'	tietä
näytös	'act'	näytöstä
pullo	'bottle'	pulloa
kollega	'colleague'	kollegaa

a. The data exhibit allomorphy. Identify the allomorphs and describe how they're conditioned. (Hint: the umlaut '¨' diacritic indicates that a vowel has been fronted.)

b. Provide the correct partitive forms for the following Finnish words:

henkilö	'person'	_____
loma	'holiday'	_____
esitys	'performance'	_____
maa	'country'	_____

Exercise 2.3

Yoruba. Consider the following data from Yoruba, a language spoken in Nigeria (from Akinlabi, to appear). (Tone marks have been omitted for the sake of simplicity.)

gbona	'be warm'	gbigbona	'warmth, heat'
dara	'be good'	didara	'goodness'
won	'cost a lot'	wiwon	'costliness'
je	'eat'	jije	'(act of) eating'
ran	'sew'	riran	'(act of) sewing'
gbe	'take'		
mu	'drink'		

a. What morphological process is illustrated in the data and what is its function? Describe how the forms on the right are derived from those on the left.
b. Please provide the likely forms for the meanings '(act of) taking' and '(act of) drinking.'

Exercise 2.4

Hebrew. Consider the following data (from Glinert 1989; Simon Mauck p.c.):

tarbut	'culture'	məturbat	'cultured'
koxav	'star'	məkuxav	'starry'
pilpel	'pepper'	məpupal	'witty'
petam	'cow'	məputam	'fat'

a. What are the category and meaning of the derived word forms on the right? How were they derived?
b. How would you derive similar forms from the following bases?

lamed	'knowledge'	_____	'educated'
tipeʃ	'fool'	_____	'foolish'
kavod	'honor'	_____	'honored'

Exercise 2.5

English. Analyze each of the following words into their constituent morphemes. Show the order in which each word was derived, justifying your analysis.

a. rehospitalization
b. incomprehensibility
c. unpreparedness
d. disenfranchisement

Exercise 2.6

Consider the brief discussion on p. 90 on how children acquire inflectional contrasts in context. Try to think of scenarios by which the child could do this. For example, children might acquire case distinctions by observing a

noun spoken in different contexts in different roles in relation to an event, e.g. 'The girl-NOM was patting the dog-ACC' vs. 'The girl-NOM pulled the dog-GEN tail-ACC,' vs. 'The dog-NOM bit the girl-ACC.' Can you create some contexts for acquiring other inflections, such as tense, aspect, number, gender and mood?

Exercise 2.7

Albanian. Consider the following data (from Camaj 1984) and answer the questions below:

sjellim	'we bring'
sillesh	'you (sing.) are brought'
posjell	'I am bringing'
sillem	'I am brought'
posillem	'I am being brought'
dotəsillem	'I will be brought'
sillemi	'we are brought'
dotəsjelləsh	'you (sing.) will bring'
letəsjelləsh	'you (sing.) should bring'
dotəsjell	'I will bring'
posjellim	'we are bringing'
dotəsjellim	'we will bring'
posillemi	'we are being brought'
dotəsillesh	'you (sing.) will be brought'
letəsjellim	'we should bring'

a. Give the likely Albanian forms for the following:

'should' _____
progressive morpheme _____
future morpheme _____

b. What are the two stem forms for 'bring' found in these data and what is the difference in their meaning?

c. Give the likely Albanian forms for the following:

'I bring' _____
'you are being brought' _____
'we will be brought' _____
'I should bring' _____

Exercise 2.8

Spanish. Study the following data and answer the questions that follow:

breve	'brief'	la brevedad	'briefness, brevity'
corto	'short; bashful'	la cortedad	'shortness; shyness'
cruel	'cruel'	la crueldad	'cruelty'
enfermo	'ill, sick'	la enfermedad	'illness'
(el) hermano	'brother'	la hermandad	'brotherhood'
improprio	'improper'	la impropriedad	'impropriety'
leve	'light, trivial'	la levedad	'lightness'
liviano	'fickle'	la liviandad	'fickleness'
mal	'evil'	la maldad	'wickedness'
solo	'alone, solitary'	la soledad	'solitude'

vario	'various'	la variedad	'variety'
(la) viuda	'widow'	la viudedad	'widowhood'

a. Explain how the words on the right are derived from those on the left, accounting for any allomorphy. What is the general meaning of the derived form?
b. Why do you think all the derived words are feminine gender? (Hint: see Box 2.7.)
c. Provide the likely Spanish derived forms for the following:

ebrio	'drunk'	_____	'drunkenness'
igual	'equal'	_____	'equality'
vasto	'vast'	_____	'vastness'
(el) vecino	'neighbor'	_____	'neighborhood'

Suggestions for further reading

Bauer, L. 2003, *Introducing linguistic morphology*, 2nd edition, Washington, DC: Georgetown University Press. This book is a well-written introduction to morphology with a good balance between data and theoretical issues. The second edition includes new chapters on psycholinguistic approaches to morphology and morphological change over time within a language.

Katamba, F. 1993, *Morphology*, New York: St. Martin's Press. Although this book is currently out of print and riddled with typographical errors, it is nonetheless still the clearest, best organized introduction to morphology around, with lots of interesting data and exercises. Borrow it from your library!

Pinker, S. 1994, "Words, words, words," *The language instinct* (chapter 5), New York: William Morrow. This chapter from Pinker's popular book on language is accessible, entertaining, and filled with many interesting facts, such as: the word *glamour* comes from the word *grammar*; English verbs have four inflectional forms (e.g. *quack*, *quacks*, *quacked*, *quacking*), whereas Turkish verbs have around two million; and the average English-speaking high school graduate probably knows around 60,000 words – about four times more than the number of words used by Shakespeare.

Spencer, A. 1991, *Morphological theory: an introduction to word structure in generative grammar*, Oxford and Malden, MA: Blackwell. At a more advanced level, this is one of the most comprehensive morphology textbooks available. It provides a good overview of various schools of thought and theoretical debates in the field. The exercises are also more challenging.

CHAPTER

3 The structure of sentences

DAVID LIGHTFOOT AND RALPH FASOLD

CHAPTER PREVIEW

KEY TERMS

adjunction
anaphor
aux(iliary) movement
Binding Theory
complement
complementizer
complementizer phrase
compositionality
continuous (discontinuous) reference
coordination
determiner phrase
formal syntax
functional category
functional syntax
gender
head
inflection
inflection phrase
language organ
lexical category
merger
name
poverty of the stimulus
problematic (unproblematic) reference
projection
pronoun
recursion
specifier
Universal Grammar (UG)
WH-movement
WH-question
zero anaphora

In order to understand the subtleties of sentence structure, it is necessary to understand how phrases are built from the words they contain, how phrases are combined into larger phrases and sentences. It is also necessary to understand what can happen to phrases and sentences after they are built – namely, parts of them can be moved and deleted. Movement and deletion take place under particular restrictions, and speakers "know" these restrictions, apparently without this being taught to them. All languages share these fundamental structural properties, but the principles that underlie them are broad enough to allow considerable differences among languages. The chapter includes a sampling of these differences.

We are presenting the structure of sentences with a primary emphasis on their form. However, there has been considerable research about the effect on the structure of sentences that comes from its use by language users. The chapter concludes with an example of this research and how it compares with a more formal approach.

GOALS The goals of this chapter are to:

- explain how sentences are constructed
- explain the concept "poverty of the stimulus"
- explain the notions "language organ" and "Universal Grammar"
- present examples of subtle restrictions that limit the ways in which sentences can be constructed and interpreted
- present a few examples of differences in sentence structure in languages from around the world
- present the differences between formal and functional analysis of sentence structure

Poverty of the stimulus

Many linguists argue that the capacity to acquire key aspects of natural language is exclusively human. Human language is built on an intricate foundation of grammatical principles. People don't have to learn a lot of what they know about these principles. The grammatical principles we are talking about don't have much to do with the grammar you learn in "grammar school." In fact, you have probably never heard of them, precisely because apparently they don't have to be taught. What children know about language goes beyond what they should be able to derive from what they hear, and very far beyond anything they are explicitly taught. The idea that people display a knowledge of grammar that is deeper than what they could get from the evidence around them is called the **poverty-of-the-stimulus** argument.

The amazing robot basketball player

To get a clearer idea of what poverty of the stimulus means, imagine this science-fiction tale. Suppose somebody found an exquisitely anthropomorphic robot in a remote part of a desert. The scientists who examined the robot eventually figured out how to activate it. When they did, the robot behaved like a person, but seemed very puzzled about the world around it, as if it had never seen it before. The scientists decided to try an experiment, allowing the robot to watch several basketball games. (Please don't ask why they would do that!) The robot watched a dozen or so games very intently. These games were well played, with players seldom committing infractions, and then only unintentionally. However, the games were loosely officiated, with some violations called, but most not.

The robot then was suited up and put in a game. Much to everyone's amazement, it proved to be an excellent basketball player who knew the game very well. At first, it seemed that the robot was simply a very careful observer. It was doing everything that it had seen the players do. But then a particularly observant scientist noticed something odd about the

robot's game. Although it did most of the things it had seen the players do, it never broke any rules. For example, it never started to dribble again after it had stopped, and it never tried to run with the ball without dribbling. The scientist thought this was notable because, although the robot had rarely seen human players doing these things, it would be very advantageous to do them. Why did the robot conclude that "the double dribble" and "traveling" were not allowed, rather than assuming that the human players simply weren't taking advantage of some very useful strategies?

Perhaps, thought the scientist, the robot concluded that anything it hadn't seen was against the rules, but this turned out not to be true. The robot had never seen anyone dribble the ball behind his back, but on several occasions it did just that. How did it know this is allowed? Even more remarkably, on two occasions it leaped into the air and fell on its back while slapping the ball forward to keep it from going into the defensive half of the court after its team had brought it into the offensive end. This prevented a back-court violation, but how did the robot know? In the games it had watched there were never any back-court violations.

The scientist decided to interview the robot. She set up a monitor and showed a carefully selected video of a number of basketball plays – some showing violations of the rules and some showing unusual but legal plays – and asked the robot if what it saw was against the rules or not.

What happened next was remarkable. The robot accurately pointed out the plays that were allowed and the ones that were not, with almost no errors. In the case of rule violations, the scientist would ask the robot what rule was violated. But the robot could not tell her the rule, just that the play wasn't allowed. When the scientist asked, the robot affirmed that it had never seen anyone play basketball before the experiment. The scientist reluctantly came to a conclusion that was as unavoidable as it was improbable. *The robot had been preprogrammed with the rules of basketball without being aware of it!*

Applying the metaphor to the structure of sentences

It is just this sort of reasoning that has led many linguists to the conclusion that people are "preprogrammed" with principles of grammar. In the first few years of life, children develop cognitive systems that characterize the structure underlying their linguistic capacity. We call these systems "grammars." A grammar, in this view, is biological – a **language organ** (Anderson and Lightfoot 2002). We don't know just how a grammar is *physically* represented in an individual's brain. But the systematic behavior that people display when they speak, and when they are asked whether or not certain structures are allowed, makes it possible for us to describe the grammar quite specifically. Just as the robot became intuitively aware of the rules of basketball as it watched some games, the grammar *emerges* when children are exposed to particular experiences, and its emergence is closely guided by genetically

encoded principles. For example, English speakers have grammars that allow *Kim loves herself*, but not *People around Kim love herself.* English speakers can tell you that the first example is fine and the second isn't, but they can't tell you why. If they try, they are likely to get it wrong. They almost certainly were not taught anything about these examples in school.

Compositionality

Grammars, in this sense, have certain defining properties. A fundamental property of the grammar of every language is that it is **compositional**: sentences are made of clauses and phrases, which in turn are made up of smaller clauses and phrases or words.

Projection

Composing the structure of a sentence begins with words that belong to categories like noun, verb, preposition, adjective, and determiner. These words are the **heads** of phrases. Let's take a sentence like *Those children want a puppy* as an example. In this sentence, *those children* is a determiner phrase (DP). Its head is the determiner *those* and it also includes the noun phrase *children*. *A puppy* is another determiner phrase; its head is the determiner *a*, and it includes the noun phrase (NP) *puppy*. *Want a puppy* is a verb phrase (VP) headed by the verb *want* and it includes the determiner phrase *a puppy*. *Children* and *puppy* are phrases too. They are noun phrases that consist of their heads only. The sentence is composed of phrases and the phrases are composed of words (even though some phrases are composed of only one word).

Current research on grammar assumes that phrases are constructed "bottom-up": words are drawn from a **lexicon** – a mental dictionary stored in people's brains rather than in a book – and **merged** into structures one-by-one. Here's how this works in detail. Every word is a member of a category. There are two types of categories, **lexical categories** and **functional categories**. Words that belong to lexical categories are semantically rich and contribute primarily to the meaning of the sentence. The common "parts of speech" – nouns, verbs, adjectives, adverbs, and prepositions – are lexical categories. Words in functional categories are semantically weak, and contribute more to the structure of the grammar of sentences than to the meaning. Articles are a subset of determiners, a functional category.

To build phrases and sentences, we start with words, which then **project** phrases of the same category. You can think of a word as a kind of seed and its projection as a stalk that it sends out. Let's start with the simple sentence *Those children want a puppy* and see how this works. The grammar starts by selecting the nouns *children* and *puppy* from the lexicon and letting them project phrases of the same category – i.e. noun phrases (NPs).

Box 3.1 Table of grammatical categories

Type	Word category	Categories: Word-level symbol	Examples	Phrasal category	Phrase-level symbol
Lexical	Noun	N	puppy, park, happiness	Noun Phrase	NP
Lexical	Verb	V	take, run, be	Verb Phrase	VP
Functional	Determiner	D	this, the, a	Determiner Phrase	DP
Lexical*	Preposition	P	in, on, between	Prepositional Phrase	PP
Lexical	Adjective	A	good, red, big	Adjective Phrase	AP
Lexical	Adverb	Adv	happily, fast, initially	Adverb Phrase	AdvP
Functional	Inflection	I	can, PRES, PAST	Inflection Phrase	IP
Functional	Complementizer	C	that	Complementizer Phrase	CP

* Preposition is sometimes considered a functional category.

In the sentence we are constructing, these noun phrases have no further structure within them, so the grammar can move on to other categories.

Want is a transitive verb – a verb that isn't complete without a direct object. So if the grammar is projecting a phrase from a transitive verb, it sends up a stalk that branches. The projection from *want* is shown in Figure 3.2. The figures we are constructing are called **tree diagrams** because, as they become more detailed, they begin to look like upside-down trees.

In other words, *want* projects a phrasal node of its own category (a verb phrase, or VP) and also branches out to another node for the category of its **complement**, the phrase it needs to be complete. (We show the complement phrase as a determiner phrase (DP) here, but it could also be an NP.) In a transitive verb phrase, the verb is the head of the phrase.

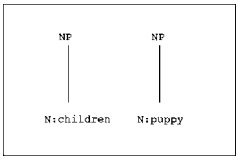

FIGURE 3.1
Projection of two simple NPs

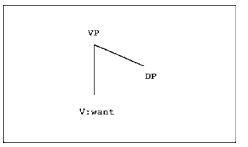

FIGURE 3.2
Projection of a transitive VP from *want*

Merger

Once the complement node of the transitive verb is projected, it has to be combined with a phrase of the type it needs. To do that, the grammar must have a mechanism that combines phrases. This is done by **merging** one phrase with another.

Each of the determiners in our example – *those* and *a* – projects a determiner phrase, and each of these determiner phrases needs a noun phrase as its complement. The grammar has already projected the two NPs in

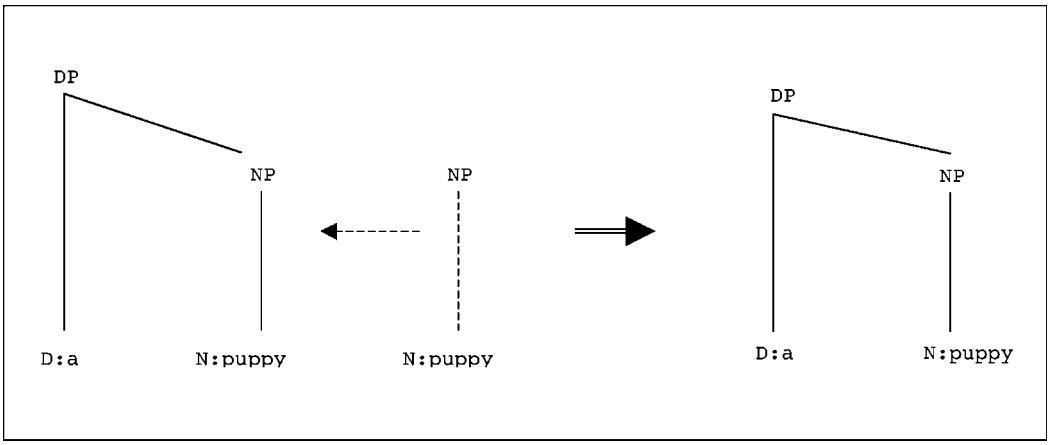

FIGURE 3.3
NP *puppy* merges with the complement NP node projected from D *a* to create DP *a puppy*

Figure 3.1. All it has to do is merge the two NPs into the two DPs, and the DPs will be complete. The way this works is illustrated in Figure 3.3. The same kind of merger produces *those children*.

Notice that at the beginning of a merger, there are two NP nodes, but only one NP afterwards. In a merger, the two nodes blend together to form one. Recall that in Figure 3.2, the verb phrase whose head is *want* needed a DP complement, which can now be merged into the VP as illustrated in Figure 3.4.

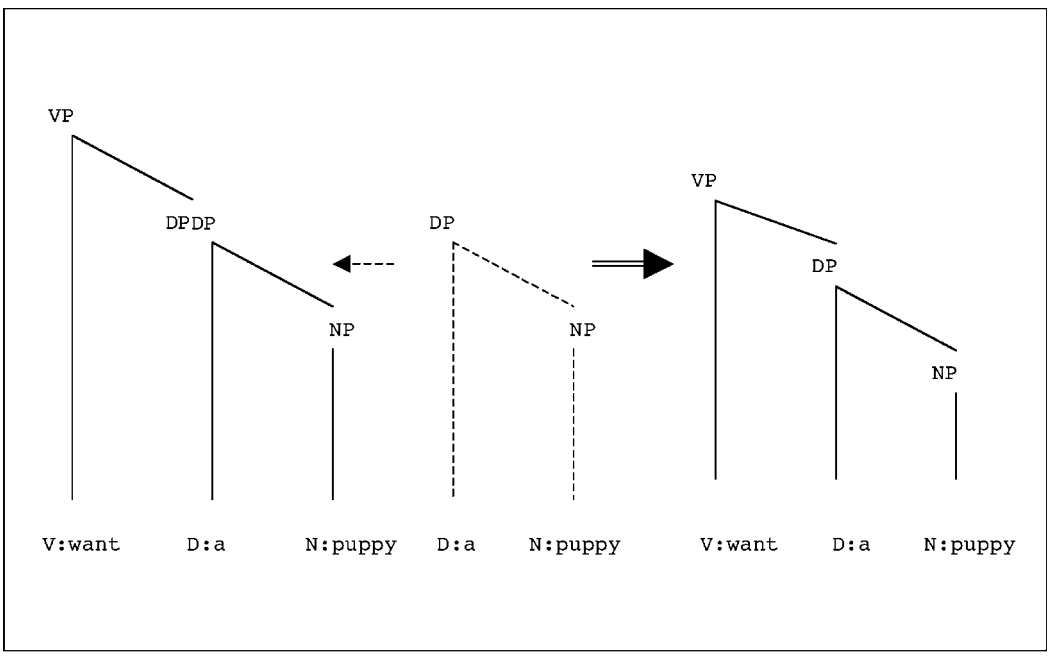

FIGURE 3.4
DP *a puppy* merges with complement DP node projected from V *want* to create VP *want a puppy*

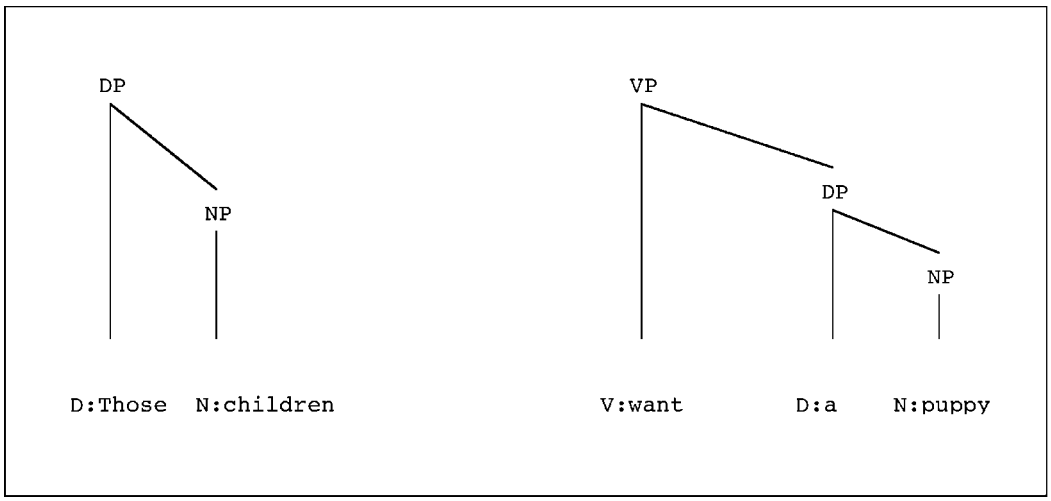

FIGURE 3.5
Two phrases resulting from mergers

So far the grammar has created the two phrases in Figure 3.5 by merger. A new functional category is needed to put them together. This category, called **inflection**, is responsible for the tense of the sentence (among other duties).

An inflection phrase (IP) is projected from abstract elements, such as present and past tense, as well as modal auxiliaries, like *can* or *may*. Its role in the structure of sentences is somewhat different from the role of the phrases we have projected so far. Inflection phrases provide the central "scaffolding" for a sentence, a structure to which the more meaningful lexical category phrases will be attached. In our sentence, the head of the inflection is not a word you can hear, but the abstract element PRES (present tense). The resulting structure is shown in Figure 3.6.

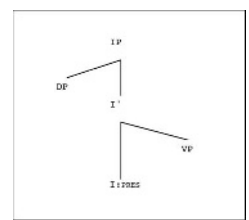

FIGURE 3.6
Projection of IP from PRES

Even for a functional category, inflection is special. Like transitive verbs and prepositions, it projects a complement branch; to be complete it has to have a verb phrase attached. In other words, you can think of the verb phrase as analogous to the object of a prepositional phrase; in our example, it provides a way for present tense to express its meaning. But inflection phrases are unusual because they are also required to project a **specifier position**. A specifier is an extra level of structure with a particular configuration. There is an intermediate node between the lowest category (I) and the IP. This intermediate node is called I′, pronounced "I-bar." The complement verb phrase branches from I′ and the specifier branches in the opposite direction from the higher IP. This upper piece of structure is called a "specifier" for a reason that was explainable in earlier versions of syntactic theory, but now is simply a conventional technical term. The specifier of an inflection phrase is the subject of the sentence, here a determiner phrase. Inflection phrases typically have determiner phrases as their specifiers, although other phrases are possible.

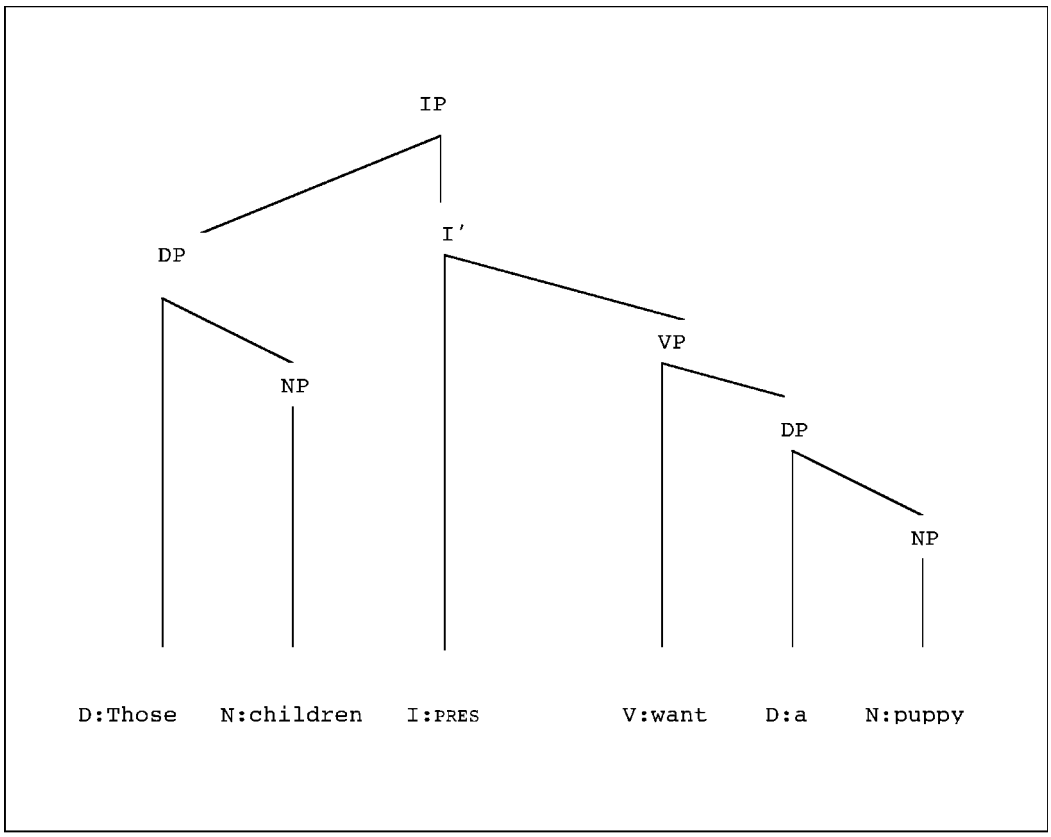

FIGURE 3.7
Merger of DP and VP into the specifier and complementizer positions of IP

Notice that in Figure 3.6, DP and VP have no content. To give them content, the grammar merges the structures in Figure 3.5 with the one in Figure 3.6. The DP containing *those children* is merged with the DP in the specifier position of IP (this is our first example of specifier merger) and the VP containing *want a puppy* is merged with the VP. The inflection phrase combines and relates the content of the DP and the VP. The result is shown in Figure 3.7.

The generation of the sentence by the grammar is almost complete. There has to be a mechanism to get the present tense (PRES) attached to the verb. In this sentence, PRES is not audible so you can't hear whether it attached or not. But if our sentence had been *My child wants a puppy*, we could have seen the present-tense marker in the suffix *-s*. Similarly, if we had chosen the past tense, the sentence would have come out *Those children wanted a puppy*, also with the tense-marking suffix on the verb. The mechanism that combines verbs and their tense turned out to be surprisingly complicated for syntactic theory, so we'll just state here that there is a way to get tense inflections where they belong.

Our grammar must have one more functional category, the **complementizer**. Complementizers are words like *that* in *I heard that those children want*

a puppy, *for* in *She hopes for good things to happen*, and *if* in *He wondered if it would rain*. Complementizers project high-level phrases that take inflection phrases as complements. Complementizer phrases (CPs) are functional categories that allow clauses to be embedded in other clauses. For example, in a sentence like *I heard that those children want a puppy*, the clause *that those children want a puppy* is embedded in a higher clause as the complement of the verb *heard*. Complementizer phrases also are necessary to understand the structure of questions and relative clauses (e.g. *The man who came in was angry*), as well as indirect quotations (e.g. *He said that those children want a puppy*). A Complementizer phrase looks like Figure 3.8.

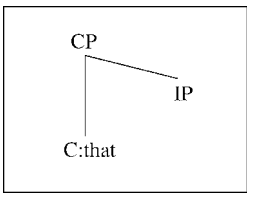

FIGURE 3.8
Projection of CP from *that*

In a sentence like *I heard that those children want a puppy*, the complementizer phrase would look like Figure 3.9.

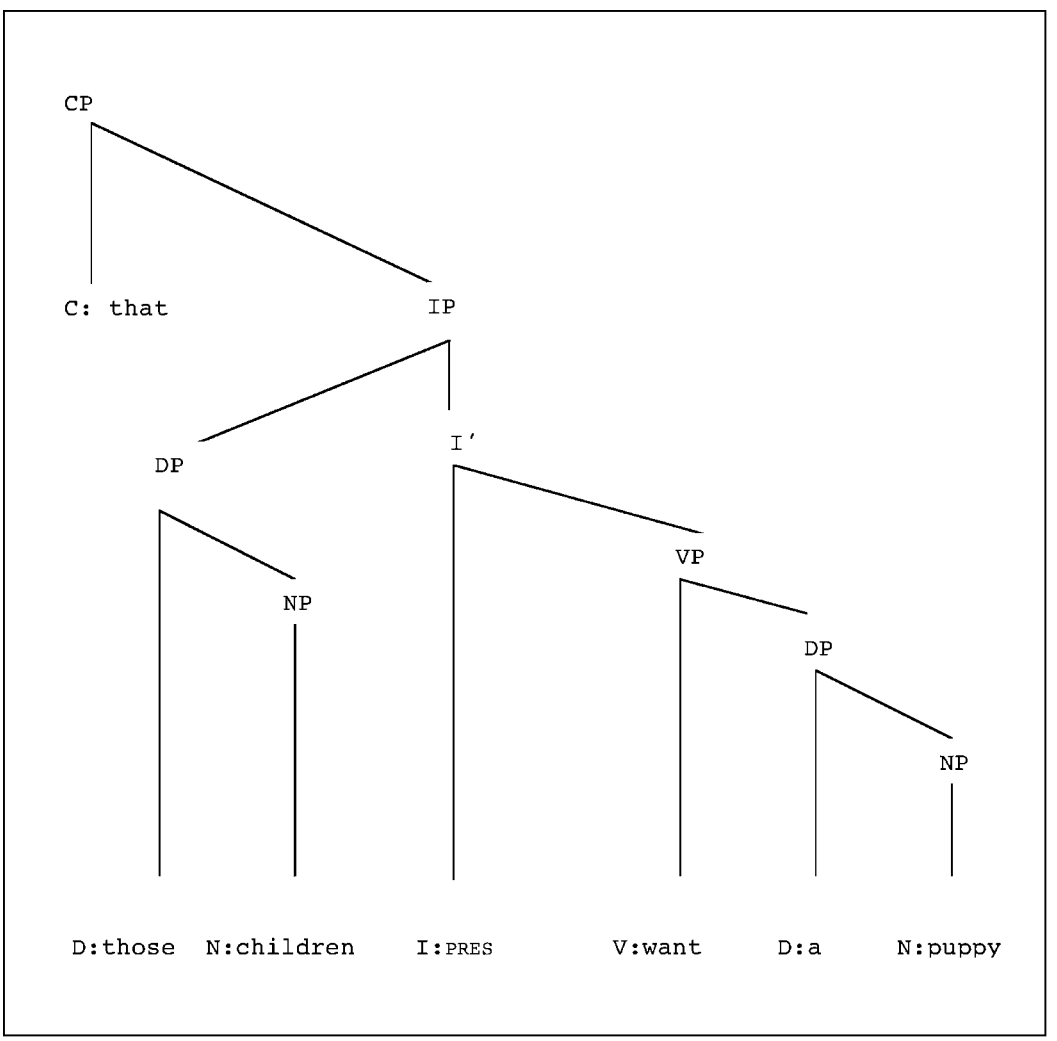

FIGURE 3.9
(I heard) that those children want a puppy

Box 3.2 Determiner phrase and possession

We have introduced three categories that we have called **functional**. The functional nature of inflection phrases is easy to see, since they serve as the "scaffold structure" for sentences, holding the inflection and the positions for the subject determiner phrase and the main verb phrase. Similarly complementizer phrases have the function of providing the structure for embedded clauses. But what about determiner phrases? From the examples we have shown, determiner seems just like a lexical category, projecting a phrase and a complement branch from a head with some semantic content.

One clearly functional duty of determiner phrases is to provide the scaffolding for possessive constructions, like *The cowboy's hat*. This possessive construction is a DP projected from an unusual determiner, the possessive morpheme spelled *'s*. The possessive determiner projects an NP complement, as usual, and this time it also projects a specifier (like inflection phrases do). The projection looks like this:

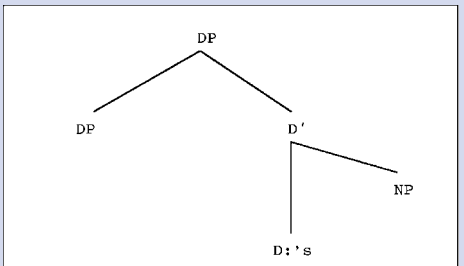

Projection of a possessive DP from *'s*.

As you can see, this structure looks very similar to a typical IP structure. When the DP *The cowboy* is merged into the specifier position and the NP *hat* into the complement position, we get:

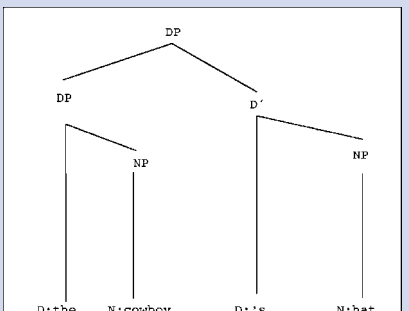

Structure of *the cowboy's hat*.

Adjunction

There are two methods for building phrases: projection and merger. We have seen how merger joins phrases by placing one phrase into the complement or specifier of another phrase. It is also possible to merge phrases that are not complements or specifiers of their host phrase. This is called **adjunction** and it adds *modifiers* to phrases. Heads, complements, and specifiers make up the core meaning of a phrase, while adjuncts add extra description. Since adjoined phrases are different from heads, complements, and specifiers, adjunction creates a site for merger by extending the phrasal node of the host phrase.

Figures 3.10, 3.11, and 3.12 illustrate adjunction. Suppose we want to add modifiers to the subject noun phrase and verb phrase of our example sentence so that it reads: *Those little children in the park want a puppy badly.* We would **adjoin** the (minimal) adjective phrase *little* to the noun phrase *children*. First, the grammar would project a new adjective phrase (AP) from the adjective *little*. The grammar would create the "roof-shaped" structure circled in Figure 3.10 that extends the NP *children*. The

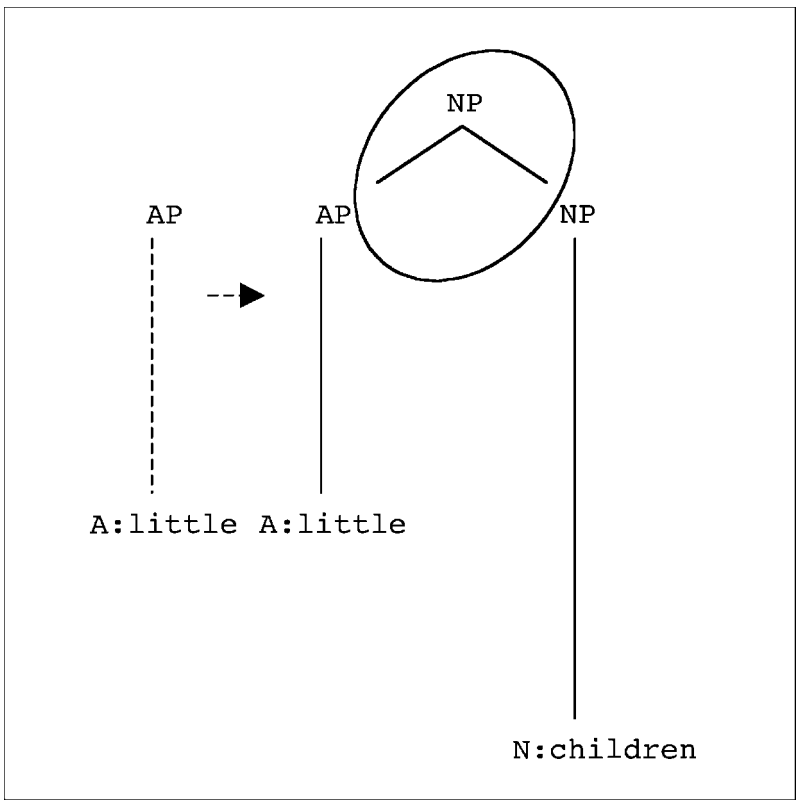

FIGURE 3.10
(Those) little children (in the park want a puppy badly)

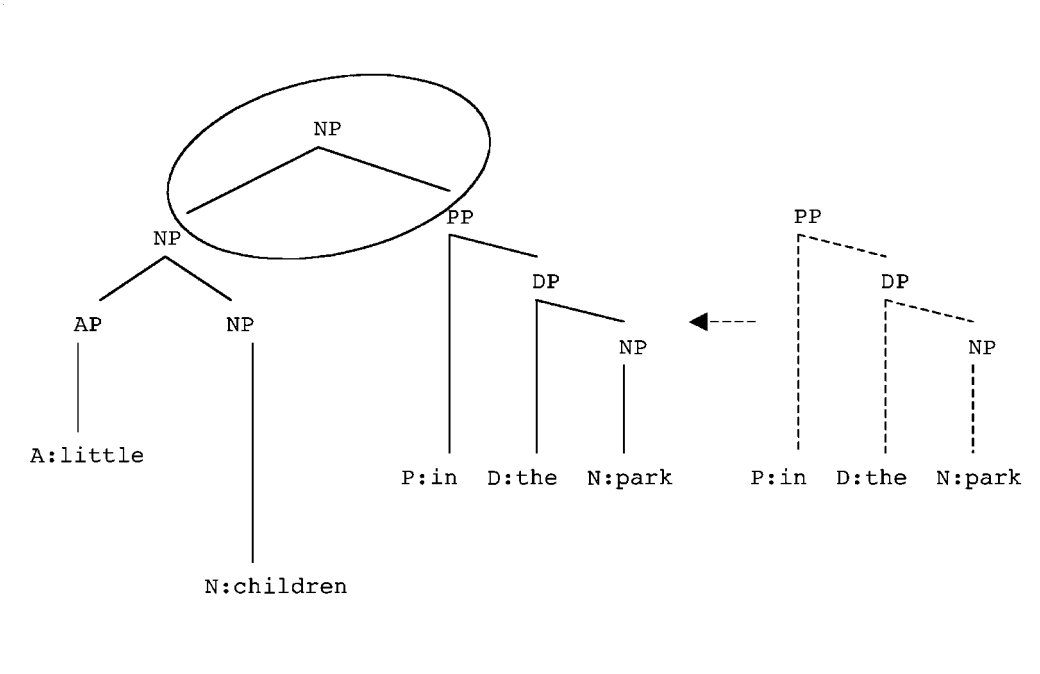

FIGURE 3.11
(Those) little children in the park (want a puppy badly)

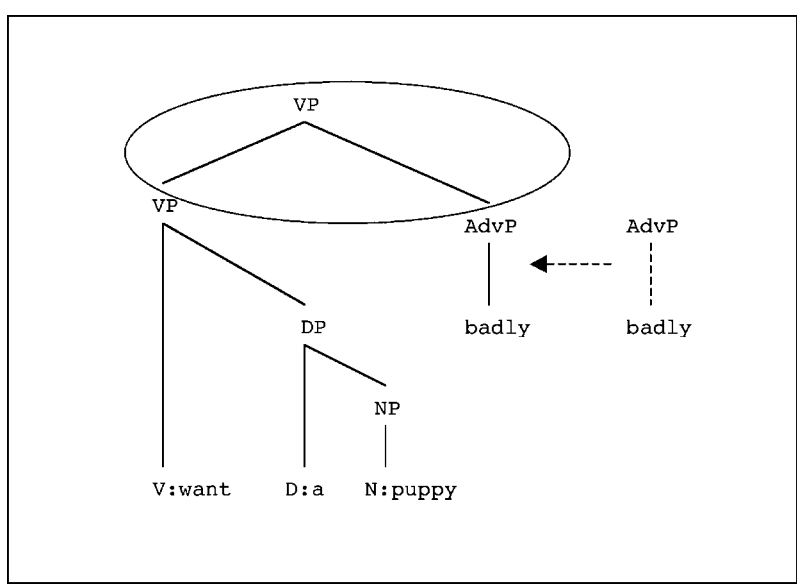

FIGURE 3.12
(Those little children in the park) want a puppy badly

new AP would be merged into this new structure. This is shown in Figure 3.10.

The merger of *in the park* is similar. The noun *park* would project an NP node which would be merged with the DP projected by the determiner *the*. This DP would in turn be merged with the prepositional phrase (PP) projected from the preposition *in*. The NP in Figure 3.10 would be further extended with a new "roof-shaped" structure, and *in the park* is merged with that. Figure 3.11 shows the process.

Similarly, the adverb *badly* would project an adverb phrase (AdvP), which the grammar would then merge with the verb phrase *want a puppy*, using the same mechanism. Figure 3.12 shows the merger of the adverb phrase *badly* with the VP *want a puppy*. Notice that in Figures 3.10 and 3.11, only the head is attached to the lowest NP node and in Figure 3.12, only the head and the complement are attached to the lower VP node. All the adjuncts are attached to the extended structures. In this way the integral parts of a phrase – head, specifier, and complement – are distinguished from adjuncts.

Building the rest of the structure for *Those little children in the park want a puppy* would proceed in the following way. First the NP in Figure 3.11 would be merged with the complement branch of the DP headed by *those*, as in Figure 3.3, to create the DP *those little children in the park*. The resulting DP and the VP in Figure 3.12, *want a puppy badly*, would be merged into IP as before to create the whole sentence (IP).

So far, we have represented phrasal structure using tree diagrams, which are graphically useful for showing the hierarchical relationships between words and phrases, but take up a lot of space. Linguists often use a linear format, called a **labeled bracketing**, to represent sentences, as in (1).

(1) $[_{IP}[_{DP}$ those children$][_{VP}$ want $[_{DP}$ a puppy$]]]$

In this format, brackets surround everything underneath a node, and the bracket is labeled with the same label a node in a tree would have. The entire sentence, within the outermost brackets, is an IP. Within the IP there is a DP (*those children*) followed by a VP (*want a puppy*). The verb phrase consists of *want*, which we have chosen not to label explicitly, and another DP (*a puppy*). When they use this format, linguists leave out much of the internal structure of each of the constituent phrases, unless it is relevant to a particular analysis.

Movement and deletion

Once phrases have been built by projection and merger (including adjunction) the grammar can apply further operations to them. Besides building phrase structures, syntax can also move parts of phrase structures around, by detaching them from the position in which they were projected or merged, and merging them somewhere else, or delete them. Movement functions by copying an item into a new location, leaving a copy in the

Why copy and delete?

It may seem strange that movement is done in two steps, copying and deletion of the original, instead of just moving an element. Syntacticians have discovered that grammar behaves as if something is left behind after a movement. That "something" seems to be a copy of what was moved. If movement is understood as copying and deleting, then it is a combination of merger (of something from somewhere else in the structure, instead of from the lexicon) and deletion, thus using two operations already needed rather than adding a third kind of operation, movement.

original position. This copy must later be deleted. We will illustrate movement with two common operations, **aux(iliary) movement and WH-movement**, which generate questions.

Auxiliary movement

Aux movement comes into play when we want to generate a simple question like *Does the man like movies?* Using the projection and merger functions we've already described, the grammar first generates the sentence in (2).

(2) [$_{CP}$ e [$_{IP}$ [$_{DP}$ the man] [$_{I'}$ PRES [$_{VP}$ like movies]]]]

Example (2) shows a complementizer phrase headed by an empty complementizer (that's what the "e" stands for), and it has an inflection phrase as its complement. The IP has the determiner phrase, *the man*, in its specifier position, its head is PRES, and its complement is the verb phrase *like movies*. Aux movement takes whatever I is (here, it is PRES) and moves it to the previously empty head position, C, leaving a copy behind. The copy is later deleted. The result is (3).

(3) [$_{CP}$ PRES [$_{IP}$ [$_{DP}$ the man] [$_{I'}$ ~~PRES~~ [$_{VP}$ like movies]]]]

The abstract present-tense morpheme is at the beginning of the sentence, far from its verb. Under this condition, the verb *do* is inserted into the structure so that the tense has a verb to work with. The result is (4).

(4) [$_{CP}$ do PRES [$_{IP}$ [$_{DP}$ the man] [$_{I'}$ e [$_{VP}$ like movies]]]]

Now, all that has to happen is for *do* and PRES to combine to form *does* and we have the sentence we want: *Does the man like movies?*

WH-movement

There's another kind of question called the **WH-question**. WH-movement, in questions, follows Aux movement. In WH-questions, there is a word like *who*, *when*, *where* or *how* at the beginning of the sentence. These WH-words start life within the sentence and are then moved to the specifier position of CP. The grammar would generate the question *What would you like?* as in (5).

(5) [$_{CP}$ e e [$_{IP}$ you [$_{I'}$ would [$_{VP}$ like what]]]]

The head of IP in (5) is the auxiliary word *would* (rather than the abstract PRES in (2–4)). We see that CP has two empty positions. The structure for CP that you saw in Figures 3.8 and 3.9 has only a head position, which is empty in (5), and a complement position, which is occupied by *you would like what* in (5). Where does the other empty position come from? For structures like this one, CP has a specifier position, just as IP

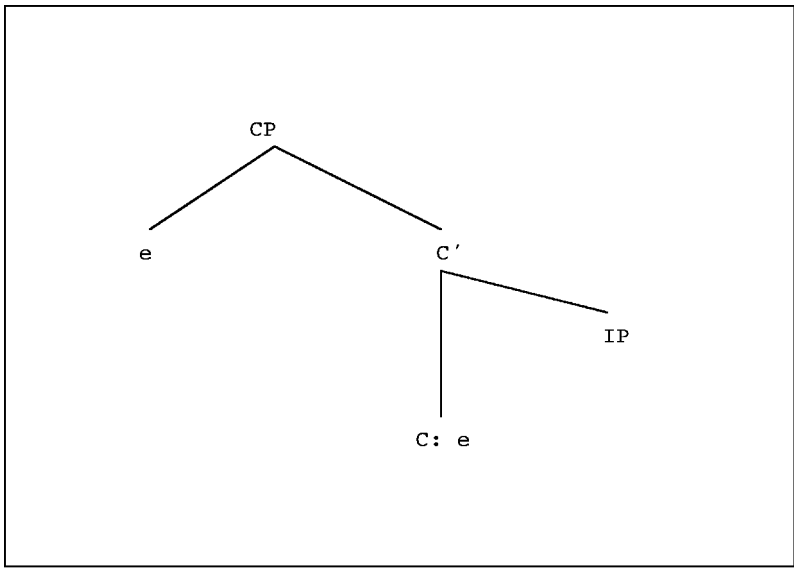

FIGURE 3.13
Complementizer Phrase with a specifier

always does, and as DP does in possessives (see Box 3.2). In principle, branching structures are never projected unless they are needed. No specifier is needed for the structure in Figure 3.9, so none is projected. Here, a specifier *is* needed. A CP with a specifier has the structure in Figure 3.13.

When a CP has a specifier, there is a C′ level like I′ in an inflection phrase and D′ in possessive DPs. The specifier position of CP can be filled by a variety of phrase types, but in Figure 3.13 it is shown as empty, as it is in (5). The empty position represented by the first of the two instances of *e* in (5), then, is the specifier position.

What appears after the verb *like,* as its complement. *Like* is a transitive verb and requires a direct object complement, and *what* stands in for the missing thing that *you would like*. First, Aux movement applies as it did in (3) above. This is shown by the shorter arrow, copying *would* to the head of CP. (The deletion of the remaining copy of *would* is assumed in this diagram.) This time, though, *would* is an audible morpheme, so we will have no need of *do*. WH-movement applies next, copying *what* to the specifier position CP, leaving the original NP position empty, once *what* is deleted. The result is (6).

(6) $[_{CP}$ what would $[_{IP}$ you $[_{I'}$ ~~would~~ $[_{VP}$ like ~~what~~]]]]

Grammars are finite; language is not

Although capable of remarkably complex processes, human beings' brains are finite. There are a limited number of cells in the human brain, and

therefore a limited (though large) number of connections between those cells. Since grammars are part of people's brains, grammars must be finite, too. But an individual human being has the capacity to understand and produce an infinite number of sentences. People say and hear completely new, or novel, sentences all the time. Take any sentence from the last chapter. Although you had never heard it, read it, or said it before in just that form, you were able to understand what it means. How can a finite grammar have an infinite capacity for producing and understanding sentences?

Recursion

Regardless of where we were raised and whether we grew up in some English-speaking community or in Tokyo or Ulan Bator, speakers of every language have **recursive devices** – means by which the same grammatical processes can apply more than once – in their grammars that make it possible for them to produce sentences of indefinite length.

Multiple adjunction

One thing this means is that we can insert words of the same kind repeatedly with no principled limit by repeated adjunction. You can easily imagine someone saying, *She fell in love with this intelligent guy.* This is accomplished by adjoining the adjective *intelligent* to the noun phrase *guy*. But it would be possible to go a lot beyond that; you could say, *She fell in love with this really intelligent, handsome, considerate, romantic, thoughtful, adorable guy* simply by adjoining more adjectives. In principle, you could go on until you exhausted all the applicable adjectives in the English language, then you could continue by coining new adjectives. Of course, nobody ever does this, because they would get tired of talking and – more to the point – everybody else would get tired of listening. But grammar would not prevent you from going on and on, because of the recursion feature.

Embedding

A more complex recursive device allows any sentence to be placed in the larger frame of another sentence. Start with what you think is the longest sentence you can imagine; you can make it longer simply by putting *He said that . . .* in front of it. If your original sentence was *The woman in New York's dress was black,* it would be lengthened to *He said that the woman in New York's dress was black.* You could lengthen that sentence by using the same structure again. Even if you have already added *He said that*, you could add something similar: *She said that he said that the woman in New York's dress was black.* You could keep reframing your developing sentence indefinitely: *Bill knew that Jill thought that Tom said that. . . .* You might get tired, but your grammar would never limit the length of your sentence.

Another kind of embedding involves the use of relative clauses. Maybe you remember the children's song, "The House that Jack Built":

This is the house that Jack built.
This is the cheese that lay in the house that Jack built.
This is the mouse that nibbled the cheese that lay in the house that Jack built.

If you had the patience and lived long enough, you could string relative clauses together indefinitely: *This is the cow that kicked the dog that chased the cat that killed the rat that caught the mouse that nibbled the cheese that lay in the house that Jack built.*

Coordination

Another recursive device found in all languages is **coordination**. Coordination links two (or more) sentences (or phrases, or words) together on equal terms, using coordinating conjunctions like *and*, *but*, and *or*: *You can try or you can give up.* Sentences can be coordinated indefinitely as well: *Rick went to a movie, and Ellen went to the store, and Stuart worked, but Mike slept and Sue read a book . . .* These various recursive devices can be combined to make an even greater variety of sentences: *This is the house that Sue knew that Jack built and this is the mouse that Bill saw eating the savory, delicious, yellow, gourmet cheese . . .*

The significance of recursion

Because grammars have recursive devices which permit expressions to be indefinitely long, they can produce an infinite variety of sentences. Since recursion is an integral part of the grammar, it follows that no one can learn a language by memorizing all the sentences of that language. There must be some other explanation for how human beings are able to learn them.

Restrictions

It turns out that deletion and movement are restricted in quite subtle ways. It turns out that English speakers "know" the restrictions on deletion and movement in a way analogous to the fictitious robot who "knew" rules of basketball that it had never been taught. The combination of watching basketball and the robot's internal programming made it possible for it to play the game by the rules. In the same way, English speakers' experience with English in childhood, combined with the linguistic genotype, **Universal Grammar,** that they have inherited, makes it possible for them to show knowledge of the language, including its untaught principles.

You can do without *that,* but not always

Let us take one example. In English, complementizers like *that* are optional, unlike equivalent words in Dutch, French, Spanish, Italian, and other

languages, where complementizers are required. So a child raised in an English-speaking home might hear the sentences of (7a–c) pronounced with or without the complementizer *that.* These experiences would show them that *that* is optional so they will learn a complementizer-deletion rule like (7d).

(7) a. Peter said [$_{CP}$ that/0 [$_{IP}$ Kay left]].
b. Kay doesn't believe [$_{CP}$ that/0 Ray is smart].
c. It was obvious [$_{CP}$ that/0 [$_{IP}$ Kay left]].
d. That → 0

So far, we can see that children learn different things about their native languages, depending on the input they are exposed to. However, children also know things about their language that they could never have learned in this way. In fact, *that* is not always optional; it is required in the contexts of (8), where the asterisk next to the zero means that deletion is not possible.

(8) a. Peter said yesterday in Chicago [that/*0 Kay had left].
b. Fay believes, but [$_{IP}$ Kay doesn't [$_{VP}$ [$_{V}$ e [$_{CP}$ that/*0 Ray is smart]]]].
c. [that/*0 Kay left] was obvious to all of us.

Something in the English grammar – and therefore in the English speaker's brain – requires the complementizer *that* in these kinds of sentences. But no adult or child ever says *Kay left was obvious to all of us*, so no one can point out that this is an ungrammatical sentence in English. This is called **negative evidence**, information about what does *not* occur in a language. They hear sentences like (8) only with complementizers, but they can't possibly learn from what they hear that the complementizer is *not* optional in just the contexts in (8). Children acquiring English "know" that the cases in (8) are impossible, but since they don't have evidence to learn this from, this knowledge must be coming from somewhere else. A straightforward answer is that the mental organ for language – which many linguists call Universal Grammar or UG – must be playing some role. This conclusion matches the scientist's conclusion that the robot must have been programmed with basketball rules, because it hadn't seen enough evidence to learn them by watching basketball games.

The facts that we are presenting may seem too commonplace to require explanation. "Of course you can't say 'Kay left was obvious to all of us'," you might think, "Anybody can see that." But in Isaac Newton's time, everybody knew that apples fall down, and nobody thought it was anything that needed to be explained. Newton advanced the field of physics and made a place for himself in history by asking "Why?" Linguists are doing the same thing when we ask questions about *that*-deletion. The results we get are surprising in the same way that the discovery of gravity was, though they are not yet as widely known.

The problem of *that* complementizer deletion is that *that* can readily be deleted in English in some contexts, but not in others. To see why, let's

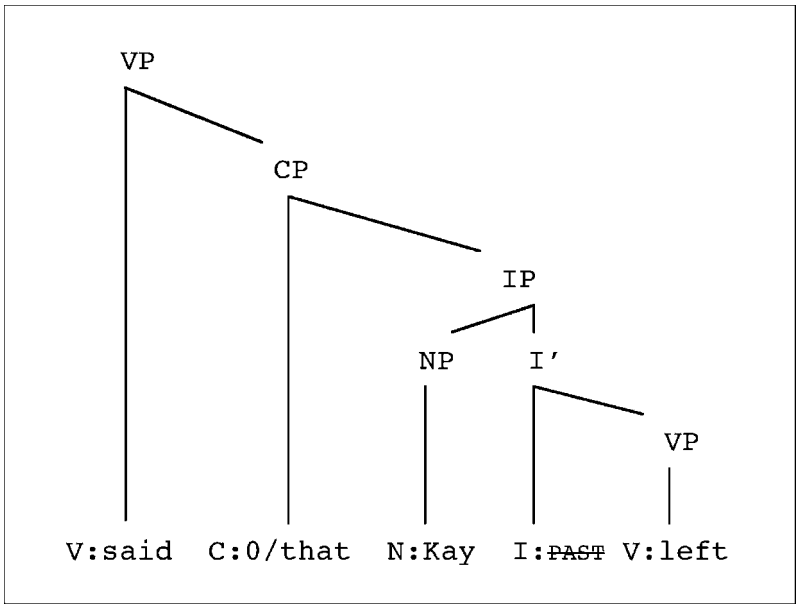

FIGURE 3.14
That Key left as part of the complement of *said* and adjacent to *said*

consider (7a) in detail. Figure 3.14 is the verb phrase part of (7a). We're assuming that the abstract inflection PAST has already become part of the verb *leave* giving the past-tense form *left*. The same thing has happened with *said*. If you look carefully at Figure 3.14, you will see that [$_{CP}$ (*that*) [$_{IP}$ *Kay left*]] is the complement of the verb phrase, and also directly adjacent to the head of that verb phrase, *said*.

The complementizer *that* can be omitted if its clause is the *complement* of an overt, adjacent word. An **overt** element is pronounced, like *that* in *Peter said that Kay left*, as opposed to *Peter said Kay left*, in which the complementizer is structurally "there," but not pronounced. The same is true of the other examples in (7), where *that* deletion is possible. In (8a), on the other hand, the clause is still the complement of *say*. But the complement is not adjacent to its head, *say*. *Chicago* is not the head of the verb phrase and the CP containing *Kay left* is not the complement of *Chicago*. In (8b) the CP containing *Kay left* is the complement of the elided but understood verb *believe*, which is not **overt** (the *e* again means empty), and in (8c), the clause is the complement of nothing. Children cannot learn all this, because they would have to be able to use data about what does not occur (negative data).

If we just look at this narrow range of facts, it seems that the following principle is part of Universal Grammar.

(9) *That* complementizers, if they are deletable, can be deleted if they are adjacent to an overt head and they are in that head's complement.

With this principle, we solve this particular poverty-of-stimulus problem. Children learning English learn that the *that* complementizer can be

deleted because they hear some kinds of sentences in which *that* may appear, but doesn't have to. In other types of sentences, like the ones in (8), they always hear *that*, never examples of its deletion. But the conditions under which *that* cannot be deleted depend on some very subtle properties of grammar, involving notions like a clause being a complement of some audible word, and nothing like this is ever part of children's experience. Since children are sensitive to it anyway, we are forced to conclude that they "know" this principle already.

Heavy determiner phrase movement

There are a number of other instances in which the principle about adjacency to a complement is operative, including "heavy" DP movement. This applies when the grammar moves to the right a long, complicated DP from a position where it would occur if it were shorter. For example, while it would be natural to say *I saw a bird on a twig*, if you wanted to describe that bird more fully, you might switch things around – *I saw on a twig a most amazing, multicolored, long-feathered bird* – moving the long complement of *saw* all the way to the right end of the sentence. Under the approach we are presenting, that means copying and moving the "heavy" DP and then deleting the copied element.

(10) a. I introduced [~~all the students from Brazil~~] to Mary [all the students from Brazil].
I introduced to Mary all the students from Brazil.
b. I expect [[~~all the students from Brazil~~] to be at the party][all the students from Brazil].
I expect to be at the party all the students from Brazil.
c. *[[**all the students from Brazil**] are unhappy] [all the students from Brazil].

In (10a) *all the students from Brazil* is the complement of *introduced*, and it can be deleted after it is copied and moved. In (10b), the resulting sentence is a bit awkward but still grammatical. It is grammatical because *all the students from Brazil* is adjacent to *expect*. In this example, however, the complement of *expect* is *all the students from Brazil to be at the party*, not just the copied phrase *all the students from Brazil.* But the copied phrase is adjacent to *expect* and is part of the complement, so it can be copied and deleted. This shows us that we should expand the principle in (9) so that it applies to more than just complementizers. We should refine (9) just a little, giving us (11).

(11) An element, if it is deletable, can be deleted if it is adjacent to an overt head, and it is either *in* that head's complement or it *is* its complement.

In (10c), on the other hand, the element which needs to be deleted (the clause in boldface type) is neither the complement nor contained in the complement of anything, so it cannot possibly conform to the principle in

(11). So the deletion is impossible and the result sounds very bad if we try: **Are unhappy all the students from Brazil.*

The problem illustrated by the example in (10c) thus becomes another poverty-of-stimulus problem to be solved by our UG principle about phrasal deletion. From what they hear, children simply learn that heavy DPs may be copied to the right. The existence of an innate UG principle, like that in (11), explains why children "know" not to produce sentences like (10c) with no further learning or experience needed. We are beginning to see how children attain a complex mature grammatical capacity through the *a priori* "knowledge" of one simple UG property interacting with some simple operations learned from exposure to simple utterances of a kind that any child hears.

We postulate that children can learn from their environment that complementizers may be omitted and that heavy DPs may be displaced. This much has to be learned, because it is not true of all languages. However, children learning English do not *learn* the restrictions. They don't have to learn that complementizers cannot be omitted in (8), or that *all the students from Brazil* cannot be displaced in (10c). They don't learn these things, because they are unlearnable. There is nothing relevant in the primary data that children have access to, the things they hear. Furthermore, they do not need to learn the restrictions, because they already "know" the UG condition represented in (11), which comes from their genetic heritage, not their experience. Postulating the simple UG condition (11) permits us to say that children learn a few simple things from their experience and, once they learn these things, they have a surprisingly complex capacity that enables them to omit complementizers and displace heavy DPs appropriately. And much more, as we shall see.

Other restrictions on movement and deletion.

Not all movements and deletions are subject to the restrictions we have described here. WH-movement under some circumstances, for example, is subject to restrictions we do not have the space to illustrate.

The Binding Theory

Linguists have long been intrigued by the constraints on where pronouns, reflexive pronouns, and full noun phrases can appear. These constraints are about how these elements can be interpreted, depending on where they appear in a structure, rather than being about movement or deletion. There are also many poverty-of-stimulus problems in the relations among pronouns, reflexive pronouns, and full noun phrases. Consider the facts in (12) and (13). The pronouns *she* and *her* may refer to Kim in (12a,b) but not in (12c,d).

(12) a. Kim_i loves $\text{her}_{i,j}$ mother.
b. Kim_i expected $\text{she}_{i,j}$ would win.
c. Kim_i expected her_j to win.
d. Kim_i loves her_j.

Sentences (13a,b) may be statements about one person named Kim (so the same index is possible), but (13c,d,e) may only be interpreted as statements about two people of the same name (distinct indices). We know this independently of context or of any aspect of the speech situation; we seem to know this simply from the form of the expression.

Use of subscripts to indicate coreference.

We express coreference by using the subscripted indices i and j; elements referring to the same person or referent will have the same subscripted index. In (12a,b) Kim *may refer to the same person (and have the same index i) as* her *and* she, *but* her *and* she *may also refer to some other (female) person, indicated by the index j. In contrast, in (12c,d)* her *must refer to somebody else than (and has a different index from)* Kim.

(13) a. Kim_i's father says $\text{Kim}_{i,j}$ is happy.
b. Kim_i's father loves $\text{Kim}_{i,j}$'s mother.
c. Kim_i loves Kim_j's mother.
d. Kim_i says Kim_j is happy.
e. Kim_i says Kim_j's mother is happy.

How does the form of the expression convey all this information? Why, for example, may *her* refer to Kim in (12a) but the two *Kims* in (13c) may not refer to the same person? Why may *she* refer to Kim in (12b) but not *her* in (12c)? Here is another area where children acquire a system which goes far beyond the input they receive. Again we have elaborate subconscious knowledge, which is not acquired through instruction of any kind; most readers would not have been aware of these distinctions until they read the last paragraph. A child may hear (12a) in a context where *her* clearly refers to Kim or in a context where *her* refers to another woman unnamed in this expression, perhaps the queen of England. On the other hand, sentence (12d) is heard only in a context where *her* refers to another woman – in fact to any woman other than Kim. Children come to know this without being supplied with evidence or being told that *her* cannot refer to Kim, unlike in (12a). If there is no learning here, then that would explain why we do not observe children making errors in the reference of names and pronouns (except in a very specific domain, which we will not discuss) – they show no signs of learning by trial-and-error. But how can we say that there is no learning? Is there an alternative?

Twenty-five years ago the well-known linguist Noam Chomsky proposed the **Binding Theory** as a solution to these poverty-of-stimulus problems (Chomsky 1981). The Binding Theory permitted a dramatic simplification of these facts, if we assume that it is a component of UG. The Binding Theory is stated, somewhat roughly, in (14). Like all components of UG, it is available to humans in *advance* of experience, actually enabling us to *interpret* our experience. The Binding Theory divides **nominals** (a cover term for nouns and pronouns) into three types: **anaphors** like *himself*, *themselves*, **pronouns** like *she*, *her*, *their*, and **names** (everything else).

(14) Binding Theory
Principle A: Anaphors are coindexed with the subject of their Domain.
Principle B: Pronouns are not coindexed with the subject of their Domain.
Principle C: Names are not coindexed with the subject of *any* clause located to their left in the sentence in which the name appears.

Each nominal has a Domain, which is roughly its clause. The Binding Theory determines whether each nominal is coindexed with the subject of its clause or a clause to its left, or not. If the nominal *is* coindexed with

the subject of its own or another clause to the left, it is said to be **bound**. That is how the Binding Theory gets its name.

(15) a. [$_{IP}$ ***[$_{DP}$ Kim's mother]*** washed herself].
b. [$_{IP}$ ***[$_{DP}$ Kim's mother]*** washed her].
c. [$_{IP}$ [$_{DP}$ Kim's mother] said that [$_{IP}$ ***[$_{DP}$ the doctor]*** washed her]].
d. [$_{IP}$ Kim said that [$_{IP}$ ***[$_{DP}$ the doctor]*** washed her]].
e. [$_{IP}$ ***Kim*** said that [$_{IP}$ ***[$_{DP}$ the doctor]*** washed Kim]].

Let's see how these three simple principles predict the facts of English nominals. For each sentence, you should check to see what each nominal can and cannot refer to. Next, we have to first determine which of the three principles in (14) applies, and then check whether the theory matches your judgments.

In (15a), who can *herself* refer to? *Herself* is an anaphor, so it should obey Principle A, meaning it is coindexed with (and refers to) the subject of its own clause – its inflection phrase (IP). The boldface, italicized determiner phrase in (15a) is the subject of this clause, so *herself* must be coindexed with *Kim's mother*. It can't refer to just *Kim* because *Kim* is only *part* of the subject. So the only possible interpretation of *herself* in (15a) is that it refers to *Kim's mother*.

In (15b), *her* may refer to *Kim* or to some other female person not mentioned in the sentence. *Her* is a pronoun, so Principle B applies. Principle B says that *her* cannot refer to the subject of its own clause. If *her* referred to *Kim's mother*, the subject of its clause, it would be *bound* within its Domain, like *herself* is in (15a). But the rule for pronouns is the opposite of the rule for anaphors; it is forbidden for pronouns to be bound within their Domains. So *her* cannot refer to *Kim's mother*. *Her* can very well refer to *Kim* in (15b), because *Kim* is only part of the subject. Simply put, the pronoun *her* in (15b) cannot refer to its own subject. Beyond that, it may refer to any female person that makes sense in context, even if she is not mentioned in the sentence.

(15c) introduces an additional complexity, an embedded clause. *The doctor* is the subject of the clause in which the pronoun *her* is located, so according to Principle B *the doctor* cannot bind *her*; they can't refer to the same person. Although *Kim's mother* is the subject of an IP, it doesn't matter, since the Domain for *her* is only the lower IP. Whatever relationships we find beyond the pronoun's IP don't matter. In fact, *her* could refer to *Kim's mother*, to *Kim*, or to some other female person. The Binding Theory doesn't tell us what *her* can refer to – just that it can't refer to the subject of its own clause – *her* is free to refer to anyone else. Similarly, in (15d), *her* may be coindexed with *Kim,* but doesn't have to be, because *Kim* is outside of the Domain of the pronoun *her*.

The considerations in (15e) are different. *Kim* is a name, so Principle C of the Binding Theory applies. Principle C is the only one that doesn't refer to a Domain. It says that names must not be bound by *any* subject to their left in a sentence, period. The DP *the doctor* is a leftward subject, so

it cannot refer to *Kim*. But the other *Kim* is excluded as well. It's not in the same clause as the lower *Kim*, but it is still a leftward subject. As a result, (15d) necessarily refers to two Kims; the lower *Kim* must be a different person than the higher *Kim.*

Box 3.3 Exceptions that prove the rule

Occasionally, you might encounter examples like the one in the Bo Nanas cartoon that clearly violate the Binding Theory. The last clause in what Bo says contains a violation of Principle B, that a pronoun cannot refer to the subject of its own clause. The last clause in the cartoon has the structure in (i).

(i) $[_{CP}\ [_{IP}$ I didn't $[_{VP}$ like me very much]]]

The pronoun *me* refers to Bo Nanas, who is also the referent of *I*. *I* is the subject of the clause containing *me*. The Binding Theory says that the anaphor *myself* is required here (and that *me* here cannot refer to the subject of its clause). But this exception actually *proves* the rule. Bo Nanas has sent himself a gift of leftover Halloween candy to make himself feel special. He wants it to seem that the package has been sent by someone else. His comment on the shabbiness of the gift is in the context of the pretence that the "I" who sent the gift is someone other than the "me" who received it. In that sense, Bo Nanas the sender is not the same person as Bo Nanas the receiver, and the cartoonist, John Kovaleski, exploits the Binding Theory to make that point. If there were no Principle B, the use of *me* would not be nearly so effective.

Notice also that (i) is the complement of the verb *thought*. *That*, which could have appeared adjacent to *thought*, has been deleted, as permitted by the principle in (11). The use of *myself* in the *if*-clause in the cartoon is an example of the adverbial use of an anaphor, which we don't cover in this chapter, although it turns out to conform to the Binding Theory.

The Binding Theory cannot be learned from the language that young children hear. It is part of UG, part of what children bring to their initial experience with language. As part of UG, children already "know" the three categories of nominals and the restrictions on their coreference. Children need only learn from language experience which words in their language fit into these categories – which words are anaphors, which are pronouns, and which are names. This they can do from positive evidence in the sentences that they hear. Once a child has learned that *themselves* is an anaphor, *her* a pronoun, and so on, all the appropriate indexing relations follow, with no further learning required. The result is a simple system that covers an infinite number of possibilities.

Box 3.4 Noam Chomsky

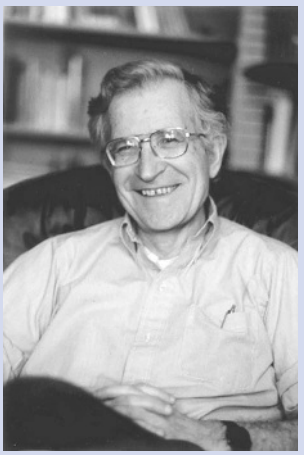

Noam Chomsky is the leading scholar in the study of the structure of sentences. He is responsible not only for the Binding Theory, but for developing the argument from the poverty of stimulus, the concept of Universal Grammar, and the entire approach to syntactic analysis that is central to this chapter. The appearance of his monograph *Syntactic Structures* in 1957 revolutionized the study of syntax. The ideas presented there formed part of a more extensive work, *The Logical Structure of Linguistic Theory,* published in 1975 but widely circulated over the previous twenty years. His current work, represented by *The Minimalist Program* (1995) and related work over the past ten years, focuses on deriving the fewest and simplest principles that can explain the unlearnable aspects of linguistic grammars.

Chomsky is currently Institute Professor in the Department of Linguistics and Philosophy at the Massachusetts Institute of Technology. He has lectured at many universities in the United States and in other countries, and is the recipient of numerous honorary degrees and awards. He has written and lectured widely on linguistics, philosophy, intellectual history, contemporary issues, international affairs, and US foreign policy. His most recent books are *A New Generation Draws the Line*; *New Horizons in the Study of Language and Mind*; *Rogue States*; *9-11*; *Understanding Power*; *On Nature and Language*; *Pirates and Emperors, Old and New*; *Chomsky on Democracy and Education*; *Middle East Illusions*; and *Hegemony or Survival.*

Summary

We have examined several areas of people's grammars and seen that we can account for complex arrays of distinctions between structures that people readily recognize as well-formed or ill-formed. This ability comes from simple information at the level of UG that interacts with simple, grammar-specific information that is readily learnable from a child's linguistic experience. The picture we have painted of children learning the grammars of languages is a more nuanced version of the fantasy about the basketball-playing robot. The robot didn't know anything about basketball until it experienced people playing the game. But once it had the opportunity to watch and play the game, it turned out to have knowledge of basketball that went beyond anything it could have learned from its experience. In other words, its knowledge of basketball was a poverty-of-stimulus problem. Just as the scientist concluded that there must have been some preprogramming in the robot's case, linguists who pursue the line of research we have shown you also conclude that children are born prepared in specific ways to learn languages. They are born with a biological language organ that linguists call Universal Grammar. If grammars have these kinds of properties and if we can discover more about them by

reasoning from poverty-of-stimulus problems of the kind illustrated here, then we can ask how such systems might be attained by children in the first few years of life. That is where the interests of syntacticians converge with those of people working on the language of young children.

Differences in syntax across languages

So far, we have focused on Universal Grammar and some of the principles that seem to be an intrinsic part of the language organ. On the other hand, within those principles, there is room for plenty of interesting differences in syntax from one language to another. We are now going to turn to some of those differences. Some of them touch on aspects of syntax that we have presented. For example, we have seen that complements appear to the right of their heads in English. That is not true of all languages. We discussed WH-movement in English, but not all languages have movable WH-phrases.

We are going to present another syntactic phenomenon that we haven't yet mentioned, but that turns up in different ways. This is the feature of some languages that places nouns in various categories, with syntactic consequences. We are going to call these **genders**, because that is what they are called in European languages that categorize nouns as feminine, masculine, and sometimes neuter. But languages elsewhere in the world put their nouns in many more genders than just two or three. We'll look at the ten or more genders in Swahili (also called Kiswahili) and the scores of genders in Thai.

Head–complement order in Hindi

In English, we saw that the complements of verbs are attached to the right. We didn't emphasize prepositional phrases, but they too have complements that branch to the right, as you can see in (16).

(16) $[_{PP}$ in $[_{DP}$ the house]]

In Hindi, verb phrases have complements branching to the left, as in (17a). There are no *pre*positional phrases in Hindi; instead Hindi has *post*positions with complements to the left, like verb phrases, as in (17b).

(17) a. $[_{VP}[_{DP}$ wah pillaa] caahnaa]
that puppy want

b. $[_{PP}[_{DP}$ ghar] mɛ̃ɛ̃]
house in

In Hindi, the structure of *The children in the house want that puppy* would be as in (18).

(18) $[_{IP}[_{DP}[_{PP}[_{DP}$ ghar] mɛ̃ɛ̃] baccee] $[_{VP}[_{DP}$ wah pillaa] caahtee hɛ̃ɛ̃]]
house in children that puppy want are

Notice that there are no overt words for *the* in Hindi; this is common in many of the world's languages. The Hindi present tense is made up of a participle and a form of *to be*, something like *is going* in English, but with a straightforward present-tense meaning. Children learning Hindi get evidence from the language around them that the head–complement relationship is left-branching. However, although word order varies across languages, the complements of verbs and pre- or postpositional phrases behave like complements in any language.

Immobile WH-words in Thai

The treatment of question words in many languages is a lot less complicated than it is in English. In English and in many other European languages, question words are moved to the front of their sentences, obeying the restrictions we have described. In other languages the words that would be translated as WH-words in English simply remain in their original positions. This is true of Thai, where questions are asked as in (19) (examples from Haas and Subhanka 1945: 58–59).

(19) a. satăaniiródfay jùu **thîinăj** khràb
station-railroad be-located where (honorific)
'Where is the railroad station?'

b. khun jàag cà? sýy **?àraj**
you want will buy what
'What do you want to buy?'

c. ródfay cà? ɔ̀ɔg **mýaray**
train will leave when
'When will the train leave?'

Thai children hear these question words in the same location in which they are understood. Although UG gives them the restrictions they would need for movement, these restrictions are not needed for syntax of elements like *thîinăj*, *?àraj*, and *mýaray*, since they don't move.

Gender in languages

Many languages have genders associated with all their nouns which affects syntax. English refers to gender only in third-person pronouns, and then gender is "natural" gender, with *he* and *him* referring to people and animals that are sexually male, *she* and *her* referring to females and *it* for everything else (except for the outmoded practice of referring to ships and maybe automobiles as *she*). In a language like German every noun has a masculine, feminine, or neuter gender. Sometimes the genders are just the ones you would expect, so *Frau* 'woman' is feminine and *Mann* 'man' is masculine. Usually, though, the assignment of gender is quite arbitrary. For example, of the common eating utensils, *Messer* 'knife' is neuter, *Gabel* 'fork' is feminine, and *Löffel* 'spoon' is masculine. When these words are put into phrases, determiners and adjectives have to have endings that match the gender of the nouns, so that you get the distinctions in (20).

(20) a. [$_{DP}$ ein [$_{NP}$ klein*es* Messer]]
a small knife

b. [$_{DP}$ ein*e* [$_{NP}$ klein*e* Gabel]]
a small fork

b. [$_{DP}$ ein [$_{NP}$ klein*er* Löffel]]
a small spoon

These gender assignments have to be learned by anyone learning German. Generally you can't guess what gender a noun belongs to. There is nothing particularly feminine about a fork, for instance, or masculine about a spoon.

The task of learning genders in Swahili, a Bantu language widely spoken in East Africa, is much more daunting. Swahili has some ten genders, which are usually called noun classes. Bantu experts count the genders in different ways, so the number of genders is given as anywhere from seven to eighteen. Like the genders in German and other European languages, Swahili genders (noun classes) have meanings, but the assignment of nouns to genders is sometimes arbitrary, like the German words for knife, fork, and spoon. For example, the *m*/*mi* class (so called because of the prefixes on singular and plural nouns, respectively) is sometimes called the "tree" class, because many of the nouns in it refer to trees and plants, but there are words for tools and other things that have this gender, as well. Swahili genders have a pervasive effect in the syntax of sentences. Not only do adjectives have to have matching prefixes (not endings, as in German), but the nouns themselves and the verbs they agree with also have to have the right prefixes. Some of how this works is illustrated in (21) (examples adapted from *Languages of the World 7*: *African Language Families*).

(21) a. M-tu m-moja m-refu a-li-anguka
M-person m-one m-tall m-PAST-fall
'One tall person fell.'

b. Ki-kapu ki-moja ki-kubwa ki-li-anguka
Ki-basket ki-one ki-large ki-PAST-fall
'One large basket fell.'

In (21a), *mtu* is Swahili for 'person' and the root *-tu* belongs, naturally enough, to the "person" (*m*/*wa*) gender. The words for 'one' and 'tall' modifying *mtu* share the "person" gender prefix. So does the verb root *-anguka* 'fall,' even though the prefix is pronounced *a-*. *Kikapu* 'basket' has the "inanimate object" (*ki*/*vi*) gender. As (21b) shows, its modifiers and the verb in its clause all share the appropriate gender prefix. In the case of transitive verbs, the verb agrees with both its subject and its complement.

(22) [$_{IP}$ [$_{NP}$ Wa-toto] [$_{VP}$ wa-na-ki-soma [$_{NP}$ ki-tabu]]]
Wa-child wa-PRES-ki-read ki-book
'The children are reading the book.'

In (22), the root *-toto* 'child,' also in the *m/wa* gender, has the plural prefix for its gender, as does the verb. But the verb also agrees with its object *-tabu* from the *ki/vi* gender. Similar agreement patterns exist for the remaining eight or so genders. In these examples, the meaning of the noun matches the meaning of the gender, but this is not always the case. The *ki/vi* "inanimate object" gender also includes a few words for people and animals and some body parts.

In Thai, nouns are categorized in a gender-like way, but in Thai, there are literally scores of genders. Thai is not a language with prefixes or suffixes, so the genders show up in a very different way. Thai does not have endings for plural, so whether you are talking about one or many things is derived from the context. If you want to specify a specific number of things, it is necessary to use a *classifier* along with the number. Consider the Thai examples in (23).

(23) a. phǒm rúucàg khruu
I know teacher.
'I know a teacher.'
'I know teachers.'

b. phǒm rúucàg khruu khon sɔ̌ɔŋ
I know teacher Clf. two
'I know two teachers.'

In (23a), only the context would determine whether I know a teacher or a number of teachers. In (23b), I specify that I know two teachers, and the classifier used for people has to be used. If I wanted to say that I see two clouds, I would use a different classifier, as in (24).

(24) phǒm hěn mêeg kɔ̂ɔn sɔ̌ɔŋ
I see cloud Clf. two
'I see two clouds.'

The classifier *kɔ̌ɔn* means things that are perceived as lumps, including stones and lumps of sugar, as well as clouds. Thai classifiers (genders) often refer to physical properties. For example, the classifier *baj* represents the gender for fruits and *phy̌yn* is the gender including pieces of cloth in a form than can be used, like towels and curtains. But like all gender systems, there is considerable unpredictability about what nouns are in what gender. The gender represented by *baj* includes not only fruits, but also eggs and containers. Another gender, calling for the classifier *khan*, covers umbrellas, forks and spoons, streetcars and automobiles (Haas and Subhanka 1945).

Languages have a tendency to group nouns in categories that sometimes have a relationship to the meanings of the nouns that are included, but are often entirely arbitrary. These categories are called genders in European languages, and we have used the term for the categories in Swahili and Thai to emphasize what they have in common. Gender has an effect on the syntax of the languages that have gender, but these effects are quite different from one type of language to another.

Functional syntax

This chapter has concentrated on one aspect of syntax, an exploration of structure at an abstract level. This approach has contributed considerable insight into what grammar is like. On the other hand, it doesn't tell the whole story about how language is organized. So far, we have taken what is generally known as a **formal** approach to syntax. Other linguists are developing a variety of approaches that are sometimes referred to together as the **functional** approach. The term *functional* in this context has an entirely different meaning from its use in our earlier discussion of functional categories. Advocates of functional syntax believe that the formal approach leaves out too much that is important. Specifically, it leaves out what the various structures are *for*, how they contribute to the usefulness of language for thought and communication. Functional syntacticians take the position that, if you focus on how syntactic structures are used, you will not need an abstract formal theory. Formal syntacticians, for their part, acknowledge that functional considerations play a role, but say that they are about language *use*, which is guided by a different set of principles from the ones that guide language *form*. For functionalists, separating form and use is artificial and misleading.

Although functional syntax takes this position at a philosophical level, in practice most work in functional syntax is about aspects of language that formal syntax never addresses. By the same token, formal syntax generally takes no interest in the issues that fascinate the functionalists. The result is that there are very few studies of the two kinds that analyze exactly the same phenomenon. You might expect that each kind of study might reveal interesting insights about the structure of sentences that the other does not. This turns out to be true.

A functional analysis of pronouns

Sometimes a functional analysis takes up a syntactic phenomenon and goes beyond what you will find in a formal study. Take, for example, the issue of the use of pronouns. We have seen a formal analysis that is only about what positions within a sentence can and cannot hold the referent of a pronoun. But this analysis says nothing whatever about the principles that guide the choice of a referent from among those that *are* allowed.

To illustrate this kind of analysis, we'll take an example presented by an eminent functionalist syntactician, Talmy Givón (1993: 235), given in (26).

(26) a. . . . after **the queen** said that,
b. **the king** went into a royal sulk.
c. **He** retired into the throne chamber,
d. **Ø** lay on the floor,
e. **Ø** quit eating,
f. and **Ø** refused to talk.
g. Finally **the queen** had had enough,
h. so **she** gave **him** a piece of **her** mind . . .

Consider (26b) and (26c). First, we'll check to see if they conform to what we've already said about pronouns and names, the term we've used for noun phrases that are not pronouns or anaphors.

(27) a. $[_{IP}[_{DP}$ the king$][_{VP}$ went into a royal sulk$]]$
b. $[_{IP}[_{DP}$ He$][_{VP}$ retired into the throne chamber$]]$

The relevant structure for these two cases is given in (27). The structure is essentially the same for both. In order to see if this structure permits either a pronoun or a name under the Binding Theory, we have to check Principles B and C, which we've repeated here as (28).

(28) a. Principle B: pronouns are not coindexed with the subject of their Domain.
b. Principle C: names are not coindexed with the subject of *any* clause located to their left in the sentence in which the name appears.

In (27a), we have a name, so it should not violate (28b). The name expression, *the king*, is the subject of the leftmost and only clause in the sentence. Therefore, there are no subjects of clauses to its left in the same sentence, and no possibility that *the king* could be improperly coindexed with them. So there is no Binding Theory violation in (27a). In (27b), we are dealing with a pronoun in the same structural position, so we have to check it against (28a). As a pronoun, *he* is not permitted to be coindexed with the subject of its own clause. Since *he* actually *is* the subject of its clause, it cannot be coindexed with it, except in a trivial sense not intended by the Binding Theory. As a result, the Binding Theory allows either a name or a pronoun in either (26b) (=27a) or (26c) (=27b).

Although Binding Theory allows either a name or a pronoun in (26b) and (26c), in fact, the name and the pronoun must appear exactly in the order of (26), or the story would have a distinctly different impact, as (29) shows.

(29) . . . after the queen said that, **he** went into a royal sulk. **The king** retired into the throne chamber, lay on the floor, quit eating, and refused to talk.

First of all, notice that in both the version in (26) and the version in (29), there is no guarantee that *the king* and *he* refer to the same person. We might infer that they do, but the Binding Theory, which is only about reference within sentences, doesn't tell us who or what each refers to. Now, notice that if *he* appeared before *the king*, as in (29), we would have no doubt that, if *he* and *the king* refer to the same person, the king has already been introduced into the story before the point at which (29) begins. This is not necessarily the case in the original version in (26). Also, in the version in (29), we infer that using the name *the king* after using a pronoun to refer to him is intended to evoke some special meaning. Perhaps the author wants to emphasize the fact that he is the king and cannot be

So what about "her"?

Notice that in (27h), it appears that the possessive pronoun her *violates Principle B (it's a pronoun that refers to the subject of its clause). But English does not have possessive anaphors (there is no* herself's*), and the Binding Theory allows a pronoun to take the place of an anaphor if an appropriate anaphor is not available in the language.*

stopped from indulging in childish behavior. We don't get any such additional implied meaning from the version in (26).

The Binding Theory – and any theory focusing solely on the sentence – cannot explain these facts, but Givón shows that there is a functional principle that can account for them. He points out that **continuous** or **discontinuous** reference plays a crucial role. A referent is continuous if it is mentally activated, usually (but not always) by being referred to in the preceding clause. Otherwise, the referent is discontinuous. In (26b), the reference to the king is discontinuous, since he hasn't been referred to in the preceding clause. In (26c–h), the reference to the king is continuous, so he is referred to either by an unstressed pronoun or zero anaphora. Notice that the reference to the king remains continuous in (26h), even though there is no explicit reference to him in (26g). This is not a cause for concern in functional analysis; it often looks for general tendencies rather than hard-and-fast laws. On the other hand, the reference to the queen has become discontinuous between (26a) and (26g), so the reference to her in (26g) is by a name, *the queen*. In (26h) the reference to both is continuous, so we have unstressed pronouns referring to both.

Givón's analysis makes a further distinction between **stressed** and **unstressed** pronouns, by which he means whether or not they are pronounced with extra emphasis (stressed pronouns are in all capital letters). Both are pronouns, so both are used for continuous reference, but the difference is whether or not identification of the referent is **problematic** or not. The difference can be seen in (30) (Givón 1993: 236), where the capital letters mean that the word is stressed.

(30) a. Mary told Suzy, then **she** told Sally.
b. Mary told Suzy, then **SHE** told Sally.

In both sentences, both Mary and Suzy are available as continuous references. In (30a), though, the continuous reference is unproblematic. We know that *she* refers to Mary because Mary is the subject of the first clause and continues as the subject of the second. In (30b), we are equally certain that *SHE* refers to Suzy, because the stress on *SHE* signals that the reference, though continuous, is problematic. It is problematic because of the subject switch between the two clauses of the same sentence.

Contrasting formal and functional analyses

Formal and functional syntax most often address different phenomena, making comparisons difficult. It happens that in the example we have just given, a standard formal analysis and the typical functional analysis given by Givón do contrast, at least in emphasis, with respect to the **zero anaphora** examples in (26d–f). Givón, like other functionalists, uses the terms *anaphor* and *anaphora* differently from Binding Theory, to refer more generally to any nominal that "looks back" in a text for its reference. Zero anaphors, represented by the boldface symbols **∅** in (26d–f), are not pronounced, but Givón (1993: 300, n1) considers them a kind of pronoun, functioning like unstressed pronouns to indicate continuous, unproblematic

referents. Givón (1993) does not discuss the difference between the use of unstressed pronouns and zero anaphora, but a functional analysis of the difference is no doubt possible.

A formal syntactician would deny the existence of zero anaphors here. In a formal analysis, (26d–f) would be considered an example of coordination at the verb phrase level, and an example of one source of recursion. The structure would be the one in (31).

(31) $[_{IP}[_{NP}$ **He**$]$ $[_{VP}[_{VP}$ retired into the throne chamber$]$ $[_{VP}$ lay on the floor$]$ $[_{VP}$ quit eating$]$ and $[_{VP}$ refused to talk$]]]$.

If you examine (31) carefully, you will see that *he* is the subject of a large verb phrase that consists of four other verb phrases conjoined with *and*. Each individual conjoined verb phrase is part of a single verb phrase in an inflection phrase with a single subject. In this analysis, there is no need to assume that there are any zero anaphoric elements serving as subjects of the internal verb phrases because, as members of the larger conjoined verb phrase, they share the same subject.

A formal syntactician would agree that it would be possible to have overt pronoun subjects with each of the four included verb phrases. If there were, the result would be (32).

(32) $[_{IP}[_{IP}[_{NP}$ He$]$ $[_{VP}$ retired into the throne chamber$]]$ $[_{IP}[_{NP}$ he$]$ $[_{VP}$ lay on the floor$]]$ $[_{IP}[_{NP}$ he$]$ $[_{VP}$ quit eating$]]$ and $[_{IP}[_{NP}$ he$]$ $[_{VP}$ refused to talk$]]]$.

Here, the formal analysis simply reflects conjunction of inflection phrases, instead of verb phrases, as in (31). Now we have one large inflection phrase consisting of four smaller, conjoined inflection phrases. Since they are inflection phrases, each one needs its own subject. The formal analysis takes the difference between (31) and (32) to be a consequence of the level of structure at which the conjunction takes place.

Formal and functional syntacticians might never agree on what would count as evidence for or against the existence of zero anaphors of this kind. However, a functional analysis that accounted for why a speaker or writer would use zero anaphora rather than unstressed pronouns could directly be applied to the choice between conjoining at the VP level rather than the IP level. Even where functional syntax and formal syntax give substantially different analyses, a complementary analysis using insights from each could be developed.

Chapter summary

Perhaps the most startling thing about the structure of sentences is not about structure at all, but the fact that much of grammatical structure does not have to be learned. People "know" a lot about what is or isn't a possible grammatical structure without having been taught, or even having had the right kind of experience to

have learned it. Instead, there appears to be a language organ that encompasses a person's language ability, with its own intrinsic properties. These properties determine much of what the ultimate structure of someone's grammar will be, independently of their experience. This line of reasoning is called the poverty-of-the-stimulus argument.

One principle of the structure of sentences is **compositionality**; the fact that sentences are composed of clauses and phrases, which in turn are made up of smaller clauses and phrases or words. Compositionality is achieved by projection of simple phrases from words from the mental lexicon. The phrases projected take on the lexical and functional categories of the words that project them. Some phrases have complement and/or specifier branches which merge with phrases that have been projected from other words. A special kind of merger is called adjunction, which allows modifiers (such as adjectives and adverbs) to be included in phrases. Once phrases are constructed by projection and merger, they can be further modified by movement and deletion.

Grammars are finite but are capable of producing an infinite number of sentences. This is achieved by recursion, which allows the same grammatical processes to apply repeatedly, with no principled limitation on how often they may apply. This means that there is no longest sentence in any language, and consequently no limit to the number of sentences it has. Three recursive devices that all languages have are multiple adjunction, embedding, and coordination.

From birth, children seem to "know" structural principles restricting the movement and deletion of phrases. Without being taught, they also "know" the principles of the Binding Theory, which limits the possibilities of coreference between nominal expressions.

This chapter has emphasized what is in common across languages, but there are remarkable differences among languages as well. The word order between heads and complements can vary and not all languages have movable WH-phrases. Grammatical gender is far more varied and important to the syntax of many languages than it is in English.

The approach to the structure of sentences presented here, formal syntax, is based on deducing abstract grammatical principles from observing what sentences are possible and not possible, without regard to how they are used. An alternative approach, functional syntax, emphasizes the influence of language use on its structure. It is possible that a more inclusive approach, complementing structural and functional analysis, would lead to a more thorough understanding of sentence structure than either type of analysis can offer on its own.

Exercises

Composing tree structures

Exercise 3.1

This series of steps will take you through the steps necessary to compose the tree for the sentence: *The snow melted.* Show all your steps.

a. Project a determiner phrase from the determiner *the*. Don't forget the complement.
b. Project a noun phrase from the noun *snow*.
c. Project a verb phrase from the verb *melt*. (Note: *melt* is intransitive.)
d. Merge the NP projected in b into the DP projected in a.
e. Project an Inflection Phrase from the inflection *ed*. Be sure you include all the required structure.
f. Merge the DP resulting from d into the IP projected in e.
g. Merge the VP projected in c into IP projected in e.
h. Make the final adjustment needed to produce the target sentence. (Hint: You should end up with a crossed-out copy of the head of IP.)

Exercise 3.2

Follow these steps to compose the sentence: *Those young people will drink lattes at the café*. Show all your steps.

a. Project a DP from *those*.
b. Project an NP from *people*.
c. Project an adjective phrase from *young*.
d. Merge the AP projected in c to the NP projected in b by adjunction. Be careful to use adjunction structural features. (Note: Show just the result of the merger, not the process.)
e. Merge the NP created in d into the DP projected in a.
f. Project a VP from *drink*. Don't forget the complement. (Note: It will be NP, not DP.)
g. Project an NP from *lattes*.
h. Merge the NP projected in g into the VP projected in f.
i. Project a DP from *the*.
j. Project an NP from *café*.
k. Merge the NP projected in j into the DP projected in i.
l. Project a prepositional phrase from *in*. Don't forget the complement.
m. Merge the DP created in k into the PP projected in l.
n. Merge the PP created in m to the VP created in h by adjunction.
o. Project an IP from *will*.
p. Merge the DP created in d into the specifier position of the IP projected in o.
q. Merge the VP created in n into the complement position of the IP projected in o.

Exercise 3.3

Follow these steps to create the tree structure for: *I believe that the snow melted*. Show all your steps.

a. Project an NP from the pronoun *I*. For the purposes of this exercise, call *I* a noun.
b. Project an VP from *believe*. (Note: Don't forget the complement, here CP.)
c. Project a complementizer phrase from *that*. (Note: This CP will not need a specifier branch, but will, as always, need a complement IP.)

d. Merge the tree you created in Exercise 3.1 into the complementizer position of the CP projected in c.
e. Merge the CP created in d into the VP projected in b.
f. Project an IP from PRES. (Note: Here the specifier is NP, not DP.)
g. Merge the NP projected in a into the IP projected in f.
h. Merge the VP created in e into the IP projected in f.
Note: Since PRES is abstract, we will not bother to move it in this exercise.

Movement and deletion

Exercise 3.4

In this exercise, you will create a tree for which you don't actually have a model in the text. If you have done Exercises 3.1–3.3 successfully, this one should not be too difficult. Show all your steps.

a. Project a CP from an *empty* complementizer *e*. You will need a specifier branch on this CP, as well as the usual complement branch. Leave the specifier branch without a label for now.
b. Project an adverb phrase (AdvP) from *where*.
c. Modify the tree you created in Exercise 3.2. The tree will be exactly the same, except that you will replace the PP *in the café*, with the AdvP projected in b.
d. Merge the IP you got in c into the complement position of the CP projected in a.
e. Apply Aux movement to the inflection *will*. (Note: *e* will be replaced at this step.)
f. Apply WH-movement to the AdvP *where*. (Notes: (1) WH-movement involves merger into the specifier position of CP. (2) Your tree will now contain a new copy of the AdvP, but the original will remain where it was, with its head crossed out.)

Recursion

Exercise 3.5

List each of the recursive devices in each of the following sentences in the order in which they appear. The first one is done for you.

a. I think that you should take a long hard look at this and do it soon.
   Embedding: *that you should take a long hard look at this and do it soon.*
   Multiple adjunction: *long hard*
   Coordination: . . . *and do it soon*
b. The fat yellow cat ate John's juicy oversized hamburger.
c. I know that you think that I don't know much about chemistry.
d. Kim would like to buy the car that your uncle who lives in the townhouse has for sale.
e. I would rather go but you would rather stay.
f. That wise old woman who lives in the village that holds a market every week knows that dishonest people think they can cheat her but she won't let them.

Restrictions on movement

Exercise 3.6

Consider the tree you constructed in Exercise 3.3. An acceptable variant of this sentence is *I believe the snow melted*. Explain how this variant can be permissibly derived from the structure you built in Exercise 3.3.

Exercise 3.7

Consider the following examples:

a. I punished the three young rapscallions severely.
b. I punished severely the three young rapscallions.
c. The three young rapscallions took their punishment bravely.
d. *Took their punishment bravely the three young rapscallions.

We can derive b from a by heavy NP shift of *the three young rapscallions*. Explain why we cannot derive d from c by the same movement.

Binding Theory

Exercise 3.8

Consider the following sentence:

$[_{IP}\ [_{NP}$ Jane's sister] $[_{VP}$ reminded her $[_{CP}$ that $[_{IP}$ the new graduate had promised to improve herself]]]].

a. Of *Jane*, *Jane's sister*, *her,* and *the new graduate*, which one is *herself* required to refer to, according to the Binding Theory?
b. Explain what principle of the Binding Theory predicts this coreference.
c. Of *Jane* and *Jane's sister*, the Binding Theory permits *her* to refer to one, but not the other. Which one is *her* allowed to refer to?
d. Explain how a principle of the Binding Theory predicts this result.
e. According to the Binding Theory, *the new graduate* can refer to one of *Jane* or *Jane's sister*, but not the other. Which one is it?
f. Explain how a principle of the Binding Theory allows one coreference, but not the other.

Word order in "Hindlish"

Exercise 3.9

"Hindlish" is a made-up language consisting of English words and Hindi word order. Translate the following from English to Hindlish. The first one is done for you. (Notes: (1) Hindlish has no determiners meaning *the* or *a/an*. (2) The whole verb phrase, including its complement, is last in Hindlish.)

a. English: The children in the house want a puppy.
   Hindlish: House in children puppy want.
b. English: I read a book in the garden.
   Hindlish:
c. English: The men at work build cars in a factory.
   Hindlish:

Gender in "Enghili"

Exercise 3.10

"Enghili" is a made-up language that has genders that work like gender in Swahili, but has English words. Also, tense in Enghili works the way it does in Swahili, as a morpheme that comes somewhere before the verb root. Translate the following Enghili sentences into English. The first one is done for you. The morpheme for the present tense in Enghili is *pre* and the morpheme for the past tense is *pa*. As in Hindlish, there are no determiners meaning *the* or *a/an*. They will have to be supplied where English needs them. Also in Enghili, noun modifiers follow their nouns. The following is a list of Enghili nouns and genders that you will need:

Noun	Gender prefixes (Singular/plural)
Child	m/wa
Book	ki/vi
Car	ki/vi
Tree	m/mi

a. Enghili: mchild mprekiread kibook.
   English: The child reads a book.
b. Enghili: kicar kipamstrike mtree.
   English:
c. Enghili: wagirl wapamkiss mbaby.
   English:
d. Mitree mitwo mipafall.
   English:

"Thailish" gender

Exercise 3.11

"Thailish" is yet another made-up language with English words. Thailish has genders that work as in Thai, and lacks endings for singular/plural or tense, as Thai does. Like Hindlish and Enghili, there are no words for *the* or *a/an*. The nouns you will need are assigned to their gender classifiers in the table below. Translate the following English sentences into Thailish. The first one is done for you.

Gender classifiers

Person	**stick**	**globe**
Child	pencil	apple
Girl	bat	ball
Student		
Man		
King		

a. English: The children ate two apples.
   Thailish: child eat apple two globe.
b. English: Three girls chased two balls.
   Thailish:
c. English: One bat touched two balls.
   Thailish:
d. English: The players used one bat.
   Thailish:
e. English: Each student has five pencils. (Hint: *Each* works like a number in Thailish.)
   Thailish:
f. English: The men bowed to the king.
   Thailish:

Functional and formal analysis

Exercise 3.12

Here is a brief story containing various kinds of reference. Answer the following questions about the story, using functional or formal analysis, as called for. Items in all capital letters are stressed.

(1) **The little girl** was calling **her cat**,
(2) but **SHE** paid no attention.
(3) **She** darted quickly into the woods
(4) to see what **she** could find.
(5) **She** thought that **the girl** might be angry at **her**,
(6) but like all cats everywhere, **she** wasn't worried.
(7) Suddenly **the little girl** appeared in front of **her**.
(8) **She** scooped **her** up, saying, "Bad kitty, bad kitty."

a. Using the concepts continuous/discontinuous and problematic/ unproblematic reference, give a functional explanation for the use of the full noun phrases *the little girl* and *her cat* rather than pronouns in (1).
b. Using the same concepts, explain the stressed pronoun *SHE* in (2).
c. Give a functional account for the unstressed pronouns in (3) and (4) and the instance of *she* in (5). (Note: The same account applies to all.)
d. Give a functional account for the use of the full noun phrase *the girl* in (5). (The same explanation would apply to *the little girl* in (7).)
e. In (5), (7), and (8), the unstressed pronoun *her* refers to the cat, even though the girl has been reintroduced into the story and a reference to the girl would be continuous. Choose *one* of these instances and give a formal explanation using the Binding Theory for why *her* cannot refer to the little girl.

Suggestions for further reading

Pinker, Steven 1994, *The language instinct*, New York: HarperCollins Publishers. An engagingly written and prize-winning book for nonspecialists explaining the concept of Universal Grammar, which Pinker calls the language instinct.

Baker, Mark C. 2001, *The atoms of language: the mind's hidden rules of grammar*, New York: Basic Books. Another book for nonspecialists about Universal Grammar, but focusing on the differences among languages, with examples. *The atoms of language* is somewhat more challenging than *The language instinct*.

Newmeyer, Frederick J. 1998, *Language form and language function*, Cambridge, MA, and London: MIT Press. Comparison of formal and functional analysis of syntax by a committed formalist with sympathy for functional work. Chapter 1 is an excellent overview of the issues involved.

Radford, Andrew 2004, *Minimalist syntax: exploring the structure of English*, Cambridge University Press. Clearly written, but detailed and technical, presentation of the theory of syntax used in most of this chapter.

References

Clanchy, J. & Ballard, B. (1991) *Essay writing for students* (2nd ed.). Sydney: Longman Cheshire.

Learning Centre, University of NSW. *Questions for the end of the day.*

Learning Development, University of Wollongong. *Essay Writing: The process: The multi-faceted nature of essay writing.*

Learning Development, University of Wollongong. *Essay Writing: The argument: Development and structure.*

Learning Development, University of Wollongong. *Essay Writing: The mechanics of essay writing.*

Learning Development, University of Wollongong. *Editing your assignments.*

Learning Development, University of Wollongong. *Annotated models of disciplinary essays: Annotated philosophy essay.*

Learning Development, University of Wollongong. *Annotated models of disciplinary essays: Annotated modern languages essay.*

Learning Development, University of Wollongong. *Annotated models of disciplinary essays: Annotated sociology essay.*

Learning Development, University of Wollongong. *Annotated models of disciplinary essays: Annotated engineering essay.* Learning Development, University of Wollongong.

MacLean, P. *Essay Writing.* Learning Skills Unit, University of Melbourne.

MacLean, P. *Essay Writing for Education Students.* Learning Skills Unit, University of Melbourne.

MacLean, P. & Starkey, R. *Towards a critical review of the literature.* Learning Skills Unit, University of Melbourne.

Oshima, A. & Hogue, A. (1991) *Writing Academic English* (2nd Ed). California: Addison Wesley.

Woodward-Kron, R., Thomson, E. & Meek, J. (2000) *Academic Writing: a language based guide* (CD-ROM), University of Wollongong.

Best, M., Tucker , J., Oliver, N., Palmer, E., Lyne, S. and Neufeldt, A. 1995 (updated 10 November 2000, cited January 2001) *Essay Writing.* University of Victoria hypertext writer's guide. http://www.clearcf.uvic.ca/writersguide/Pages/EssaysToc.html