

HOW LANGUAGE WORKS



Michael Gasser
Indiana University

Book: How Language Works (Gasser)

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (<https://LibreTexts.org>) and like the thousands of other texts available within this powerful platform, it is freely available for reading, printing, and "consuming."

The LibreTexts mission is to bring together students, faculty, and scholars in a collaborative effort to provide an accessible, and comprehensive platform that empowers our community to develop, curate, adapt, and adopt openly licensed resources and technologies; through these efforts we can reduce the financial burden born from traditional educational resource costs, ensuring education is more accessible for students and communities worldwide.

Most, but not all, pages in the library have licenses that may allow individuals to make changes, save, and print this book. Carefully consult the applicable license(s) before pursuing such effects. Instructors can adopt existing LibreTexts texts or Remix them to quickly build course-specific resources to meet the needs of their students. Unlike traditional textbooks, LibreTexts' web based origins allow powerful integration of advanced features and new technologies to support learning.



LibreTexts is the adaptable, user-friendly non-profit open education resource platform that educators trust for creating, customizing, and sharing accessible, interactive textbooks, adaptive homework, and ancillary materials. We collaborate with individuals and organizations to champion open education initiatives, support institutional publishing programs, drive curriculum development projects, and more.

The LibreTexts libraries are Powered by [NICE CXone Expert](#) and was supported by the Department of Education Open Textbook Pilot Project, the California Education Learning Lab, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation nor the US Department of Education.

Have questions or comments? For information about adoptions or adaptations contact info@LibreTexts.org or visit our main website at <https://LibreTexts.org>.

This text was compiled on 04/14/2026

TABLE OF CONTENTS

Licensing

1: Introduction

- 1.0: Introduction
- 1.1: What We Study
- 1.2: How We Study Language
- 1.3: What We Don't Do- Prescribing and Evaluating Language
- 1.4: Dialects and Languages
- 1.5: A Stand and Two Themes
- 1.6: Why Study Language
- 1.7: Problems

2: Word Meanings

- 2.0: Introduction to Word Meanings
- 2.1: Reference and Proper Nouns
- 2.2: Categories and Common Nouns
- 2.3: Word Senses and Taxonomies
- 2.4: Metaphor and Metonymy
- 2.5: Deixis and Person
- 2.6: Lexical Differences Among Languages
- 2.7: Learning Meaning
- 2.8: Problems

3: Word Forms - Units

- 3.0: Introduction
- 3.1: Phonemes
- 3.2: Iconicity
- 3.3: Vowels (Part 1)
- 3.4: Vowels (Part 2)
- 3.5: English Consonants
- 3.6: Consonants in Other Languages
- 3.7: Syllables
- 3.8: Problems

4: Word Forms - Processes

- 4.0: Introduction to Word Forms
- 4.1: Phonetic Contexts
- 4.2: Assimilation
- 4.3: Distribution of Phones
- 4.4: Learning Phonology
- 4.5: English Accents
- 4.6: Phonological Change
- 4.7: Phonology in the Wild
- 4.8: Problems

5: Composition - Combining Words

- 5.0: Introduction
- 5.1: Attributes and Attribution
- 5.2: Modification
- 5.3: Compositionality and Idiomaticity
- 5.4: Problems

6: Sentences

- 6.0: Introduction
- 6.1: States and Events
- 6.2: Situation Schemas and Semantic Roles
- 6.3: Constituency and Noun Phrases
- 6.4: Subjects
- 6.5: Direct Objects
- 6.6: Adjuncts
- 6.7: Sentence Functions
- 6.8: Problems

7: Grammatical Categories

- 7.0: Introduction
- 7.1: Morphemes
- 7.2: Grammatical Categories and NPs
- 7.3: Grammatical Categories and Verbs
- 7.4: Problems

8: Derivation

- 8.0: Introduction
- 8.1: Derivational Morphology
- 8.2: Foregrounding and Backgrounding
- 8.3: Active and Passive Voice
- 8.4: More Verb Derivation
- 8.5: Problems

9: Appendices

- 9.0: Phonetic Symbols
- 9.1: Languages Cited in this Book
- 9.2: References

[Index](#)

[Glossary](#)

[Index](#)

[Glossary](#)

[Detailed Licensing](#)

Licensing

A detailed breakdown of this resource's licensing can be found in [Back Matter/Detailed Licensing](#).

CHAPTER OVERVIEW

1: Introduction

- 1.0: Introduction
- 1.1: What We Study
- 1.2: How We Study Language
- 1.3: What We Don't Do- Prescribing and Evaluating Language
- 1.4: Dialects and Languages
- 1.5: A Stand and Two Themes
- 1.6: Why Study Language
- 1.7: Problems

This page titled [1: Introduction](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.0: Introduction

Students studying linguistics and other language sciences for the first time often have misconceptions about what they are about and what they can offer them. They may think that linguists are authorities on what is correct and what is incorrect in a given language. But linguistics is the *science* of language; it treats language and the ways people use it as phenomena to be studied much as a geologist treats the earth. Linguists want to figure out *how language works*. They are no more in the business of making value judgments about people's language than geologists are in the business of making value judgments about the behavior of the earth.

But language is a cultural phenomenon and we all have deep-seated, cultural ideas about what it is and how we ought to use it, so knowing where to begin in studying it scientifically is not a trivial matter at all. Issues arise that would not if we were geologists figuring out how to study earthquakes or the structure of the earth's crust. For this reason, before we dive into the study of language, we will need to examine some of the biases that we all have concerning language and to set some ground rules for how we are going to proceed. Because there is more than one way to begin, it will also be useful to establish a basic stance to guide us. Finally, because human language is an enormously complex subject, the book will focus on a narrow range of topics and themes; there will be no pretense of covering the field in anything like a complete fashion. This first chapter is designed to deal with these preliminary issues. But first, you will need to know about the various conventions that I will be using in the book.

This page titled [1.0: Introduction](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.1: What We Study

A science takes phenomena of one kind or another as its subject matter and attempts to describe and explain them objectively. Scientists gather particular kinds of data, analyze them, and create theories that account for the data. Here are some sciences and informal descriptions of the phenomena they are concerned with.

Science	Subject matter
chemistry	how substances combine to form other substances
psychology	how individuals behave
sociology	how people behave in groups
cultural anthropology	how human cultures resemble and differ from each other
linguistics	how language works

How would linguists and other language scientists go about objectively describing and explaining **how language works**? What kind of data would they examine? How would they analyze the data? What would it mean to "account for" the data with a theory? We will look at these questions in this section.

Social Science and Objectivity

Exercise 1.1.1

You're a female American anthropologist specializing in the study of family structure and the roles that different family members play. You're studying a remote ethnic group, and you discover that the people believe that baby girls need to be brought up more harshly than boys to prepare them for a harsh life and that fathers have no particular role to play in the bringing up of their children. How might it be difficult for you to be objective in your study of these beliefs and behaviors?

Of the sciences listed in the table above, all but chemistry are concerned with *human* behavior. For all of these sciences, the work of the scientist is complicated by the attitudes that we have toward the behavior; that is, it may be difficult to be objective in our study. It is probably cultural anthropologists who face this difficulty most often. The anthropologist described in the box above, having grown up among highly educated people in modern US society, almost certainly believes that boys and girls should be treated equally and that fathers should play an important role in the raising of their children. But she must somehow put these beliefs aside in her work. Instead of labeling the people's practices as "wrong", she must attempt to see how they fit in with the other practices in the society, whether she notices consistent patterns in their behavior. If she is to make any value judgments about the behaviors, they should be based on whether the behaviors contribute to the stability of the society, not on her own feelings about them. Her job is to describe and explain the society, not to judge it.

As another example, consider the case of another anthropologist specializing in the scientific study of religious beliefs and practices. To do this objectively, he has to put aside whatever beliefs he himself has on questions such as the existence of God and life after death. There is no way he can pretend not to have such beliefs; he just has to try to keep them from getting in the way. Note that this is not a problem for chemists; in their work they don't have to worry about their prior attitudes toward particular chemical reactions.

When we attempt to study human phenomena such as religion or family roles or language scientifically, we do not deny the inevitability or even the value of our attitudes toward the phenomena. It's just that *evaluating* the phenomena and trying to convince people to behave or believe in a particular way is not the business of scientists (at least not the main business); it's somebody else's. A priest or minister has a very different purpose from an anthropologist specializing in the study of religion (although the priest or minister can probably benefit from the insights of anthropologists who study religion). A family counselor who advises a husband

to allow his wife more freedom is not doing science (although the counselor can probably benefit from the insights of anthropologists who study the family).

It is important to come to grips with some of our preconceptions about language before we begin to approach language as the object of scientific study. The next section is about some of those preconceptions and where they come from.

Attitudes Toward the Speech of Others

Exercise 1.1.2

You overhear the following conversation.

1. A: *Did you hear those two girls talking? "He don't mean nothin'." "I seen it." "Me and him fought." Can't they learn to speak English?*

B: *I know what you mean. They're just lazy, if you ask me.*

What do you think of comments like this? Do speakers from some regions or speakers belonging to some social or ethnic groups tend to be lazier than others in their speech?

As you certainly know, people are quite conscious of how they differ from people from other regions, social groups, or ethnic groups. They notice differences in dress, in food, in patterns of social interaction, in which qualities are valued or attract attention. And it is natural to evaluate these features of other groups, to think of their dress as fashionable or weird, to think of their food as tasteless or gross, to think of their social behaviors as friendly or offensive. The same is true for language. People hear speech that differs from their own and they may find it sloppy, elegant, or monotonous. These impressions may also be associated with the languages of particular groups rather than (or in addition to) the people themselves: we may find a certain language more expressive, more logical, even more masculine. What's the source of these impressions? Are they accurate?

Differences Between Languages

Undeniably communities of people do tend to differ. To take an obvious example, food preparation is more important in some cultures than others; some cultures are famous the world over for their cuisine. For language, the differences are again obvious to anyone. It's not just that languages sound different. Some languages make distinctions — in sounds, in words, in grammar — that others don't; in fact most of this book is about just this topic. And people learning a second language often have trouble making the distinctions that aren't part of their first language. What we naturally notice, as speakers of a particular language, is what is "missing" in other languages and what kinds of mistakes second-language learners make in trying to speak our language. This may lead us, consciously or unconsciously, to think there is something deficient about the other language or even about the speakers of the other language. It is very difficult for us to see it from the other perspective, to see that we also fail to make distinctions that matter in the other language and have trouble making them when we try to learn that language.

For example, as speakers of English, we may be surprised to find that Japanese has no words corresponding to English *a* and *the*, words that are so basic to English we may almost take them for granted. And we may be struck by the errors that Japanese learners of English make in trying to master these words. Similarly, we are struck by the confusions Japanese learners may have in pronouncing English words with the sounds that we write with *l* and *r*, a distinction not made in Japanese. But these same Japanese speakers may be surprised when they first learn that English has only one word for 'you' (Japanese has at least six possibilities) and struck by the tendency of English-speaking learners of Japanese to always use the same word for 'you'. And they are similarly struck by the difficulty English-speaking learners of Japanese have with distinctions in vowel length and pitch change, distinctions that don't exist in English.

In fact there is no evidence that people in some cultures speak in sloppier or more elegant or more monotonous ways than people in other cultures. And while languages do differ in striking ways, these different features seem to balance each other out. As far as we know, all languages are equally expressive, equally logical. If you're not already convinced of this, I hope you will be after you have read this book.

But the example in the box above doesn't concern two different languages; it concerns a single language, English, and its speakers. The fact is that there is also considerable variation within English (or any other major language); that is, English has dialects. I'll have a lot more to say about dialects in the section on dialects and languages. For now, the main point to be made is that what

linguists have learned about the essential equality of languages applies to dialects as well. Though it is often even harder for people to accept this fact for dialects than for languages, as far as anyone knows, there is nothing inherently inferior or superior about any dialect of any language.

Linguistic Chauvinism and Intolerance

So if impressions like those of the speakers in the box above have no basis in fact, where do they come from? There are three possibilities. First, these people may have been told by an authority, for example, an English teacher, that certain usages are just plain wrong. Clearly, the reasoning would go, anyone who knows this should not be using those forms. We'll return to this issue in the next section. The fact is that what is "wrong" is all relative. The girls quoted in the conversation in the box would almost certainly find it wrong to say *it doesn't mean anything* when speaking to each other. If they wrote *it don't mean nothin'* in a school essay, on the other hand, that would be another matter (though it would still not be reason to call them lazy; it would just be evidence that they had not learned the rules of the variety of English that is appropriate in school). Second, these people may have a stereotype concerning the group in question, and they may be transferring that stereotype to the speech of that group. Third, what they hear differs from the English they speak, and people may be quite intolerant when it comes to speech. Especially if they belong (or believe they belong) to a political, economic, or intellectual elite, their view may be something like the following: "the way I speak the language is the right way; any other way is wrong".

Whatever the reason for the impressions of A and B in the box, a linguist would respond to them by saying that the two girls were simply speaking a different dialect of English, a dialect with its own grammar differing from the grammar of the dialect of A and B.

Describing and Explaining Language

You're a linguist who wants to study the English spoken in the Caribbean island nation of Grenada. How do you go about this? What sort of data do you gather?

Data for Research on Language

Linguists and other language scientists, unlike A and B in the box in the last section, are interested in what people do, not what somebody thinks they should do. The English of the girls overheard by A and B is just as legitimate an object of study as the speech of any other group. To carry on their study, clearly researchers need to gather examples of language. There are two sorts of ways to get these.

- By collecting naturally occurring language, either written texts or spoken language.

Linguists usually study spoken (or signed) language because it is more basic than written language. Most of the human languages that have existed have not been written at all, and among those that are written, many people do not read or write them. In addition, though language learning continues throughout life, most of the basic patterns of a language are probably mastered by the time a child is six years old. So the written form of the language has little or nothing to do with this fundamental early learning of language.

- By eliciting language by asking people particular questions or by doing experiments that call for language.

Linguists use both kinds of data. For example, once you'd arrived in Grenada, you might get permission to record phone conversations, then transcribe the conversations, perhaps using a special notation that shows the speakers' pronunciation. Or you might recruit one or more willing speakers to help you in your study by translating words or sentences from your English into theirs or by telling you whether certain sentences are possible in their English.

Experiments on Language

Rather than using words and sentences produced by speakers (or writers), linguists and (even more often) other language scientists sometimes gather other kinds of data. For example, they might record the acoustic properties of speech or the movements of the tongue, lips, and jaw during speech. Or they might present people with words, sentences, pictures, or movies (the stimulus) and see how they respond or how long it takes them to respond to them. The responses in experiments like these could involve

- saying something, for example, the first word that comes to mind or a sentence that describes a visual stimulus.
- pressing a button, for example, to make a choice concerning the stimulus.

Given some data, a linguist or other language scientist has to do something with it. Much of this book will be concerned with what they might do, so what follows will be just an introduction to this topic. For every kind of research on language, there are two things to be considered: what aspect of language is being studied and what the research is supposed to accomplish (and how it does this). We'll start with the first and look at the second in the next section.

The Content of Research on Language

Language, even a particular individual language, is far too complex a subject to be studied in its entirety by any one researcher. As already mentioned in the overview of the book, there is a higher-level distinction we can make concerning what research is supposed to be about, between the study of *language as system* and the study of *language behavior*. In either of these two cases, the language scientist is normally studying only some aspect of the phenomenon, one or more of four sorts of things about a language (or dialect): its sounds, its words, its grammar, and its use in context. A researcher interested in the sounds of the language (or dialect) might try to figure out what the basic sounds of the language are, how they combine to form words, or how speakers produce the sounds. Many researchers in this area believe that it is possible to study the sounds of a language more or less independently from the other aspects of the language.

A second area of research is the words of a language, usually thought of as organized in some sort of abstract dictionary, referred to as the **lexicon**. A researcher interested in words might study how speakers find words when they are formulating sentences or how abstract meanings build on simpler meanings (how is *over* in *get over a problem* related to *over* in *jump over the puddle*?).

Another possible kind of research would try to characterize what counts as a possible sentence in the language, that is, what's **grammatical** in the language. It is not as easy as it might seem to define this concept. We must be careful to avoid any bias on the speaker's part based on what they have heard from teachers of their language in school because what we care about is what people actually say, not what someone tells them they should say.

But we also cannot just treat any sentence that occurs as grammatical because people make **speech errors**. By "errors", we do not mean that they break rules that apply to dialects other than theirs (for example, by saying *ain't* or *he don't*). Instead we mean slips of the tongue, false starts, and hesitations. For example, the following example includes several speech errors.

2. *Well, I think you ... I mean, the ... the ... um, this isn't coming out right at all.*

People produce such "sentences" all the time, but they clearly also know that there is something wrong with them. That is, linguists probably do not want their descriptions of a language to include such sequences. So grammatical sentences are possible sentences that do not contain speech errors.

This is not to say that there is nothing interesting about speech errors. In fact, like human errors more generally, they can give us lots of insights about the underlying mechanisms. There is a whole community of researchers that take speech errors of one kind or another as the data they try to explain.

Degrees of Grammaticality

But there is another complexity; grammaticality doesn't seem to be an all-or-none matter. That is, while some forms may be completely acceptable to all of the speakers of a given dialect all of the time and other forms may be completely unacceptable to all of the speakers all of the time, there may also be intermediate cases that are not so clear. For example, some English speakers use *a chalk* for a stick (or piece) of chalk; others would be less comfortable with this (though they would not find it as unacceptable as, say, *a clay* for a lump of clay). There may even be variation within a single speaker. For example, an English speaker may say *for my wife and me* on some occasions and *for my wife and I* on others.

Returning to the examples ridiculed by A and B in the box above, we see that by the definition of grammaticality that linguists work with, such sentences as this one may be perfectly grammatical for the girl who said it.

3. *He don't mean nothin'.*

A complete description of the grammar of these girls (and the community of speakers that they belong to) would have to specify just what counts as a grammatical sentence (for example, sentence 3) and what doesn't (for example, *he don't nothin' mean* or *he don't meant nothin'*), possibly singling out areas of grammar where there is disagreement and variation among the speakers. This description would obviously have to say something about word order and about which forms can go with which other forms (*meant* is a perfectly good word in their dialect, but not following *don't*).

But some linguists are not satisfied with just describing the grammatical sentences because this says nothing about what those sentences are for. Instead these linguists are concerned with describing how meanings and functions of language relate to words and grammatical sentences (for more on this idea, see this section). So an account that includes sentence 3 above as a grammatical sentence (for some English speakers) doesn't help us understand how this sentence conveys information about some person familiar to the hearer (*he*) and about the speaker's belief about that person's intentions. This book follows this second position on what we should be describing, that is, that we should be saying how language accomplishes things for speakers and hearers.

Finally, those same linguists who are interested in how sentences convey meaning may also be interested in describing a fourth sort of aspect of the language, a sort of "correctness" that is different from grammaticality. A sentence can be grammatical and meaningful — that is, the words and grammatical patterns in the sentence can sound right and correctly describe some possible situation in the world — but the sentence can still be inappropriate. Consider the following sentence.

4. *There is no life on the moon.*

This sentence makes perfect sense and describes a true state of affairs. But if you walked up to a stranger on the street and said it, they'd think you were crazy. It would not be an appropriate way to begin (or end, for that matter) a conversation with a stranger. Just as speakers of a language have knowledge about what is grammatical in their language, they also have knowledge about what is **appropriate**.

Learners' Errors

Of course not everyone who uses a language or dialect (see this section for the difference between dialects and languages) knows how to do so grammatically and appropriately in all situations. In particular, language learners have only imperfect knowledge of the language or dialect they are learning, and they can be expected to make errors. Children learning English as a first language may say "doos" for *juice* or *me up* when they want to be picked up. Teenaged speakers of English as a first language may still commit errors of appropriateness, using informal expressions such as *bigtime* in formal contexts. And adults learning English as a second language may say "diss" or "dees" for *this* or *I make the homework* for *I am doing the homework*. We consider these to be errors, but there are only errors from the perspective of the system defined by the behavior of the adult native speakers of the language or dialect that is being learned. Such examples can also be seen relative to the learner's own linguistic system, which has its own pronunciation, vocabulary, grammar, and patterns of usage. Researchers studying second-language learning often find it useful to treat the learner's knowledge of the second language as a sort of language in its own right, what they call "interlanguage".

This section has looked at the kinds of topics that interest linguists and other language scientists and the kinds of data that they might look for to help them in their research. But we haven't thought much about what the outcome of the research is. What would it mean to describe or explain language? We'll look at these aspects of research on language in the next section.

This page titled [1.1: What We Study](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.2: How We Study Language

Accounts, Generalizations, and Theories

In addition to content, we can look at research on language from the perspective of what it is trying to accomplish and how it does that. First, we need to go back and consider again what science is all about. As we've seen, a scientist starts with a phenomenon of interest, gathers some data on the phenomenon, and attempts to come up with a description or explanation of it. This may take the form of a discussion in some language (such as English or Chinese), a set of equations, or an algorithm, that is, a precise description of a set of processes to be carried (by a person or a computer). Whatever form it takes, this result of the scientist's research may be referred to as an "account" or an "analysis" of the phenomenon. A scientific account is expected to include some sort of **generalization** about the phenomenon, that is, to go beyond simply listing the data.

A generalization can be relatively specific, applying only to a very narrow range of data. For example, an anthropologist studying kinship relations in some ethnic group might conclude that the relation between a mother-in-law and daughter-in-law matters more than other in-law relations in the culture. However, most often, an anthropologist wants a particular culture to be seen only as an example of a more general phenomenon. That is, the goal is not just to describe or explain one culture but to say something general about human culture. So the anthropologist might want to state a generalization about in-law relations across all ethnic groups belonging to some type or to all ethnic groups. In the most general cases, the generalization is usually stated in terms of a particular **theory**, which is a general set of principles for understanding phenomena of a particular type. A theory is supposed to offer an *explanation* of the phenomena, not just a description. For example, kinship theories start with a set of basic categories that are supposed to be sufficient both for describing and explaining the role of the different possible kinship relations in all societies. A theory is like an "account" or an "analysis", only more general.

How Scientists Describe and Explain: Generalizations and Theories

To take a linguistic example, a researcher might be interested in describing the way present-tense verbs are negated in English dialects like the one referred to by A and B in the box in the last section:

- *I/you/he/she/it/we/they don't go*

The linguist could make the low-level generalization that for all but a few verbs in these dialects, the present tense negative is formed by putting *don't* before the base form of the verb (*go, mean, eat*, etc.). This generalization would apply only to the dialects being investigated.

Or the researcher could note that there is a distinction in the affirmative in these dialects that is not made in the negative, the same distinction that is made in standard English dialects:

- *I/you/we/they go*
- *he/she/it goes*

Explaining Languages, Explaining Language

Going beyond these English dialects, the linguist might then discover that something similar happens in many languages, for example, among the languages that this book focuses on, in Japanese and Amharic, making the generalization that languages tend to make more distinctions in the affirmative than the negative. The researcher could then go even further and try to place negation within a set of other forms in terms of how likely they are to favor distinctions; for example, the researcher might discover that fewer distinctions are made in verbs in subordinate clauses (for example, in *the places they go*) than in independent clauses (for example, in *they go many places*) in many languages. (Don't worry if you don't understand what "subordinate", "independent", or "clause" means.) Even more abstractly, the linguist could try to explain why there would be such tendencies across many languages. For example, they might propose that it is more difficult to produce and understand negative forms than affirmative forms and that this puts pressure on the users of the language, and hence on the language itself, to compensate by making some other aspect of the grammar simpler. This kind of proposal would be a theory that is designed to describe and explain a set of grammatical phenomena across many languages.

To summarize this section so far, we see that a scientist trying to understand language makes generalizations about data. These generalizations can apply only to one language, or they can apply to language in general. A linguist or other language scientist often

works within the context of a particular theory of language. I'll have more to say about the role of theories in this section. But we've said nothing so far about what the goal of a description or explanation is. There are two kinds of possibilities, related to the two fundamental ways of looking at language that are mentioned in the overview of the book: focus on the **product** (sounds, words, sentences, etc.) and focus on the **process** (speaking, reading, understanding, etc.).

Product and Process

Exercise 1.2.1

Consider this line from the English comedy show "The Two Ronnies":

1. *She left her husband for the garbageman.*

The joke, in case you missed it, is based on the **ambiguity**, that is, the multiple possible interpretations, of the verb *left*. The sentence could mean that the woman abandoned her husband in favor of the garbageman, or it could mean that she put him outside for the garbageman to carry away. It turns out that ambiguity is quite common in language and is the basis of puns such as this one. What sort of a problem does this present for a description or explanation of how language works? What aspects of language would an explanation of ambiguity need to refer to?

All instances of language are obviously the result of processes: speakers, writers, and signers (in the case of sign language) produce something that we call language and hearers, readers, and sign observers attempt to understand something that has been produced. Over a longer time scale, what a person knows about a language (that is, how to produce and understand it) changes; this is the process that we call **learning**, **development**, or **acquisition**. Over an even longer time scale, every language also changes; the English of today is not the same as the English of 1900. This slower process is called **language change**. Finally, over the longest time scale that is relevant for language, we know that at some point in the distant past, for example, 200,000 years ago, the ancestors of modern humans did not have anything like what we call language; the process that resulted in the kind of system we have now is called **language evolution**.

Most linguists choose to ignore all of these processes and focus instead on the products, the words, sentences, and entire discourses that are produced and understood by people at a particular point in historical time. In their research they attempt to describe and explain these products, to generalize about what must appear, what may appear, and what may not appear. Other language scientists, especially psycholinguists and computational linguists, focus instead on the processes themselves. In their research they attempt to describe and explain these processes, to generalize about what is going on during language behavior, language change, or language evolution, for example, when a speaker pronounces a word, when a child learns how to combine words into sentences, when a population of agents "invents" grammar in the process of evolving a communication system. Their accounts and theories often take the form of algorithms and are often implemented in the form of computer programs. They may be called **processing accounts** or **computational models**.

There are at least two important differences between these two ways of looking at language. First, the product-oriented perspective deals with *static* objects; even though it took time for the words or sentences to be produced or understood, the things being studied have no real time in them, except in that certain parts come before other parts. In the process-oriented perspective, on the other hand, time cannot be ignored. Processing happens in real time, and the implementations of processing accounts as computer programs obviously run in real time. These accounts differ a lot in terms of how serious they are about the time course of human language processing, but they are all in some sense *dynamic*.

Production and Comprehension

Second, processing accounts are *directional*. At its most basic, language processing is either production or comprehension. Even processing accounts that are concerned with the slower processes of learning or evolution are based on some idea of how production or comprehension takes place. I'll have more to say about production and comprehension later in this chapter. For now the important idea is that these processes occur in opposite directions.

In production, a speaker, writer, or signer starts with something to be communicated (perhaps to themselves) and then goes through a process that results finally in an instance of language (a spoken, written, or signed utterance of some kind). When I produce the sentence *take the garbage out*, I start with something that isn't language at all, something that doesn't include the words *take* and

garbage but is more like my mental representation of some situation in the world, either one I'm experiencing (seeing the garbage piling up in the house) or one I'm imagining (seeing the hearer taking the garbage out). Then somehow I get from this thought to the utterance itself. In comprehension, a hearer, reader, or sign observer starts with an instance of language and then goes through a process that results in some sort of approximation to what the speaker, writer, or signer wanted to communicate. When I hear somebody say the sentence *take the garbage out*, I start with some sounds and from these eventually figure out what the speaker wants and what I'm supposed to do, and of course what I'm supposed to do isn't language at all; it's an action. The product-oriented perspective ignores this directionality, treating it as irrelevant to what makes language work.

So who is right? Does the product-oriented perspective or the process-oriented perspective give us more insight into how language works? To some extent, the answer depends on what we're after. If we want a way to describe languages in as efficient and understandable a way as possible, then it may be that we can confine ourselves to product-oriented research. The outcome of this research could be an archive of many languages in a form that would allow researchers from different fields and people who want to teach or learn the languages to consult the archive. Of course if our goal is to understand what people are doing when they are actually using language or if we want to write programs that allow computers to use language, then we will need to rely on process-oriented research.

But what if we are interested in the more abstract and theoretical question of *why* language is the way it is? Which perspective is the right one, or do we need both? There is a lot of controversy within the language sciences on this point, with one camp, associated especially with the famous linguist Noam Chomsky, claiming that we can learn what makes language special by studying products alone. The idea is that processing is something separate, something to be understood in its terms, but not something we need to refer to to understand how language works. The opposing camp, associated with the theories in linguistics known as cognitive linguistics and functionalist linguistics and with many language researchers in fields outside of linguistics proper, takes the view that the nature of language is intimately tied to the way it is used, that if we want to understand *why* language is the way it is, we need to refer to processing. As I discuss more in this section, this book belongs more in the second camp. I will assume that both product-oriented and process-oriented perspectives can help us understand how language works.

Let's return to the issue of ambiguity, illustrated in the box above, to see how thinking about process as well as product can help us figure out what is going on. For the sentence in question, a product-oriented approach might simply include the information that *left* (or *leave*) has at least two meanings, though of course it would have to be more precise about what is meant by "have two meanings". A process-oriented approach would look at the processing of the sentence from both directions. This would make clear that ambiguity is a "problem", that is, a potential challenge for a person or a computer, in the comprehension direction but not necessarily in the production direction. It is such a problem in fact that a very large body of research from the processing perspective has looked at what is called **disambiguation**. The problem is that when a word is ambiguous, a listener (reader, sign observer) has to figure out which meaning is intended (or, in the case of a joke like the one in the box, to see that the sentence has a possible interpretation for each of the meanings). It is easy to imagine how a person hearing a sentence like the one in the box would know that the word *left* has more than one meaning. What is hard to explain is how the person knows which meaning is the right one, or, in this case knows that both are right. Theories of disambiguation are designed to explain the process. We'll return to disambiguation later; the important point for now is that it only becomes an issue when we look at language from a processing perspective.

In this section and the last, we've seen that the business of linguists and other language scientists (and the subject of this book) is trying to describe and possibly explain particular languages, language in general, and how people use languages. But this is not the only way we might treat language. Some people who are not scientists and may not be particularly interested in the scientific study of language do talk about language as a part of their work. And one of their concerns may be deciding what people should and should not say and trying to enforce these decisions. This kind of work, some of its consequences, and how linguists sometimes comment on it are discussed in the next section.

This page titled [1.2: How We Study Language](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.3: What We Don't Do- Prescribing and Evaluating Language

Standards

Exercise 1.3.1

You're in a creative writing class discussing short story writing, and you say

1. *You shouldn't introduce a new character unless they have an important part to play in the story.*

The instructor of the course corrects you: "'He or she has', not 'they have'." How do you feel about being corrected in this way? Do you think the instructor was justified?

Just as it is normal to evaluate the religious beliefs and practices and the family behavior of particular individuals within our culture, it is also normal to evaluate the linguistic behavior of people. We may treat some speech and writing patterns as acceptable or unacceptable, superior or inferior, appropriate or inappropriate. Attitudes such as these must be based on a **standard**, some idea of what counts as desirable behavior. People who are concerned with defining and maintaining linguistic standards are **prescribing** (rather than describing) language. Usually they prescribe aspects of grammar, and in this role they are referred to as **prescriptive grammarians**. There are three kinds of questions we can ask of linguistic standards and language prescription.

1. What is the purpose of the standard? What use is it put to?
2. Where does the standard come from?
3. How is the standard enforced?

Why a Standard

One reason for a linguistic standard is that people within a community (for example, a nation) will be better able to understand one another if they agree on a set of words and a set of rules for pronunciation, spelling, and grammar. The process of defining the rules (and sometimes the set of words as well) is called **standardization**. Once a standard has been agreed on, it can be used in the media and taught in the schools.

Japan provides a good example of this process of standardization and promulgation of the standard. Children all over Japan, speaking widely divergent dialects of Japanese, are all taught to speak and write Standard Japanese in school, and radio and television announcers are all expected to be familiar with the standard vocabulary, pronunciation, and grammar. If these announcers were to speak in the southern dialect of Kagoshima or the northern dialect of Aomori, people in most of Japan would have difficulty understanding them. Something like this happens in the United States, though more informally. Children across the country are taught the same written standard vocabulary and grammar, though they may not be taught a single standard pronunciation, and radio and television announcers tend to speak in a single standard.

A related, though often unstated, goal of a standard may be to eliminate diversity, which may be perceived as a threat to national unity. Sometimes the diversity is reflected in different, though closely related, dialects, but sometimes it is reflected in completely different languages. It is one thing to teach a standard in the schools across a country. It is another to discourage people from speaking their local dialects or languages (for more on dialects and languages, see the next section). Sometimes the native speakers of the non-standard dialects or languages become willing participants in this process in their desire to be integrated into the society.

In Japan, for example, standardization sometimes had the effect of eliminating diversity. In the case of the Ryukyu Islands, the southernmost region of Japan, the dialects that are spoken there are so different from the standard that they may be considered a separate language (or languages). Today, following the repressive policies of school administrators in the first half of the twentieth century as well as the economic and social pressures for Ryukyans to conform, this language is on the verge of extinction. This is a familiar phenomenon in many countries with minority populations. In the US and Canada, immigrant and American Indian students have often been prevented from speaking their home languages in schools. These educational policies were most notorious in the case of American Indian students, who were sometimes sent to schools which kept them separated from their communities and were punished for speaking in their native languages. These policies were one of the major factors leading to the extinction and near extinction of the majority of North American Indian languages. For more on the topic of language death and "endangered

languages", that is, languages in danger of dying, see the websites of Terralingua and the Foundation for Endangered Languages, two organizations dedicated to linguistic diversity.

Finally a (never stated and perhaps often unconscious) purpose of a standard may be to exclude certain groups from power. If the standard is based on the speech of one group, either from a particular region or a particular class, then this gives people in that group an advantage when it comes to jobs and ultimately power.

Standard Languages and Politics

Clearly language standardization is a political issue, and as such it is not really the business of linguists (though it is studied by sociologists interested in the social and political aspects of language). However, as linguists often become closely involved with the people whose languages they study, they may become advocates for these groups when their languages or their well-being are threatened because of the language policies of governments.

Other reasons are often stated for prescribing language, though these may mask the ultimate political reasons. These reasons include the supposed illogical or ambiguous nature of some constructions used by people, and linguists have also sometimes gotten involved in the debate in these cases because of their expertise. An example is the use of English *they*, *them*, and *their* to refer to a single person, usually of unknown or unspecified gender, as in the example in the box above. The idea is that since *they* is supposed to be plural, it should not be allowed to refer to one person. In cases like this, prescriptive grammarians are trying to actually improve the language or perhaps to preserve what they suppose to be an earlier, purer form of the language. The problem is that linguists and other language scientists, experts on the "logic" and the degree of ambiguity in language, can find nothing inherently wrong with any of these constructions, at least not with the usual examples of "bad" English constructions. In fact languages seem to have their own built-in "prescriptive" mechanism which weeds out whatever patterns don't work. It's a survival-of-the-fittest sort of arrangement, and it implies that those common patterns such as "singular *they*" (which has been around in the spoken and written language for at least 500 years) are quite fit linguistically.

Where Standards Come From

If we are to define a set of standards for a community, they have to come from somewhere. One possibility for the source is from other languages, and, strange though this may seem, standardizers have sometimes resorted to it. Take the "split infinitive", which prescriptive English grammarians sometimes argue against. An English infinitive is an expression consisting of *to* followed by a verb stem (don't worry now if you don't know what this is), for example, *to go* or *to sing*. So according to this rule, it is incorrect to say or write *to boldly go* or *to not sing*. The main argument seems to have been based on the fact that Latin and German and other familiar European languages cannot split their infinitives. But at least nowadays it seems just silly to try to force one language into the patterns of another.

Only slightly less extreme is the attempt to impose patterns from earlier in the history of the language onto present-day speakers. An example of this is the distinction between the first consonant in *which* and the first consonant in *witch*. For most speakers of American English and the English of England, there is no distinction, but there used to be (and there still is in some dialects, for example in Irish English and Southern (US) English). There are still school teachers in the United States who tell their students that it is "wrong" to pronounce *which* like *witch* (my daughter had such a teacher). As we will see at various points throughout this book, language is always changing, and there is little we can do to stop this. In any case, if the loss of the distinction between these two consonants interfered with comprehension, it wouldn't have happened.

It ain't about Good and Bad Language; it's about Class and Education

If the standards come instead from among the speakers of the modern language, the question is which speakers. Language usage varies both with region and with social group, so if we want to standardize, we need to choose a region or group. As noted above, there are sometimes political considerations involved. In many countries, standard usage comes from the speech and writing of well-educated or upper-class people from a particular region, though the regional origin of the standard is not so clear in the case of the United States. For example, the current English standard prefers *you aren't* to *you ain't* because the latter form, although common in regional dialects all over the English-speaking world, is associated with uneducated speakers.

Actually this argument only holds today. The history is more complicated. *Ain't* was once used by upper-class speakers, but prescriptive grammarians who felt that it was somehow "lazy" or "illogical" succeeded in mostly obliterating it from upper-class

speech in the English-speaking world. See this short history of *ain't* from the *American Heritage Dictionary of the English Language* (2000).

Another example is the pronunciation of pairs of words like *pin/pen* and *since/sense* as the same. This is a feature of the speech of a large region of the United States, including the South and some areas bordering on it, but it would not be considered a feature of standard American English.

How Standards are Enforced

Given reasons for creating a standard and agreement on what constitutes the standard, there is the question of how the standard patterns are to be spread through the population. The obvious venue for this is the schools, and in many countries considerable attention is devoted to making sure students are familiar with the standard language of their community. Much of this involves simply exposing students to examples of (mostly written) texts in the standard, and this often has the desired effect, at least in the writing of the students. Sometimes the teaching of the standard involves attempts to prevent or undo frequent non-standard patterns of usage. Here we are sometimes dealing with standards that don't correspond well to the usage of *any* native speakers of the language, including educated adults. For example, teachers may try to get their children to stop using "singular *they*". Not surprisingly, efforts like this are mostly futile; children are likely to find it impossible to change their grammar in a way that doesn't match what they hear around them.

Trying to Teach Grammar to Native Speakers

Although their efforts rarely seem to have an effect on children's speaking, in a limited way they may affect the children's written language. This is at least possible when the "rule" amounts to the prohibition of a particular form, for example, *ain't*. When it is more complicated, the prescribing may backfire, leading to behavior that is not what the prescribers had in mind. An interesting example is the use of "subjective" pronouns (*I, he, we*, etc.) vs. "objective" pronouns (*me, him, us*, etc.) in English. In an attempt to prevent usages such as *him and me are friends*, teachers have created a situation in which many speakers now say *with he and I*, instead of *with him and me*, not at all in accordance with the original intention of the teachers. This is an example of **hypercorrection**, which occurs when a prescriptive rule is applied in too many cases. For more on the issue of hypercorrection with English pronouns, see this interview with English professor Jack Lynch, which explains the usage but also takes a prescriptivist approach.

In summary, we've seen two sorts of usages that people try to enforce, usually attempting to prevent an "incorrect" form. In one situation, there is a form perceived as incorrect, such as *ain't*, that is characteristic of speakers of some dialects but not normally part of the speech or writing of speakers of the standard dialect. In another sort of situation, there is a form perceived as incorrect, such as "singular *they*", which is used by many or most of the speakers of the standard dialect as well as by speakers of other dialects. In this second sort of situation, there is often disagreement about what should count as the standard, and it is this kind of case where linguists sometimes become involved because of their interest and expertise in what people actually say or write. (My advice is not to get involved in one of these arguments. For some reason it often seems hard for people, at least for English speakers, to be rational about what counts as "correct" or "incorrect".)

I have focused on English. For a brief history of English "usage", that is, the concern for what is right and what is wrong in standard English, see this passage by E. Ward Gilman from the 1989 edition of *Webster's Dictionary of English Usage*. Not surprisingly, there are similar discussions and debates concerning usage in some other languages, though this sort of concern does not appear to be at all universal.

Other Evaluation of Language

So far we have looked at attempts to enforce linguistic standards on people. But people, especially composition and creative writing teachers, are also involved in evaluations of language based on other kinds of criteria. One is appropriateness, discussed in the last section. Appropriateness is one aspect of language that we continue to learn as adolescents and adults, and it is often explicitly taught. For example, a student might be corrected for using a word or phrase perceived as too casual in certain contexts, say, *be into* in the sense of 'be interested in' in an essay. Or an employee might be corrected by a colleague for calling his boss "dude".

Another important kind of evaluation concerns the **effectiveness** of language. A usage may be grammatical and appropriate to the context but still not accomplish the speaker's or writer's goal. For example, language may be needlessly ambiguous, as in the following example.

2. *Clark broke open the coconut with the screwdriver and then put it back in his pocket.*

A composition teacher might criticize the student's use of *it* in this sentence; does *it* refer to the coconut or the screwdriver? Or a use of language may fail to accomplish a higher-level goal of the speaker or writer. Consider the following example at the beginning of an argument.

3. *Income tax should be abolished. For one thing, people could use the money they pay as taxes for food or housing.*

A teacher might find the writer's point trivial and obvious and suggest leaving it out.

These are all legitimate reasons for evaluating language of course. While linguists and other language scientists are not in the business of evaluating people's language, they can be of help by studying what appropriateness is, what makes an expression interpretable by a hearer or reader, and how the parts of a text relate to one another.

This page titled [1.3: What We Don't Do- Prescribing and Evaluating Language](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.4: Dialects and Languages

Idiolects and Dialects

Exercise 1.4.1

Two Americans are talking about a couple they have just met.

- *She sounded English to me, but he doesn't seem to have any accent at all.*

Two English people are talking about the same couple.

- *He sounded American to me, but she doesn't seem to have any accent at all.*

What's going on here? Who has the accent?

What I know about my language and how to use it is called my **idiolect**. It almost certainly varies in minor ways from the idiolects of all other speakers. But what is an idiolect? That is, what kinds of things do I know? In one sense, this whole book is an answer to that question, but we need to have a first cut at the answer here to help us get started.

- I know words. I have a vocabulary, a set of words which I know how to pronounce and use appropriately. For example, I know how to say the word *apple*, I know that it refers to a particular type of fruit, I come up with this word when I want to refer to a particular apple, and I understand it when I hear it.
- I know how to pronounce words and combinations of words more generally. That is, there are aspects of pronunciation that go beyond individual words. For example, I know to pronounce the ending that we spell *-ed* like /t/ in words like *picked* and *watched* but to pronounce it like /d/ in words like *signed* and *burned*.
- I know how to put words together into sentences in meaningful ways. For example, I know that if I want to ask when a particular train leaves I can say *when does the train leave?*, but not *when leaves the train?*.
- I know how to use language appropriately to achieve my goals. I know that if I want a friend to lend me \$100, it is better to say *I was wondering if you could lend me some money* than to say *give me \$100*.

I'll be much more careful later on about how each of these types of knowledge is described, but for now I'll say (informally) that my idiolect involves knowledge about vocabulary, pronunciation, grammar, and usage.

Of course no one is really interested in describing idiolects. Linguists and other language scientists study the speech of communities of people, not of individuals. More specifically, they study the knowledge of vocabulary, pronunciation, grammar, and usage that is *shared* by the members of a speech community. Because the members of the community agree on this knowledge, because it differs (at least in some ways) from the knowledge shared by other communities, and because it is mostly arbitrary, I will refer to the knowledge as linguistic conventions.

But what is a speech community? I will use this term to refer to any group of people that shares a set of linguistic conventions differing in some noticeable way from the conventions found elsewhere. You may know that in the United States people in some cities have some characteristic features in their pronunciation, although they are easily understood by people elsewhere in the United States. For example, people native to Pittsburgh are known for using *you uns(or yinz)* to mean 'you plural'. Here's an example from the (partly tongue-in-cheek) "Pittsburghese" website: *if yinz wants served, raise your hands*. The number of conventions that distinguish Pittsburghers from other English speakers in the northeastern United States is actually pretty small, but because there is such a set of conventions, we can consider these people to be a speech community. The speech patterns, that is, conventions of vocabulary, pronunciation, grammar, and usage, of a speech community are called a dialect, so we can speak of a "Pittsburgh dialect".

Note that a dialect may not be defined entirely on the basis of its physical location. Cities often contain a variety of ethnic and social groups with different speech patterns. For example, the African-American population of many US cities (for example, Pittsburgh) often has a quite different dialect from the Euro-American population of the same cities.

Which Dialect do you Speak? There may be a Number of Possible Answers.

What about larger communities? Pittsburghers share some speech conventions with speakers in other cities of the northeast and north midwest, for example, their pronunciation of the *a* in a word like *hands*, as in the example above (more on this pronunciation later on). And people in that larger region share some conventions with people in an even larger region encompassing speakers in most of the northern and western United States, for example, their pronunciation of the long English vowels (*bite*, *beat*, *bait*, *boat*, etc.). And people in that even larger region share many conventions with English speakers all over North America, including most of their grammar and usage conventions, as well as a number of pronunciation conventions, for example, the tendency to pronounce the words *latter* and *ladder* in roughly the same way.

This idea of larger and larger communities, each sharing fewer and fewer conventions, is an over-simplification in one sense. The fact is that the boundaries of the communities overlap in many ways. If we look at particular vocabulary, we may find a region with one boundary, whereas if we look at other vocabulary or at some pronunciation convention, we may find another boundary. For example, Pittsburghers tend to say *pop* (as opposed to *soda* or some other word) for carbonated drinks, and they share this convention with many speakers in the northern midwestern cities who also share their pronunciation of the vowel in *hands*, but *not* with speakers to the east of them, in New York City, for example, who share the pronunciation but not the word. (New Yorkers tend to say *soda* rather than *pop*.) Thus where we draw the boundaries around a dialect depends on which convention or set of conventions we're looking at. For more about *soda* vs. *pop*, see this interesting website.

Another way what I've said so far is an over-simplification is that there is great variation within any of these regions. Some of this variation has to do with the constant contact between dialects that is a fact of life in most communities. Some of the variation also has to do with the fact that people often know a range of ways to say things and they may sometimes avoid their local dialect in favor of a standard (see below) in certain situations.

Each of these shared sets of conventions, whether at the level of a small village, a subculture within a city, or a larger region, is a dialect. And a linguist can be interested in describing any level and any aspect of the dialect at any level (pronunciation, vocabulary, grammar, usage). The pronunciation associated with a dialect is called an accent.

Languages

Exercise 1.4.2

What is a language? How would you tell someone (say, an alien with no knowledge of human culture) what English is, without using the word *language*?

We can of course extend the boundaries in our example even further, beyond North America to include England, Scotland, Wales, Ireland, Australia, New Zealand, a number of Caribbean countries, and communities within many other countries. This large speech "community" is not really a community in the usual sense of the word, but it does share many conventions. For example, in all of these places, speakers make a question from a sentence like *he ate potatoes* by inserting the word *did* and changing the form of the verb *ate*: *did he eat potatoes?*, and of course speakers in all of these places share the word *potato* for referring to a class of tuberous vegetables. The conventions of this large "community" are what we refer to as "English", which we consider a language. Thus in one sense a language is a set of dialects. In another sense it is (like a dialect) a set of conventions shared by a speech community.

Two Dialects of One Language or Two Separate Languages?

But how do we decide when a collection of dialects is a language and not just another, more general dialect? As we've already seen, a dialect can also be a set of dialects (the North American English dialect consists of Southern dialect, New England dialect, Canadian dialect, etc.). What makes English a language and not just another very general dialect? What makes Canadian English a dialect of English and not a language in its own right?

The answer to this question is complicated. In fact there is no clear answer because the words *dialect* and *language* are used in different ways for different purposes. There are two completely different kinds of criteria related to the distinction between dialect and language, linguistic criteria and social or political criteria.

Linguistic Criteria

Given two overlapping sets of linguistic conventions associated with two different speech communities, for example, Mexican Spanish and Argentine Spanish, how do we decide whether they should count as two dialects or two separate languages? One criterion is the degree of overlap: how similar are the vocabulary, the pronunciation, the grammar, and the usage? Unfortunately there's no simple way to measure this overlap, at least no way that researchers would agree on. One way to have a sense of the overlap, though, is **mutual intelligibility**, the extent to which speakers from the two or more speech communities can understand each other. Mutual intelligibility is also not easy to measure, and it is often based on the impressions of speakers and hearers, how much they understand when they encounter members of the other group or how long it takes them to get accustomed to the speech of the other group. We also need to establish some sort of intelligibility threshold; no two speakers can be expected to understand each other all of the time. So none of this is precise at all. The idea is simply that if two sets of linguistic conventions are similar enough so that their speakers can usually understand each other, then the two sets of conventions should count as dialects of the same language rather than separate languages. On these grounds, we call Mexican Spanish and Argentine Spanish dialects of the same language (Spanish) because speakers of these dialects normally have little trouble understanding each other.

To find out what should count as a separate language on grounds of mutual intelligibility, a good resource is Ethnologue, an online database of all of the world's known languages, 6,912 according to their current listing. The Ethnologue compilers attempt to use mutual intelligibility to decide what should count as a language. While English is listed as a single language, both German and Italian are listed as multiple languages. Each of these languages, for example, the variety of Italian called Sicilian, is usually referred to as a "dialect", but, according to the *Ethnologue* compilers, these are distinct enough to be considered separate languages. Again, the criterion of mutual intelligibility is a rough one, and some of *Ethnologue's* claims are controversial.

Social and Political Criteria

Another sort of criterion for what counts as a dialect is the social or political unity of the group in question. In Bavaria, a state in southern Germany, and in parts of Austria most people speak a dialect called Bavarian or Austro-Bavarian, which on grounds of mutual intelligibility could be considered a language distinct from the speech of Germans and Austrians in other regions. Ethnologue calls Bavarian a language. But Bavarian is clearly closely *related* to those other dialects and not more closely related to dialects of some other language, and so for mainly political reasons, it is convenient to consider it a dialect of the German "language", rather than a language in its own right. Something similar can be said about the speaking conventions of the older generation in the Ryukyu Islands in southern Japan (because these dialects are dying out, most young people do not speak them). On the basis of mutual intelligibility, we could divide the island dialects into several separate languages, each distinct from the Japanese language (as is done in Ethnologue and in the Wikipedia article on these languages). But the Ryukyu Islands are politically part of Japan, and these dialects are clearly *related* to Japanese and not related at all to any other known language (unless we consider each of them to be a language). So for political reasons, it is convenient to consider them dialects of Japanese, just as the dialect of Osaka is considered a dialect of Japanese.

Mutually Intelligible "Languages"

At the other extreme are examples like the languages spoken in the northern European countries Sweden, Norway, and Denmark. These "languages" are all related to one another, and speakers from some pairs of countries within these have little difficulty understanding one another when they are speaking the standard dialects of their languages, despite the obvious differences, especially in pronunciation. Thus on grounds of mutual intelligibility, we might consider some of these "languages" to be dialects of a single language. But Swedish, Danish, and Norwegian are official "languages" of separate countries, and there are separate spelling conventions for some of the sounds in the languages.

Actually the situation is even more complex than this because Norway has two official dialects, and a fourth related language, Faroese, is spoken in the Faroe Islands, which are administered by Denmark.

So for mainly political reasons, they are considered separate languages rather than dialects of a single language.

To summarize, the line between dialects of one language and separate languages is somewhat arbitrary. However, wherever we draw the line, three points should be clear.

- Every language has multiple dialects.
- Every speaker of every language is also a speaker of at least one dialect of that language.

- Since the pronunciation conventions of a dialect constitute an accent, every speaker of every language speaks with some accent. There is no such thing as "speaking without an accent".

Standard Dialects

Exercise 1.4.3

The following appears on the website of a person who spent some time in Pittsburgh: "probably relating to the rest of Pittsburgh's terrible dialect, which I, fortunately, did not pick up". Why would some dialects be thought of as "terrible"?

Some dialects within a language may be singled out for special status. When we're dealing with a political unit, such as a nation, in which related dialects are spoken by most people, one dialect is often treated as the standard dialect. You know something about this already from the last section of the book. The standard dialect is often the only dialect that is written, and it is the one that is taught in schools and (with some exceptions) used in the media. Thus in Germany, Austria, and the German-speaking part of Switzerland, it is Standard German that is taught in the schools and used in broadcasting, even though most people in this region are not native speakers of the Standard German dialect. This means that most people in the German-speaking countries end up **bidialectal**. The same situation holds in Japan, where it is Standard Japanese, based on Tokyo dialect, that is taught in the schools and used in the media. Note how this makes it possible to speak of a German or Japanese speech community, even when the native dialects of people in these communities are very different from one another, because all educated speakers in these communities share the standard dialect, often as a second dialect.

So what do we mean when we say "German" or "Japanese"? There are two possibilities. "German" could mean Standard German, that is, one of the set of dialects spoken in Germany and also the basis of written German. Or it could mean the collection of related dialects, some mutually unintelligible, which are spoken in Germany and other countries where Standard German is the official language (Austria) or one of the official languages (Switzerland). When linguists refer to "German" or "Japanese", without specifying the dialect, they normally mean the standard dialect.

Dialects of English in the US and England

In the United States, the situation is somewhat simpler than in Germany or Japan because the differences among most of the dialects are not nearly as great; native speakers of English in the United States have little trouble understanding each other. (An important exception is African-American Vernacular English (AAVE), spoken mainly by many African-Americans.) As in Germany and Japan, we have an (informal) standard dialect, for vocabulary, grammar, and usage, if not for pronunciation. Thus children in Pittsburgh learn in school to write sentences like *the school needs to be renovated* rather than *the school needs renovated*, which would be grammatical in their local dialect. Americans tend to be relatively tolerant of differences in *accent*, however. Teachers in schools throughout the country teach the standard grammar but use their own local pronunciation. If we have a standard accent, it is the one people associate with television announcers, the accent characteristic of much of the Midwest and the West. This accent is called General American. I will have more to say about it later.

The situation in England is similar to that in the United States, and the standard vocabulary, grammar, and usage that children learn to write in English schools are very similar to the American standard. However, in England, there is a stronger idea of a standard accent than in the United States and more pressure for children to learn this accent if it differs from their home accent. This accent is referred to as **Received Pronunciation** (RP); it is based on the speech of educated speakers in southern England. (Note that RP is standard *English* English pronunciation, not *British* English; in Scotland, there is a quite different standard accent.) I'll have more to say about RP and how it differs from General American and other English accents in this section.

"Just" a "Dialect" or a Full-blown "Language"

The existence of a single standard dialect among a set of non-standard dialects has important social implications. The non-standard dialects have less prestige, and their use may be discouraged in formal situations, not just situations in which writing is called for. Sometimes, as in the Ryukyu Islands in Japan or in some regions of France and Spain, this leads to the decline and possible death of the non-standard "dialects" (which would be considered languages by the mutual intelligibility criterion). In other situations, speakers of non-standard dialects retain pride in their local speech patterns, while recognizing that they are not appropriate in certain situations. Finally, this pride, along with other cultural differences separating the speakers of the non-standard dialect from the speakers of other dialects (non-standard or standard), may lead to pressure to have non-standard dialects given official status,

especially if they differ significantly from the standard. At this point the words *dialect* and *language* become politically charged terms because the supporters of official status for the non-standard dialect may feel the need to argue that it is not "just" a dialect of the larger language but rather a language in its own right. This has happened in the United States with AAVE (here is an essay on this topic by the sociolinguist John Rickford) and in Europe with many languages that are normally considered "dialects" of other languages (this website includes many of them as well as links to other sites concerned with the "minority language" question in Europe and elsewhere).

Language Families

Why Some Languages Resemble Each Other

We've seen how as we extend the boundaries of speech communities, we get fewer and fewer shared conventions. When we reach the level of a language such as English, Spanish, or Mandarin Chinese, we have a speech community which shares a set of conventions (in some cases a standard dialect) which allows people in the community to communicate with one another despite dialect differences. But we can go beyond a language. So for English, we could extend the boundaries to include the Netherlands, Germany, Scandinavia, and some other regions in western Europe. We'd now find a much smaller set of shared conventions. All speakers in this large "community", for example, share a word meaning 'all' which is similar in pronunciation to the English word *all*. But there would be no reason to call this set of conventions a "language" since the speakers obviously do not understand each other and do not belong to a single political unit with a single standard dialect. Instead we refer to this set of conventions, or set of languages, as a language family, in this case, the Germanic languages. The members of a language family resemble each other because they are **genetically** related; that is, historically they derived from a common ancestor language. (Note that this use of the word *genetic* differs somewhat from its use in biology; the speakers of Germanic languages are not necessarily genetically closer to one another than they are to the speakers of other languages.) The ancestor of the modern Germanic languages was not a written language, so we can only infer what it was like.

In most cases we can go even further back; the ancestor languages of two or more families themselves may have had a common ancestor language. Thus the modern Romance languages, including Spanish, French, Italian, Portuguese, Catalan, and Romanian; the modern Germanic languages; and many other languages spoken today in Europe, the Middle East, and South Asia, apparently descended from a much older (and also unwritten) language. This means we can group all of these languages into a single family, in this case the one we call **Indo-European**. Sometimes, to distinguish the lower from the higher levels within a family tree of languages, we use "language family" only for the largest grouping (for example, Indo-European) and "branch" to refer to groupings within this (for example, Germanic and Romance). Note that there may be many intermediate levels in the family tree of languages. Within Germanic, for example, there is North, including the Scandinavian languages, and West, including English, Dutch, and German.

Note also that languages may resemble each other in one way or another for reasons other than a genetic relationship. The main non-genetic source of similarity is **language contact**; when the speech communities for two language are in close cultural contact, their languages often influence one another. So modern Japanese vocabulary includes thousands of words borrowed from Chinese and uses the Chinese writing system (as well as writing systems specific to Japanese). But, except in the sense that all human languages may be ultimately related to one another, there is no evidence that Japanese is genetically related to Chinese. A more complicated situation occurred in Western Asia with the complicated cultural influences among people speaking Arabic, Persian, and Turkish. These three languages belong to separate language families (Afro-Asiatic, Indo-European, and Altaic, respectively), which are either unrelated to one another or only very distantly related, but Turkish and Persian have borrowed many words from Arabic, Turkish has also borrowed many words from Persian, and Persian borrowed its writing system from Arabic.

This page titled [1.4: Dialects and Languages](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.5: A Stand and Two Themes

A Stand

Linguistics is a relatively young science, and psycholinguistics and computational linguistics are even younger, so it's not surprising that these fields are still fraught with controversy. Though there is general agreement on some core topics, some quite basic issues are still up in the air. It is not the place of an introductory textbook to go into all of the controversies; I will only try to do this for one topic, language learning, perhaps the most hotly debated of all. But it is impossible to discuss any topic related to language without taking some sort of stand on the questions that divide language scientists. All linguistics textbooks take such a stand, though, sadly, few of them tell you that they are doing this. The stand taken in this book is that language and language behavior are not phenomena separate from the rest of human perceiving, acting, and reasoning; that we can only understand how language works by understanding how it fits into the rest of human behavior. Another way to think of this position is that we will be treating the language sciences as **cognitive sciences**.

This position is related to two of the ways of looking at and studying language that we've already discussed. As noted in the overview of the book, we can study language as a system independent of the people who use it, or we can study language behavior, focusing on use. As noted in the section on how language is studied, we can be concerned with product or process. The stand that this book takes implies that it cannot be enough to study language as a system in its own right (though it might be useful some of the time to do this) or to study product rather than process (though again it may be useful some of the time). It also implies that the border between language and non-language is not necessarily a clear one and not something we should spend much time worrying about. So for every topic covered in this book, I will try to look at it both ways, to see how the system/product perspective and behavior/process perspective both contribute to an understanding of how language works.

I didn't invent this position of course. A number of other people have defended one or another aspect of it. On the need to consider cognitive processes outside of language in order to understand language, cognitive linguists like Ronald Langacker and George Lakoff have made this case in many of their writings, in which they invoke general psychological notions such as attention, categorization, and memory. On the need to treat language as process, a very convincing argument was made in 1983 by the computer scientist Terry Winograd in his important book *Language as a Cognitive Process* (1983).

In the rest of this section, I discuss two themes that are consistent with this stand and that will guide the rest of the book.

Two Themes

Meaning, Function, and Form

The Main Theme of the book is that language associates **meaning** and **function** with **form** and that understanding how language works requires that we focus on these associations. All three of these terms will become clearer later in this book. For now we will just try to get the gist of what's intended by them.

Linguistic **form** concerns the way a linguistic expression sounds (for spoken language) or looks (for written or signed language) or how a linguistic expression is produced. There are two ways of looking at form. The usual way within linguistics is to think of form as a sequence of elements. For written language this is a natural way to treat form because written sentences *are* sequences of elements (characters). For spoken or signed language, it is less so. Treating a stretch of spoken or signed language as a sequence of elements means in effect focusing on the written transcription of the language in terms of units of sound or primitive gestures of some sort. This approach tends to play down the role of time, to look at a stretch of language as a static object. An alternative and more radical way of looking at form (at least for spoken or signed language) is to treat a stretch of language as something inherently continuous and dynamic, something that cannot (or should not) be transcribed as a written sequence of elements. In this book, I will usually discuss linguistic form in the first way, mostly because this is the way it is usually done and because it seems to simplify the analysis. But I will also insert periodic warnings about the biases that are built into this approach. One particular bias we should be aware of is one based on literacy; that is, as skilled users of alphabets, we may tend to think of spoken and signed language in this way (for an argument of this sort, see this paper by linguist Robert F. Port).

I will use **meaning** to refer to what language is *about*, the concepts that words and linguistic patterns refer to. By **concept** I will mean a unit of cognitive experience, a way people have of abstracting over their experiences in the world. For example, in a baby's experience the same face keeps appearing over it, and the baby abstracts over these different occurrences of face-appearing the

concept of daddy. In this book labels for concepts will often appear like this. If the nature of form in the study of language is controversial, the nature of meaning is even more so. Some people deny that there is such a thing as meaning or that talking about meaning is helpful or that there is a consistent way to define "meaning" or "concept" or "about." I will try to take these positions seriously in this book, but the fact is that I don't know how to even get started without looking at the pole of language that is opposed to form.

Concepts without Words

Whatever we mean by concepts, it is clear that not all of them are associated with words. Words are linguistic; concepts need not be. We all have concepts that we have no words or grammatical patterns for. For example, one concept I have is the little depressed region bounded by two vertical ridges that is found between the nose and the middle of the upper lip. I have no idea what to call that place (though I'm pretty sure it has a name in at least some languages); that concept does not represent the meaning of a word for me.

Another point to note here is that meanings do not need to be seen as thing that are "out there" in the world. Since the stand taken in this book is that language is basically a cognitive phenomenon, this would be a strange way to think about meaning. Instead I will be treating meanings as things that are "in here", cognitive entities realized in our brains, depending on our *interpretation* of what's out there and including imagined entities that aren't "out there" at all (though they are inevitably based on things that are).

I will use **function** in two senses, first, for the uses that people put language to and second, for the uses that particular words or patterns have within stretches of language. It is the former sense that concerns us here. People use language to refer, to assert, to command, to convince, to get information, to entertain, to deceive, and much more, and these uses of language obviously have something to do with the forms that they choose. (If you want somebody to lend you their computer, you don't say *you will lend me your computer*, you say something like *could I borrow your computer?*.)

Many people who study language make a distinction between function in this sense (part of what they call "pragmatics") and meaning (what they may call "semantics"). I would argue instead that the difference is a matter of degree, but that there's not much point in worrying about where the line between pragmatics and semantics is.

Constraints

A second theme of the book is that language is the way it is in part because of **constraints** coming from the nature of human biology and human cognition. A constraint is a kind of limitation on what is possible. Consider constraints that come from the nature of the human body itself. For spoken language, the physical and physiological properties of the vocal tract constrain what can be produced, and the auditory system constrains what can be perceived. For sign language, the physical and physiological properties of the hands, arms, upper body, and face constrain what can be produced, and the visual system constrains what can be perceived.

Just as important are **cognitive** constraints, constraints arising from the nature of the human mind. First, there are limitations on human **memory**. Cognitive scientists divide human memory into **short-term** and **long-term memory**. Short-term memory is used to temporarily store the information that is needed as we interpret the world around us. When we are faced with a complicated visual scene, we can't look at everything at once, and we have to scan the scene in order to come up with some kind of interpretation. But we need to temporarily store some sort of record of what we've already seen in the scene as we do this. That information is kept in short-term memory. Likewise, when we are listening to a sentence, we need to remember the parts that we've already heard as we are trying figure out the meaning of the whole. Short-term memory is used for this as well. The same applies to reading and to the visual processing of sign language. Cognitive science research has shown that our short-term memory has a very restricted capacity, and this places strong constraints on how sentences and discourses are organized.

We use our long-term memory to store information that we learn through experience. Information may remain in long-term memory for a very long time, even indefinitely, and human long-term memory is very large, larger than the memory of any current computer. But long-term memory is still limited. We obviously do not have an infinite memory capacity, and just because we have stored an item in memory doesn't mean we can retrieve it when we need it. We will see later how the finiteness of long-term memory matters for how language works.

Since languages have to be learned, limitations on human learning are obviously also relevant. This is perhaps the most complex and controversial topic of all. Everyone at least agrees that language must be learnable because children obviously learn it. That is,

it must be possible to figure out the forms, the functions, the meanings, and the associations between them on the basis of the examples of language that are available to young children. It turns out that coming up with an explanation for how this might happen has proven to be very challenging though.

Speaker Orientation and Hearer Orientation

Finally there are constraints that are specific to the two ways in which language is processed, language production, that is, speaking, writing, and signing, and language comprehension, that is, listening, reading, and interpreting linguistic signs. To simplify matters, I will be referring to a language producer as the **Speaker**, even though writers and signers as well as speakers proper are intended. And I will be referring to a language comprehender as the **Hearer**, even though readers and sign interpreters as well as hearers proper are included. (To remind ourselves of these distinctions, I will capitalize Speaker and Hearer when they are used in this way.)

For the Speaker, the main constraint is that the production of linguistic forms be easy. It is easier for the Speaker to make relatively few distinctions because the Speaker has to remember what the distinctions are and to make the extra effort to keep things distinct. For example, maintaining the agreement between subject and verb in the present tense in English (*the girl **sings**, the girls **sing***) requires an effort on the part of the Speaker. What is easy for the Speaker depends in turn on constraints from the body. For example, large movements of the tongue tip are more difficult to execute than short movements, so short movements should be preferred from the Speaker's perspective.

For the Hearer, the main constraint is that linguistic forms that need to be distinguished can be easily distinguished in comprehension. This constraint also depends on the body, specifically the parts of the nervous system that are responsible for hearing (for spoken language) and vision (for written and signed language). For example, if a language contains a large number of homophones, that is, different words which sound the same (such as *two*, *too*, and *to* in English), this may put a burden on the Hearer.

Speaker-oriented constraints and Hearer-oriented constraints often oppose each other; what simplifies things for the Speaker (for example, not making many distinctions) complicates things for the Hearer. We will see many examples in the book of these opposing tendencies.

The opposition of Speaker orientation and Hearer orientation is particularly clear as languages change. Languages change for a variety of reasons — contact with other languages, imperfect learning by children, random fluctuation — but it appears that all languages are always changing. Most of the changes can be seen as either Speaker-oriented or Hearer-oriented. For example, the grammar of a language may become simplified as some suffixes are dropped, a change that seems to result in less work for the Speaker. But the two kinds of trends always balance each other out in the end, and the simplification of the grammar in one way will probably be compensated for by an increase in complexity (from the perspective of the Speaker) somewhere else in the grammar of the language. Otherwise language would fail as a communicative device. These built-in pressures in favor of the Speaker and the Hearer apparently prevent the world's languages from moving in some general overall direction. That is, at least in recent history, it does not appear that languages have generally been getting simpler or more complicated (in any sense of these words) as they evolve.

This page titled [1.5: A Stand and Two Themes](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.6: Why Study Language

What good is the scientific study of language? Why does anyone do it? Why should you care about it? These are the sorts of questions you have a right to ask about any university course. The answer to the last question depends a lot, of course, on how you happened to end up in a course using this book in the first place and on what your interests and long-term goals are. Language is a part of everyone's life, but it is more central to some people than to others. But I happen to believe that a scientific look at language should be a part of the basic curriculum, like mathematics and history are.

Second-language Learning

Many of you have already studied one or more languages other than your first, and more of you will later on. A few of you may teach a foreign language. In either case, you are not likely to find the learning process an easy one. Some of the difficulties faced by second-language learners have to do with differences between their first and second languages, differences in pronunciation, vocabulary, grammar, and usage. You do not have to know linguistics to learn a second language; after all, people all over the world who have never heard of linguistics do this successfully all the time. However, knowing what pronunciation, vocabulary, grammar, and usage are; how they can differ between languages; and how people seem to learn them (as first- or second-language learners) can help you be aware of and understand your problems and possibly correct some of them. A second-language teacher needs to be able to focus on problem areas, for example, by giving lots of practice or by simplifying other aspects of the language being learned. It is difficult, if not impossible, both to understand the source of the problem and to come up with ways of addressing it without understanding the nature of the material being learned, that is, what linguists and other language scientists study.

First-language Learning

Some students, when beginning a linguistics course, believe it will help them improve their knowledge of their first language. In fact this is *not* something you can expect from a linguistics course. You already know the great majority of the words, the grammatical patterns, and the usage conventions that you will need to survive. Of course you can improve; you can learn new words and expressions, become more proficient with the grammatical patterns that are part of formal language, and get better at using language to accomplish your goals. Most of the improvement should come naturally as you are exposed to the complex language of academia and the workplace. But you will also face evaluations of your language by other people — teachers, colleagues, supervisors, even family and friends — throughout your life. You will be in a better position to make use of this criticism and advice if you understand what sort of problem is involved (if there really is a problem) and how it fits into the larger scheme of things. This is where linguistics can help.

Another way in which most people encounter first-language learning is in raising children. Of course you don't have to know linguistics to know how to "teach" a baby a language; babies aren't really taught language anyway. But knowing what it is that babies learn when they learn a language can make the process more enjoyable. You'll be able to better appreciate what an amazing process learning a language is, why so many people are fascinated with how this process happens, how your baby does the same kinds of things that others do and, at the same time, how your baby's learning steps are unique.

You may end up in a job that involves first-language learning more directly, as a teacher of your first language to native speakers. In this capacity, part of your job will probably be to make sure your students are competent in the standard dialect of the language. As we saw in the section on what linguists study and the section on prescription, sorting out what belongs to the standard is not a trivial matter, and a knowledge of what linguistic conventions are seems essential. As I've already said, there is a lot of confusion and controversy about what should be emphasized or even taught at all. You will also need to teach your students what counts as appropriate and effective language. Again linguistics and other language sciences can help; some language scientists devote their efforts to figuring out what makes different expressions appropriate in different situations, while others are concerned with how words and expressions are interpreted by hearers and readers.

Or you might work as a speech therapist, dealing with people with speech disabilities of one kind or another. Here the relevance of the scientific study of language is obvious; you first need to know what the norms of a language are before you can hope to address the ways in which your clients or patients deviate from these norms.

Cognition

Language is probably the best window we have on the workings of the human mind. Language gives us the extraordinary ability to describe the contents of our thoughts, an ability that no other animal has. Of course there are many unconscious aspects to cognition that we cannot talk about, but these properties are apparently also reflected in what we say. The units of language — elements of form, words, grammatical patterns, conventions of usage — are in some sense also units of cognition. The implication is that the study of what all languages share is also the study of what it is to be human, something that is certainly an important topic for any educated person.

Linguistics and Tolerance

What we can learn about the human mind by studying how languages *differ* from one another is more controversial. Linguists and other cognitive scientists disagree on how deeply the nature of a person's first language influences how the person thinks and views the world around them. I'll have more to say on this topic at various points throughout the book. But language is such an important part of our lives that learning about the languages of other people, including how those languages differ from ours, is in a very real sense learning about those people. As with any other aspect of culture, lack of knowledge can lead to intolerance. It is easy to believe that other languages are inferior in one way or another to ours, to think that some languages, especially the languages of relatively small ethnic groups in the Third World, are more primitive than others. So the realization that a language like Tzeltal, Lingala, Amharic, or Inuktitut has a set of grammatical categories and communicative options not even found in English is an eye-opening experience. Looking more closely at languages, and in particular at languages that might seem exotic to us, can make us more tolerant.

Finally, language is what we use to influence the beliefs of one another. This happens in everyday conversation, as we argue about who forgot to put the mayonnaise jar back in the refrigerator or what the results of the latest election mean. It happens in advertising, as companies do their best to get us to buy their toothpaste, shoes, and cars. It happens in education, as educators provide us with what they say are truths and try to convince us that knowing how to speak Spanish or how to do a t-test will be useful to us later in life. It happens in politics, as politicians and political activists try to get us to vote for them or to support their program. None of this is new, but the enormous quantity of information that is now available to most of us *is* new, and most of this information is designed to change our beliefs in some way. Obviously an educated person needs to be able to navigate their way through all of this, to sort out the nonsense, to see how bias and ideology are behind what is being claimed, to be a critical reader and listener, to make informed decisions. I don't believe that any of this is possible without understanding the role that language plays in knowledge, belief, and persuasion.

This page titled [1.6: Why Study Language](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.7: Problems

1.8.1 Describing and Explaining, Prescribing and Evaluating

1.8.1.1

Note

This problem set assumes some basic knowledge of English dialects, so if you are not a native speaker of English, you will probably want to collaborate with someone who is.

Each of the following includes two ways to say (at least roughly) the same thing. For at least some people, the first way (A) could be seen as "wrong" in some sense. We saw in these two sections of the book that there are various ways in which a particular linguistic form can be thought of as "wrong" (though not necessarily by linguists).

1. The use is not grammatical or otherwise acceptable (in the case of the pronunciation of a word) in any dialect of the language.
2. The use is grammatical or acceptable in some non-standard dialects, but not grammatical in the standard dialect.
3. The use is grammatical or otherwise acceptable for most speakers of the standard dialect, but some speakers may find it "wrong" for one reason or another.
4. The use is grammatical or acceptable but not appropriate in the context.
5. The use is grammatical and appropriate but not effective.

For each of the following, decide which type of "mistake" is involved in A (as compared to B), and explain your answer in a sentence or two. There may not be a single right answer for some of the problems.

1. A: *nuclear* pronounced noo-kyuh-lur
B: *nuclear* pronounced noo-klee-ur
2. A: *yellow* pronounced yell-uh
B: *yellow* pronounced yell-oh
3. A: *Mr. President, what you're saying sucks.*
B: *Mr. President, I can't accept what you're saying.*
4. A: *I come back last night.*
B: *I came back last night.*
5. A: *Next shake the test tube that you dissolved the salt in.*
B: *Next shake the test tube in which you dissolved the salt.*
6. A: *Whoever it belongs to should claim it.*
B: *Whomever it belongs to should claim it.*
7. A: *He left glass water on table.*
B: *He left a glass of water on the table.*
8. A: *When I exited the bed this morning, my hair looked like rabbit ears.* (an eight-year-old speaking to his mother; this is a real example)
B: *When I got out of bed this morning, my hair looked like rabbit ears.*
9. A: *John ordered his meal from the waiter, and he brought it to him right away.*
B: *John ordered his meal from the waiter, and the waiter brought it to him right away.*
10. A: *How did y'all like the party?*
B: *How did you like the party?*
11. A: *I would be thankful if it could be let known to me whether taking a student under your fold and tutelage for a long and hopefully eventful partnership would be acceptable to you.* (a real example from an e-mail to me)
B: *Please let me know if you can take on a new student.*

1.8.1.2

You are a linguist studying the grammar of a particular dialect of English. You have recorded a stretch of speech between two men who are speakers of the dialect, and now you are taking examples from the recording to use as data. Included in the recording is the

following statement by one of the speakers.

- *You don't even realizing it.*

You are trying to decide whether your account of the grammar of this dialect should include this sentence. Unfortunately you no longer have access to the two speakers, so you can't ask them questions. For each of the following, say whether it would be a relevant thing for you to do in order to help make your decision. Explain each answer.

1. You consult a book written by another linguist about the grammar of a closely related dialect.
2. You listen for similar sentences in the rest of your recording.
3. You ask more educated speakers from the same region as the two men whether they could say something like this.
4. You listen to what immediately follows the example on the recording to see if the speaker corrected himself.

1.8.2 How We Study Language

In this section, we saw that ambiguity is an example of a phenomenon in language that can best be appreciated from the perspective of comprehension. Consider the English words that are written *bank*, meaning roughly 'side of a river,' 'kind of financial institution,' and 'tilt while making a turn' (there are other meanings as well). Because all of these are written (and pronounced) the same, a reader (or listener) has to disambiguate the form when it appears. But this can be done in different ways, using different kinds of information. For each of the following sentences, say what kind of information a reader could use to disambiguate the word *bank*, and rank the three for how difficult they would probably be for a computer program. Imagine that the person or program knows only the three meanings of *bank* given above.

1. *I don't have much my money in my account at the **bank**.*
2. *Before landing, the plane will have to **bank** to the right.*
3. *When I climbed out of the boat, all of my money fell onto the **bank**.*

1.8.3 Two Themes

For each of the following facts about human language, say whether it derives from **Speaker (or Learner) orientation** or **Hearer orientation**, that is, whether it makes things easier for the Speaker, the Hearer, or the Learner. Explain your answer in a sentence. There may be more than one possible answer.

1. *Fos* is a more likely word in a human language than *zkt* is.
2. Languages tend not to have many homophones, that is, words that sound the same but have completely different meanings such as *bear* meaning the animal and *bear* meaning 'tolerate'.
3. Languages do not have specific names for every object; instead they have names for categories of objects.
4. In sign languages, signs are normally not made with unrelated movements of the two hands, for example, with one hand moving up and down while the other moves sideways.
5. Apparently all human languages have a way of distinguishing statements from questions.
6. Some languages have only three vowels. If so, these vowels tend to be pronounced quite differently from one another, usually similar to the vowels in *hot*, *heat*, and *hoot*.

This page titled [1.7: Problems](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

2: Word Meanings

One point all linguists probably agree on is the centrality of words to language. All aspects of language are tied in some way or other to words. In this chapter, we'll start off imagining a world with no language at all and see what is gained by adding just this one basic feature of human language. Words have two aspects, their forms and their meanings, and in this chapter we'll only look at meaning. Actually we'll only be considering words in one category, those words that refer to things in the world.

[2.0: Introduction to Word Meanings](#)

[2.1: Reference and Proper Nouns](#)

[2.2: Categories and Common Nouns](#)

[2.3: Word Senses and Taxonomies](#)

[2.4: Metaphor and Metonymy](#)

[2.5: Deixis and Person](#)

[2.6: Lexical Differences Among Languages](#)

[2.7: Learning Meaning](#)

[2.8: Problems](#)

This page titled [2: Word Meanings](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.0: Introduction to Word Meanings

One point all linguists probably agree on is the centrality of words to language. All aspects of language are tied in some way or other to words. In this chapter, we'll start off imagining a world with no language at all and see what is gained by adding just this one basic feature of human language. Words have two aspects, their forms and their meanings, and in this chapter we'll only look at meaning. Actually we'll only be considering words in one category, those words that refer to things in the world. How people use words to refer is just one aspect of the question of what it means to mean, which turns out to be an enormously complicated topic, one where linguists and other cognitive scientists still have a long way to go. In a way the book is starting off with the hardest topic of all. But the idea of meaning is at the heart of what language is, so we can't really put it off. Even just scratching the surface of this topic, as I'll do in this chapter, will lead us to look at notions that seem to be beyond language: how people categorize objects in the world and how people use one kind of situation to help them understand another kind of situation. But to say anything at all about the meanings of words seems to require an account of where those meanings come from and what good they are for us.

This page titled [2.0: Introduction to Word Meanings](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.1: Reference and Proper Nouns

A World Before Language

Imagine a tribe of people much like us. You can think of them as our hominid ancestors, or if you prefer, as a race of aliens on another planet; for our purposes, it doesn't really matter since we will only be using them to help us understand what language does. Like us, they live in a world full of regularity. There are objects around them with characteristic properties: rocks that tend to be hard but to break when they are hit hard enough; plants that tend to be soft and have parts that can be pulled off and sometimes eaten; animals that move and can be dangerous. (We'll let our people be vegetarians.) The world is also regular in terms of what happens: the regular appearance of celestial bodies and seasons; the inevitability of rain following certain kinds of clouds; the way certain animals run away and others attack when they are confronted; the way other members of the tribe respond to friendliness, flirtation, rejection, and aggression.

To survive in their world, these people, like us, have evolved nervous systems that allow them to find these regularities in their world. They can sense the basic features of the world: colors, textures, edges, movements, sounds, tastes, smells, consistencies, hardnesses. Like us, they are also expert learners. This allows them, together with their basic sensory abilities, to *discover* the regularity in their world. Learning is important for our people because it allows them to survive in very different environments; no matter where they happen to end up — the desert, the mountains, the forest, the plains, the coast — they can figure out what their world is like. Finding the regularity is important in turn because it allows them to learn what kind of behavior is appropriate. If you can identify the edible plants and the animals you need to run away from, you can survive. Even if you can't identify the edible plants and the threatening animals, you may still be able to survive by getting help from others in your tribe. But to get help from others, you have to have figured out what kinds of behaviors lead them to cooperate with you and what kinds lead them to reject or even attack you. If you can't recognize aggression in the face and body of a much stronger member of your tribe, you won't survive.

These people are like us in a number of ways. But they differ from us in one very important way. They do not have language, at least language as we know it. Throughout this book, we will be giving them language, a little at a time, at each step asking what advantages and complexities the new features bring. For now we will refer to them as "Prelings", to emphasize that they are pre-linguistic; they don't have language yet.

While our tribe of Prelings is purely speculative, you are already familiar with a real group of pre-linguistic humans. All people at birth are pre-linguistic in that they do not yet know a language. We can also learn a lot about language by taking the perspective of these real pre-linguistic beings, and we will do this in the book sometimes too.

Objects and Individuals

Exercise 2.1.1

Consider the things in the world that we refer to using the following words:

1. *box, bird, road, hand*
2. *milk, dirt, air*

How is the first group different from the second?

One of the things that both our imaginary Prelings and our real pre-linguistic infant notice right away about their world is that it contains **objects**, entities with clear boundaries around them that have a number of stable properties. For example, consider an apple, one of the things that both our Prelings and our infant happen to encounter. While some things about an apple can change quickly, its position, for example, others are relatively stable: its color, shape, size, smell, and consistency, for example. The same could be said of a cave, a cat, a bush, or a person (or another Preling). (Note that we are using the word *object* in a way that differs from the usual, non-technical use of the word in that it can refer to living things, including people, as well as inanimate things such as rocks and plants. This will happen often in the book: the use of a word that you are already familiar with in an unfamiliar, somewhat different sense.)

Actually it's an over-simplification to say that infants recognize objects. The ability to recognize objects actually develops during the first months of life. For some pointers to research on this area, see the work of psychologist Scott Johnson.

Lines Between Categories may not be Easy to Define.

But not everything in the environment is an object. The water in some portion of a river is not an object because it has no obvious boundaries. And the situation in which a lizard is resting on top of a rock is even less like an object because it is not stable; the lizard may run away at any minute. What about a river or a mountain? We can view a river as an object if we allow a boundary to be the border separating two different kinds of material, soil and water in this case. And we can view a mountain as an object if we allow a boundary to be clear on some sides (where the top and sides of the mountain are separated from the air) but not so clear on others (where the mountain is separated from the plain around it). These examples are important because they illustrate a common theme: whenever we attempt to divide up the things in the world in a particular way, for example, into objects vs. non-objects, we will find good examples — an apple is a good example of an object; some water is a good example of a non-object — and others, such as lakes, that need to be stretched to fit into one group or the other.

Prelings and babies are not only good at finding objects; they are good at distinguishing particular objects from one another. A baby learns early on to distinguish its parents from other people. Prelings need to be able to distinguish members of their tribe from one another — their mate, their children, the tribe chief — in order to know how to behave. And they need to be able to distinguish particular bushes and particular hills and bends in the river from each other so that they can know where they are when they are out searching for food. We will refer to particular objects as **individuals**. Later we'll use this word for particular instances of things other than objects.

Proper Nouns

Exercise 2.1.2

One function of words is to "point" at things. But what advantages do words have over pointing with our fingers?

Our pre-linguistic Prelings are social creatures. They survive by sharing the food they find, by helping each other take care of their young, and by protecting each other from predators. But they are limited in how much they can cooperate because the few signals they have communicate only vague meanings such 'danger!' and 'I am angry' and 'I want to mate.' They can also point to individual objects in order to draw the attention of other members of the tribe to them. But when an individual object is out of sight, pointing won't work.

Consider what the advantages would be of being able to "point" to something out of sight. If one member of the tribe was in danger, another member could let the others know about this by "pointing" to the one in danger. Of course the pointing alone wouldn't convey the danger; the others would have to infer this from the urgency of the pointer's actions. Or the Preling could "point" to a familiar, but out-of-sight, landmark in the tribe's environment, alerting the others that something important had turned up there, such as a previously unknown source of food.

Pointing with Words

Language offers several ways of "pointing" with words. One simple way is with words which have only this pointing function, words like *this* and *that*. We'll learn more about these kinds of words later. For now, note that they have the disadvantage that they say little about what object is being pointed to and in fact have to be used together with other language to make any sense: *pick that up, this is my sister*.

A more precise alternative is provided by names. They allow us to point to individuals even when they are not there. Nearly all human languages have a category of words we'll call **proper nouns** that function as names. English examples include *Mommy*, *George*, *Fluffy*, *Indiana*. In many languages the category of proper nouns is distinguished formally from other expressions in certain ways. In English, for example, we don't normally precede proper nouns with the word *the*, unless more than one individual has the name and we are trying to make it clear which one we are talking about. That is, we do not say *please give this to the Clark* (whereas we can say *please give this to the coach*). This will be a familiar theme in the book: a particular notion or function of language (for example, naming) may be represented by a type of word or linguistic pattern which is distinguished formally from other words or patterns (for example, proper nouns). This is just another example of our Main Theme, that language is about the relation between meaning/function and form. (Note that proper nouns are not the only possible names in English; we also have expressions like *the White House* and *the Grand Canyon*. We'll look at these kinds of names later.)

Reference

The pointing function that names perform is called **reference**. Most of the time I'll be using this word (somewhat loosely) both for what Speakers do with words and what the words themselves "do". The individual that is referred to is called the **referent**. I will have a lot more to say about reference, especially in the chapter on Composition. For now, we can think of reference in the following way.

1. The Speaker has an individual in mind and wants the Hearer to have that same individual in mind.
2. The Speaker produces the name.
3. The Hearer recognizes the name, and it activates the representation of the individual in the Hearer's short-term or long-term memory.

Reference by itself can never make the Speaker's complete intention clear, that is, why the Speaker would want to call the Hearer's attention to the referent in the first place. Reference is only one aspect of human language, and Speakers have access to more than just names. But even with the help of a full-blown modern human language, Speakers can never directly transfer what's in their mind to the minds of Hearers. There's *always* some guesswork involved in interpreting language; this is a topic we'll return to when we discuss how combinations of words are interpreted. In other words, what our Prelings would have to do in understanding their primitive "language", consisting only of proper nouns, is only really different in degree from what we have to do all the time.

Note that a subclass of proper nouns, names of people, have another function in modern languages. We can use them to *call* the people that they name. In English, we do this by simply saying the name of the person, often loudly. Since this function differs in some ways from what we're calling reference, it should not surprise is that it differs in *form* in some languages. The special form that a name takes when it is used for calling in these languages is called the **vocative**. Note that the calling function is in some ways even simpler than the reference function since it applies only to other people (or animals that we treat like people) and it only appropriate when the person being called is present. So we can imagine the Prelings coming up with this use of language even earlier than they come up with reference.

Knowing and Using Names

If there are only a small number of interesting things to be named, it might be efficient to have the list of names built into the Prelings so they wouldn't have to figure them out in every new generation. For this solution to work for real animals, however, evolution would somehow have to settle on it; that is, the knowledge (or whatever) that is required would have to be genetic in the end. And for this to happen, members of the species that are genetically predisposed to produce particular sounds (or gestures) and respond appropriately to them would have to have an advantage over those that aren't. The problem with this possibility is that the knowledge required to produce and understand names is quite complex, and it is very difficult to imagine how anybody could have ended up with a genetic predisposition to have this knowledge.


Let's consider what it would take for a Preling to know a particular name. First, it has to be able to both recognize and produce the *form* for the name, a sound like "Clark" or some unique gesture that refers to Clark. This alone is not at all trivial, but we'll wait until the next chapter to worry about why. For now, let's just assume that there is some sort of representation for the form of the word in the Preling's brain.

"Knowing" the Meaning of a Name

The Preling would also have to be able to recognize the referent of the name, the person, village, or mountain that is referred to. In a sense, then, it has to "know" the referent. But what is "know" exactly? Let's consider two possibilities. First, there is a single "place" in the Preling's long-term memory for the referent, a place that somehow gets activated when the Preling sees, thinks about, or (if it knows the word) hears about the referent. This place brings together all of the various features of the referent, each with its own place in memory, as well as other places in memory that are associated with particular responses the Preling makes when it encounters the referent. This is what I'll call a **localized** representation. Another possibility is a **distributed** representation of the referent. Instead of appearing in one place in memory, there are only the separate places for all of the different features and responses that are associated with the referent. These are connected with one another in a large network of relationships in such a way that when enough of them are activated, the whole set of places that are associated with the referent become active. In the distributed view, there is no single place in memory that gets activated when the Preling has the referent in mind. Does it make a difference which of these ways of representing the referent we go with? We'll see later that it does have some implications. For

now, for the sake of simplicity, we'll assume the localist position, which is the one that is assumed (though usually not explicitly) by most linguists and by many other language scientists.

The picture so far, then, looks like that in the figure below. The large rectangles are meant to stand for some person's representation of two individuals (Clark and Lois) and the forms of the words that refer to them ("Clark" and "Lois"). The arrows connecting the pairs of rectangles represent the form-meaning relationship. The small squares above the boxes are supposed to suggest all the different features associated with Clark and Lois (their height, their personality, etc.) and the different responses that this person might have to Clark and Lois. Some of these are also connected to one another because they tend to be activated together. Note that these are supposed to be features that are shared across different individuals, so one of them (maybe for hair color) is associated with the rectangles for both people.



To use knowledge like this, Hearers would have to have a way to cause the form box to be activated when somebody says the word in their presence. This would then cause the meaning box to be activated and the boxes for the different features to be activated. Speakers would have to have a way of activating the different features associated with the person when thinking about or seeing that person. This would then cause the box for the meaning of the word to be activated, which would cause the box for the form of the word to be activated. There would then have to be a way for this to cause the Speaker to actually produce the word form. Note that I have left almost all of the details unspecified in these processes, and in fact much of what actually goes on when a person produces or understands a word like *Lois* remains controversial. Even the general structure of the knowledge shown in the figure would not be accepted by everybody.

Why Words are Not Inherited

One thing should be clear; however all of this works, it is very complex. So it would be hard to imagine how the relevant knowledge for a whole set of proper nouns could evolve in the Prelings. The alternative is that this knowledge is learned, that each member of the tribe must pick up the set of words from the other living members of the tribe. There is another good argument in favor of learning. The individuals that need to be referred to will change from generation to generation, so unless the *chief* of the tribe and other important members always look the same, there could be no words for these people. Of course simply saying that the knowledge is learned still leaves many questions unanswered. In particular it doesn't tell us where the names, or the idea of naming, came from in the first place. But we'll have to leave these sorts of questions unanswered so that we can move on.

Learning Names

How might the learning of names take place? I've been arguing that names are a relatively simple, basic part of human language, so it shouldn't surprise us that proper nouns are among the first words that babies seem to learn. However, we have to be careful in assigning adult linguistic behavior to infants. When a baby utters "Mama" in the presence of its mother, does this mean that it is referring to its mother in the sense defined above? That is, does the baby want the Hearer to have its mother in mind? Almost certainly not. While it is probably true that Speakers reason about what is in or what could be in the minds of others, everyone who studies babies agrees that babies do not yet have this ability. In any case, for the moment, we won't bother with what it might mean to have a "theory of mind", as this is called. We'll just be concerned with the simpler problem of what it would mean for a baby to figure out how certain sounds or gestures are associated with particular individuals in the world.

Given the picture in the figure above, there would seem to be three parts to this process: learning the form, for example, learning what "Lois" sounds like and how to pronounce it; learning the meaning, for example, learning to identify what Lois looks like and how she behaves; and connecting these two aspects of the word. In fact, if we start with the idea that language is about the

relationship between forms and meanings, we could say the same thing for the learning of all linguistic patterns. Thinking of things this way — in terms of three separate learning processes — will often be useful, but is obviously an oversimplification because the processes can interact with each other in various ways. In particular there is the interesting possibility that people only learn the concept behind the meaning of a word (or other linguistic pattern) as they learn the word; that is, the existence of a single form for a range of different situations is what clues them into the existence of the concept in the first place. But to make sense out of this suggestion, we'll have to go far beyond names and proper nouns and look at the meanings of other types of words. In the next section we'll look at the general category of nouns.

This page titled [2.1: Reference and Proper Nouns](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.2: Categories and Common Nouns

The Limitations of Proper Nouns

Exercise 2.2.1

What would the disadvantages be of a language that contained only names, that is, words that applied to individual things (or situations)?

Names (in the form of proper nouns) buy our Prelings a lot. They can now call the attention of the members of their tribe to some of the important objects in their environment. But they can only do this for the objects that have names assigned to them. For each individual object, being named requires the following:

1. The members of the tribe can all identify the object.
2. The members of the tribe agree on the name for the object, that is, the form of the word.
3. Each member remembers the name so that the form can be recalled when the Preling wants to refer to the individual and the name can be understood when another member of the tribe uses it.

Imagine the following situation. You have been out foraging around and have discovered a previously unknown person, apparently a member of another tribe, hiding near a large rock on the opposite side of the nearby river. Your tribe has an agreed-on name for the river but, though everyone is certainly familiar with the rock, no one has felt the need to refer to it before so it has no name. And because the new person is unknown to everyone in the tribe but you, that person can't have a name. You would desperately like to refer to these two individuals, the rock and the new person, but you can't. You need something more than simple names.

But names have another inadequacy. Recall one of the themes of this book, that language is **constrained** by the bodies and the cognitive capacities of language users. Now consider what it would be like to have a name, a separate proper noun, like *Lois* or *Detroit*, for every individual object you might ever refer to. Even if the members of your tribe could agree on the names, the limits of **long-term memory** would get in the way: how would you ever remember the millions of different names that you'd need?

It was convenient to begin the discussion of words and meaning with names, but for the reasons given above, maybe we shouldn't consider them to be the most basic kind of words.

Categories

The solution to the problems associated with names is provided by another basic human (and Preling) capacity, the capacity to categorize. People group objects in the world into **categories** on the basis of their similarities. Not only can we find individual apples, cats, and mountains in the world and distinguish individual apples, cats, and mountains from one another; we can categorize objects as being apples (and not pears or grapes), cats (and not dogs or goats), or mountains (and not valleys or boulders). An apple is an apple, that is, an **instance** of the category apple, because it shares properties with other apples. Apples tend to have a characteristic shape, a characteristic taste and smell (though there are of course a number of variations), and a characteristic size (though there is again a range of precise possibilities), and they grow on a type of tree with typical properties of its own.


Categories are in the Mind, Not in the World.

Our categories are to some extent *natural*; they correspond to relatively clean divisions in the world. Thus it's not just that we *think* apples are different from pears; they actually are. However, rather than thinking of categories as features of the world, we will be viewing them as cognitive entities, as **concepts**. This is because people clearly have the capacity to come up with categories for things that exist only in their imagination and because different cultures, and different people with the same culture, may categorize the same set of individuals differently. We'll see throughout this book how the similarities and differences between languages can give us insight into how human categorization works.

What good are categories? They permit people to respond in a similar fashion to many different objects once the objects can be categorized. Thus once you know an object is an instance of the category apple, you know that you can eat it or cut it up and put it in a pie. You don't have to learn separately for each individual apple what sorts of behaviors are possible.

Using and Learning Categories

Before we go on and relate categories to language, we need to remind ourselves that we should be considering process as well as product. Rather than think of categories as things in the mind, it will usually be more useful to think in terms of the process of **categorization**, how a person figures out what category something belongs to so that they can then behave in an appropriate way for that thing (eat it, approach it, run from it, etc.). As with individuals, we can think of categories as a localized or distributed in the mind of the categorizer. For now, we'll assume that they are localized, that is, that there is a place in long-term memory dedicated to each category that the Preling or person knows. In such a system, categorization of an object of some sort would normally start with some perceptual input, that is, something seen, heard, felt, smelled, and/or tasted by the categorizer. This would then activate the places in memory dedicated to particular perceptual features that are associated with the object being perceived. If these features overlap enough with the features in a category C and don't overlap more with the features of some other category, then the place in memory for category C is activated. The figure below illustrates this process. Some of the features (small squares) are activated by the perceptual input (activated features appear in green), and these are associated with two categories, apple and pear, but more strongly with apple, so it is activated.



The processing perspective also reminds us that categories like apple are not innate; they have to be learned. So a complete theory of categorization would also have to explain the learning process. But going into that now would lead us too far afield.

Common Nouns

Now let's assume our Prelings have this powerful categorizing capacity. (There is good reason to believe that we're not the only animals that have it, by the way.) All we need to extend their primitive communication system in a powerful way is to associate words with some of the categories. Human languages have such words; we will call them **common nouns**, words like *apple*, *cat*, and *mountain*. Now when a Preling wants to refer to an individual and doesn't have a name for it, there is still an out. If the Preling can categorize the individual and can find a common noun for that category in long-term memory, then they can use that noun to refer to the object. Of course the Hearer now has a new problem, that of figuring out which instance of the category the Speaker is talking about. If the Speaker says "rock", which rock is intended? For both Speaker and Hearer, there is clearly a lot of linguistic work to be done in using common nouns. But think how much more versatile the Prelings' communication system has become. With names they could only refer to the small set of individuals that had words assigned to them. Now they have the potential to refer to an infinite number of individuals, that is, all of the instances of all of the categories they have names for.

Probably because they are so frequent in our speech and because their function is not so complex, common nouns are among the first words that babies learn, at least in many languages. Babies' utterances early in their second year usually consist of single words, and many of these are common nouns such as *juice* or *kitty*. As with proper nouns, we have to be careful in interpreting children's utterances; when these words first appear, it is not necessarily the case that children are using them to refer to objects. They may simply have learned to respond with a particular word form when they are in the presence of particular things.

What Categories Are

Exercise 2.2.2

When we want to find the meaning of a word, we often look it up in a dictionary. But remember that linguists are interested in describing the knowledge and behavior that ordinary native speakers of a language have. For linguists, why would dictionary definitions not work very well as accounts of word meaning?

The Inadequacy of Dictionary Definitions

What's the best way to describe human categories, including those that common nouns are associated with? Or alternately, what form do categories take in the mind? Or alternately, how do nouns mean? This question is one of the hardest and most contentious in the study of language, and there is no agreed-on answer. One possibility would be something like the definition in a dictionary. Here is Webster's definition for *apple*: "the fleshy usually rounded and red or yellow edible pome fruit of a tree (genus *Malus*) of the rose family". This seems to be on the right track, but it includes words like *pome*, which most native speakers (including me) don't know and the relationship between apples and roses, which, if people know it, doesn't seem central to their knowledge of apples. (In any case, a three-year-old, who may have no difficulty understanding and using the word *apple*, wouldn't normally know about this relationship.) Also the definition conveys nothing of what it feels like to bite into an apple, which for many people probably has something to do with what the word suggests to them.

Another problem with dictionary definitions as a model of what we know about words is that they are *in* the language that they are defining. It is hard to imagine how a child who still knows almost nothing about the grammar of the language would be learning meanings in the form of complex expressions in the language that is being learned. Some linguists and philosophers of language have tried to come up with meanings for words that are somewhat like dictionary definitions, expressed in a kind of "language of thought," which all children supposedly already somehow know. But there have been problems with this view since it has been difficult to come up with definitions that hold in all cases. There always seem to be exceptions. Even for something as concrete and seemingly obvious as an apple, nothing seems to be absolutely necessary. Does it have to have a peel and a core? No, a peeled or a cored apple is still an apple. A particular texture? It would strike us as strange, but something that is like an apple in every other way, but has the texture of butter, say, would still be an apple.

Another way of looking at the problem of exceptions is to ask what the necessary and sufficient conditions are for membership in a category. Necessary conditions are those that must be true for all members of the category. Thinking in terms of necessary conditions means taking the perspective of the Hearer. When a Hearer hears a noun like *apple* or *horse*, they would like to be able to know what properties the referent has. If the word is *horse*, can they infer without question that the thing referred to has four legs? Well, no, a three-legged horse is still a horse. It seems that, except in scientific contexts, when nouns need to be used in very precise ways, natural human categories cannot be defined in terms of necessary conditions.

Sufficient conditions are the set of properties that must be true of a thing before we can conclude that it is a member of a category. Thinking in terms of sufficient conditions means taking the perspective of the Speaker. Given an object to be referred to, a Speaker has to categorize it on the basis of the properties it has and then find a word that's associated with the category. It seems just as difficult to come up with a set of sufficient conditions for a category as with a set of necessary conditions. To categorize something as a horse, is it enough to know that it has typical horse features on its head and body? What if it has those features but only because what we're seeing is a horse costume worn by two people? And what exactly are the typical head and body features anyway?

"Good" and "Bad" Apples

One popular idea among linguists and psychologists is that a category takes the form of a prototype, a typical member of the category. For example, for me a prototypical apple is something like a Red Delicious: red, about 6 cm across, and relatively sweet. The prototype may include sensory features (e.g., what it looks and tastes like) and features that have to do with function (what you do with it). On this view, category membership is not an all-or-none matter; an individual is a more or less good member of a category depending on how close it is to the prototype. A Granny Smith is an apple but not as "good" an apple as a Red Delicious for me. And a crabapple is even "worse."

The prototype idea makes sense because objects take more or less time for people to label with nouns or to identify when they hear the nouns. This has been shown, for example, for birds. Americans are faster to call a robin a bird than to call a chicken, an ostrich, or a penguin a bird. I would probably come up with the noun *apple* faster when I'm labeling a Red Delicious than when I'm labeling a Granny Smith. And I would think of or find a Red Delicious faster than a Granny Smith when I hear the word *apple*. Note how the prototype idea also avoids the problem of necessary and sufficient conditions. A three-legged horse may not match the horse prototype perfectly, but it can still be considered a horse, just an atypical one, that is, not a very "good" horse.

There are still other proposals for what categories are. A quite radical one, but one that has been shown to agree with lots of data on categorization, is the **exemplar-based** theory (actually a whole family of related theories). In this approach there is no explicit

representation of categories at all. Each category is just the set of all of the instances of that category that have been remembered (many imperfectly). What is an apple on this view? It is all of the individual apples that remain in the mind of a particular person. And hearing the word *apple* results in some combination of all of these being recalled all at once.

We will be meeting categories of one type or another in every section of this book. In fact to a large extent, the study of language is the study of linguistic categories. As we discuss these categories in more detail, it will sometimes be possible to ignore the details of how categories are represented and how categorization works, but at other times we'll have to consider how a prototype theory or an exemplar-based theory would differ.

Masses, Things, and a Lexicon

Before we go on, let's extend the range of things our Prelings can refer to. Say a Preling has noticed some smoke coming from across the river, a possible sign of another tribe, and wants to report this to the others. Smoke is not an object since it doesn't have well-defined boundaries, although, like objects, it does have characteristic and relatively constant properties, its color, its smell, the way it moves in the wind. We will refer to such things as **masses**. In human languages, masses tend to be referred to by words that are similar to the words used for referring to objects. Some English examples are *fog*, *water*, *clay*, *wood*, and *soup*. Like speakers of English and most other modern languages, our Prelings will use common nouns for masses as they do for objects. We will refer to the meta-category of objects and masses as **things**. Things are characterized by a set of relatively stable properties.

Our Prelings now have the potential for a very large set of words, as many as the categories that they group the things in their world into. Modern humans may know tens of thousands of nouns for the categories they need to refer to. While the Prelings are at this stage very far from having full-blown human language, they do have the beginnings, so it is time to change their name. They now have a **lexicon**, a mental dictionary of words, so we will call them Lexies.


This page titled [2.2: Categories and Common Nouns](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.3: Word Senses and Taxonomies

Word Senses

As cultures develop, they create or learn about new categories of things, for example, tools, and they then have the need to refer to these new things. Where might the words for the new categories come from?


We have seen that words — common nouns — are associated with categories of things. I will refer to those categories that make up the meanings of words as **semantic categories**. As already noted, people also have plenty of categories that have no words associated with them. In fact which categories have labels varies from person to person and from language to language, as we will see soon. One way to summarize what we've discussed so far is shown in the figure below. So far the situation I've described looks like the following:



In the figure the word form is connected to the word's meaning, a semantic category, by an arrow that I will use to indicate a meaning relation. The arrow has two heads because a Speaker can get from a semantic category to a word form, and a Hearer can get from a word form to the word's meaning. You should already know that the figure does not correspond to everyone's view of what word meaning is because what I'm calling the "semantic category" could be distributed in various ways rather than localized as it appears here.

How Words Change

What happens when there is a new category that we need a word for? One possibility is to invent an entirely new word, but this apparently doesn't happen so often. More often the meaning of an existing word is extended to include the new category. When meaning is extended, this is done on the basis of some kind of relationship between the old meaning and the new. I will refer to this as a **conceptual relation**. The general situation involved in extending the meaning of a word, **semantic extension**, is shown in the figure below.



Let's see how semantic extension might work for our Lexies. Say they are familiar with domestic cats and have a word for them; for simplicity, let's say it's pronounced like the English word *cat*. Now they discover tigers and leopards. Each of these new categories gets its own noun, but in becoming familiar with these new animals, the Lexies see their similarities with domestic cats and develop a new category that encompasses all three categories of animals. How do they refer to this new category, that is, to the more general category of cats, what zoologists refer to as members of the family *Felidae*? (For simplicity, I'll call this category general-cat.) One possibility, which is the one used by many English speakers, is to refer to this category using the same word that is used for its most familiar sub-category, that is, *cat*. Note that the word *cat* now has two related meanings. I will refer to related meanings of a single word as **word senses**. For a word that has more than one sense, it is up to the Hearer to figure out when the word is used which of the senses is intended by the Speaker.

Here is another English example of a word with multiple senses. The noun *chicken* can refer both to a particular kind of bird (an object) and to a kind of meat made from this bird (a mass). Notice the two senses in the following sentences.

1. There's a **chicken** walking around in the back yard.
2. There's some **chicken** in the freezer.

The discussion so far might make it seem that language users, or entire language communities, extend the meanings of words consciously, but this rarely happens. Instead there seems to be a natural process by which the meanings of words change over time. As with other kinds of language change, the details of the process are not well understood. Somehow a change that starts with a small number of Speakers has to spread throughout the community and become conventional.

Generalization, Specialization, and Taxonomies

Exercise 2.3.1

A baby uses the word *truck* for cars and buses as well as trucks. How would you describe this error? How is it like a semantic extension? He also uses the word *train* only for the toy train he plays with. How is this kind of error different from the first?


The extension of *cat* to include a new sense seems reasonable because the two senses are closely related. In this subsection, we'll examine this particular conceptual relation in more detail.

Every domestic cat is a member of the cat family (*Felidae*), that is, an instance of the category general-cat. Every tiger and every leopard is also an instance of this category. All of the properties that characterize the category general-cat, in particular, a characteristic body shape and way of moving, characterize domestic cats, tigers, and leopards. I will refer to the conceptual relation between the categories domestic-cat and general-cat as **specialization-generalization**. General-cat generalizes domestic-cat; domestic-cat specializes general-cat.

One Way Knowledge of Objects is Organized in Long-term Memory


There is evidence from cognitive science that much of what we know about objects is organized in terms of the specialization-generalization relation. This is especially true for living things, for which we seem to represent the categories in the world in taxonomies with multiple levels for different degrees of generality. Here is a possible portion of a taxonomy representing knowledge of animals. There are three points to be mentioned about the figure, and about conceptual taxonomies in general.

- As indicated by the small capitals, this is meant to represent a taxonomy of concepts, potential meanings of words, not a taxonomy of words.
- Such a taxonomy is a *psychological* entity; it represents the way some person might organize their knowledge about the world, *not* a scientific account of what is actually in the world. As you may know, zoologists would include several more levels in their scientific taxonomy for animals.
- As before, we need to be careful about taking the boxes too seriously since there are theories of categories that get by without anything so localized.



As we go higher in the figure, the categories become more general; there are fewer and fewer features that characterize their instances. Thus we can say a great deal about the characteristic sounds made by domestic cats, much less about the characteristic sounds made by members of the cat family, even less about the characteristic sounds made by mammals.

Here is another example, representing how someone might organize their knowledge about fruits.




Again there are fewer and fewer characteristic properties of the categories as we go higher in the taxonomy. Apples have a characteristic smell that resembles the smell of pears but is not characteristic of the smell of fruits in general.


Note

If you're familiar with an object-oriented programming language such as Java, Python, or C++, this should all appear familiar to you. Taxonomies like those illustrated in the two figures are inheritance hierarchies, with the nodes in the trees corresponding to the classes of object-oriented programming. This relationship brings up the question of how we should implement knowledge of this sort in a program that's designed to simulate human processing of words or to be a component of a practical system that interacts with users. There are two possibilities. We can use object orientation to directly implement the taxonomy, creating explicit classes for `Food`, `Fruit`, `Berry`, etc., and using the inheritance mechanisms that are built into the language itself to make use of the knowledge that is shared. Or we can treat each category in the taxonomy as an object, an instance of a class like `Concept` and write procedures to implement inheritance between these objects. Most researchers have followed the second path, in part because it gives them control over how inheritance works, how and when new categories are created and associated with the existing taxonomy. But further consideration of these implementation issues is beyond the scope of this book.

What does all of this have to do with words? First, as we have seen from the example of *cat*, words may extend their meanings on the basis of the generalization relation. The noun *cat* came to mean not only domestic-cat, but also general-cat. The figure below illustrates this process.



The opposite is also possible; a word may take on a sense that is more specific than its original sense. Consider the following example from the history of English. In Middle English the word *gerol* meant 'young person'. Over time the meaning of the word shifted to 'young female person', which is the sense it has in its present form, *girl*. This is shown in the figure below, which illustrates the situation after the meaning had changed. The new sense was a specialization of the original sense. Notice that in this case, unlike what happened with *cat* in our first example, the original sense was lost after the word took on the new sense.



Generalization and Specialization in the Borrowing of Words

A similar process may occur when a word is borrowed from one language into another; the sense of the word in the source language may be generalized or specialized after the word enters the target language. The usual Tzeltal word for 'person' is *kirsiano* (the exact pronunciation depending on the dialect), borrowed from the Spanish word *cristiano* meaning 'Christian'. For whatever reason, in the borrowing process the meaning was generalized from 'Christian person' to (any) 'person'. The Swahili word *safari* means 'trip', but when this word was borrowed into English from Swahili, it took on the special sense of a 'trip in search of game animals'. In this case, the meaning was specialized when the word was borrowed.


There are also many examples of generalization and specialization in the speech of young children. In this case we can think of the adult sense of the word as basic and the children's extension of this as a second sense, though it may not be clear that the child has learned the adult sense. Because these uses are seen as errors from the perspective of the adult model, they are referred to as **over-generalization** and **under-generalization**. An example of over-generalization is the use of *dog* to refer to other mammals roughly

the size of dogs such as goats in addition to dogs. Examples of under-generalization are harder to observe because they require noticing that the child *fails* to use a word for a referent where an adult would use the word. An example would be the use of *dog* to refer *only* to a particular dog.

These errors made by children are interesting because they point up the possible similarity between the mechanism of language change and the mechanisms of language learning. What we see in the history of a language, the result of changes across an entire community, possibly over the course of multiple generations, resembles what we see in the development of a single child's linguistic system. We'll see more such examples in the book later on. The similarity implies that we might be able to understand language change as a process of learning taking place in the minds of many different language users, children and adults.

Concepts without Words (Again)

Finally, notice that there may be categories in a person's taxonomy that the person has no words for. Look at how the fruit taxonomy in the figure above corresponds to words for a particular English speaker (me).



While I have a category for things that are either apples or pears, that is, I recognize the similarities between apples and pears, I don't have a word for this category. Note that this does not imply that there is no word in the language for this category; botanists call the family that includes pears and apples *Pomoideae*. But I had never heard this word before I looked it up in order to mention it here, so it is (or was) not a part of my mental lexicon. Because I have words for all or most of the other categories in this taxonomy, we can consider the apple/pear category to constitute a **lexical gap** for me. We will see other examples of lexical gaps later in this chapter.

This page titled [2.3: Word Senses and Taxonomies](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.4: Metaphor and Metonymy


Metaphor

Exercise 2.4.1

Consider the use of the word *web* for the World-Wide Web and the use of the word *root* to mean 'source' (as in *the root of all evil*). What is the basis for these semantic extensions? What do they have in common?

So far we've looked at how the meanings of words can be extended, both by adult speakers and by babies learning the language, in ways that make them more or less general. In this section we'll consider two other general kinds of conceptual relations that permit word meanings to be extended: similarity and various kinds of close association.

First consider the situation that arose when computers were first outfitted with pointing devices to be manipulated in one hand by moving them across a pad and pushing one of their buttons. The noun that came to be used for these devices, *mouse*, was based on the resemblance of the devices to the animal: the general size and shape and the tail-like cable. Thus the meaning of the word *mouse* was extended on the basis of the physical similarity between one category (the animal) and another (the pointing device). Extension of a word's meaning on the basis of similarity is known as **metaphoric extension**. This figure illustrates the process.



Here is another computer-related example. When computer software started providing users with sets of choices that they could select from, the word that emerged for these lists was *menu*. The metaphoric extension was based on the similarity between a restaurant menu, a list of food and drink choices, and the list of choices that the computer user was to select from.

Metaphoric Extension from Concrete to Abstract

One frequent use of metaphor is the application of a word referring to an object category to a more abstract semantic category, something not physical at all. Consider the structure of taxonomies as in this example from the last section. If we turn the figure illustrating the taxonomy over, it resembles a tree, with the most general category as the root and the most specific categories the leaves. This is in fact how cognitive scientists refer to structures like this; *tree* is applied to the whole structure, *root* is applied to the point where all of the branches begin, and *leaf* is applied to the point beyond which there are no more branches. Note that a taxonomy is not a physical thing at all, so with metaphoric extension we have now taken common nouns such as *tree* and *leaf* outside of the realm of the physical entirely.

This example also illustrates how metaphor often operates on two entire domains, each with its own elements and internal structure. The **source domain** is the one that is being used to understand the (usually more complex) **target domain**. In this example, the source domain is tree, the target domain taxonomy. The metaphor is based on multiple similarities between the domains: correspondences (or **mappings**) between the elements (leaves and specific concepts, for example) and the relations between the elements (branches join the root to the leaves; generalization links join the most general concept to the more specific ones).

Metonymy


The word for 'language' may be related to other, less abstract words, for example, the word for 'tongue' (as in Spanish), the word for 'mouth' (as in Oromo), the word for 'voice' (as in Tzeltal). Assuming that the 'language' sense is an extension of the more basic sense in each case, what's the basis of this extension?

A somewhat more complicated possibility for extending a word meaning is based on a quite different conceptual relation, not similarity between the instances of the two categories but a strong association between them. This is referred to as metonymic

extension. Consider the association between an organization (an abstract concept), such as a sports team or a government, and its base location. While we can refer to the organization directly using its name, we often find it convenient to use the name of the location to refer to the organization.

1. *Dallas* won yesterday's game.
2. No one is sure what *Moscow's* response will be.

This figure illustrates the first example of metonymy.



Another conceptual relation that permits metonymy is that between a document and the content of the document. Thus the word *book* refers to a physical object: a collection of sheets with printing or pictures on them that is bound together. But we can also use the word to refer to the informational content of the physical book. Compare the uses of the word in these two sentences.

3. This **book** is almost too heavy to lift.
4. I don't understand this **book** at all.

In the first example, the speaker is clearly referring to the physical object, in the second example to the information contained in the physical object. In a case like this, metonymic extension allows a noun referring to a physical object to refer to something more abstract.

Insulting with Metonymy

Metonymy may also be used in situations where an alternative to an existing noun is called for, perhaps as a very informal or insulting term. Examples are the use of *wheels* to mean 'car', *brain* to mean 'intelligent person', and *asshole* to mean 'person' in an insulting context. In these examples the relevant conceptual relation is between a whole and a part. (There is much more going on than just this, especially in the last example, because the choice of the particular part is obviously also relevant!)

Metonymy may also come into play "on the fly", when a speaker is using language creatively. Here's an example from Fauconnier (1985). One waitress in a restaurant is speaking to another.

5. The **omelet** left without paying.

Of course the speaker doesn't mean that some cooked eggs left the restaurant; she is referring to a customer. The conceptual relation that is the basis for the metonymic extension in this case is the relation between a customer and the customer's order. Note that *omelet* would only be expected to get this interpretation in the appropriate context; that same person would not normally be called *the omelet*.

Metonymy in Children's Nouns?

Finally we see apparent examples of metonymy in the speech of young children. Andrés, who was learning both English and Spanish, used the Spanish word *luna* ('moon' in adult Spanish) during his second year to refer not only to the moon and crescent shapes but also to the pens or pencils used to *draw* crescent shapes. It appears that he has extended the word on the basis of the relation between an image and the instrument used to produce the image. But, as always, we must be careful in interpreting children's utterances. Most of Andrés' utterances during this period consisted of a single word. When he said "luna" apparently referring to a pen, did he really mean something more like "use this to draw a crescent"? We have no simple way of knowing.

This page titled [2.4: Metaphor and Metonymy](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.5: Deixis and Person


Utterance Contexts

Exercise 2.5.1

Consider the sentence *I love you*. What changes in the interpretation of the words *I* and *you* when different speakers say this to different hearers?

Each time a Speaker **refers**, there is a word (or more) that is uttered, that is, a form, and there is a **referent**, the thing that is being referred to. (Of course the referent may be something imaginary, but we can still talk about it existing in the mind of the Speaker.) But for each reference, there are always in addition several other things: the Speaker, one or more Hearers, the time and place of the reference, a set of things that the Speaker and the Hearer are currently aware of, and possibly some other language that has just been produced by the Speaker, the Hearer, or somebody else. A particular instance of language is an **utterance**. An utterance needs to be distinguished from a particular word, phrase, or sentence because an utterance has an **utterance context**, a particular Speaker, Hearer, time, place, available things, and recent language, in addition to its own linguistic form. I'll sometimes refer to the Speaker and Hearer as **utterance participants**. For example, we can put the English words *I*, *like*, and *it* together to make the English sentence *I like it*, but this sentence is a different utterance each time it is uttered.

Each utterance has its own context, and, as we will see below, for each context the sentence has a different meaning. That is, meaning always changes from one utterance context to another. The figure below is one way of representing the elements of an utterance context. Each of the elements is a **role** in the context, a kind of slot that gets filled by something in each different situation. For example, the Speaker role is filled by a particular person, and the Location role is filled by a particular place. We will meet the concept of role again later in this book; in fact it is one of the most fundamental notions in cognitive science.



Deixis


Exercise 2.5.2

How does the notion of utterance context help us understand how words like *I* and *you* refer?


Now we'll examine how our Lexies can use the idea of utterance context to come up with a new way to refer to some of things around them. Each person in the tribe has a name, and when referring to members of the tribe, a Speaker can use the name of the referent (a proper noun such as *Philip*). Proper nouns have an unusual property; unless there is more than one individual with the same name, proper nouns always refer to the same individual in the world, no matter when, where, and by whom they are uttered. (Note that this is not true of common nouns; we can use the noun *apple* at different times to refer to any number of different individual apples in the world.)

Referring to Someone in Terms of their Utterance Role

Now consider an approach to reference that is in a sense the opposite of that taken with proper nouns, which are completely independent of the utterance context, one in which the meaning depends *completely* on the utterance context. Say Clark is speaking to Lois and wants to refer to Lois. He could use her name, a proper noun. Or he could use a word which refers to the Hearer, *whoever that might be*. This is how the English word *you* works. Without knowing the utterance context, or at least knowing who the Hearer is, we have no idea what the referent of *you* is. The figure below illustrates the meaning of *you*. The referent of a *you* utterance is joined by the meaning arrow to the *role* of Hearer rather than to a particular individual (as for proper nouns) or a category (as for common nouns).



For each utterance of *you*, the word gets its interpretation from the utterance context, that is, who it is that fills the role of Hearer. So in our example, with Clark as the Speaker and Lois as the Hearer, the situation is as shown in the figure below. There are question marks in the Location and Time roles because we don't know what these are, and the number at the end of "utterance" is meant to indicate that this is a particular utterance, not the general prototype for utterances shown in the previous two figures.



You is an example of a deictic expression, an expression that gets its meaning directly from the utterance context, that makes reference to one or more of the roles in the utterance context: the Speaker, the Hearer, the location, or the time. The noun form of the word is **deixis**.

You may already have figured out where we're going next. Just as we have an English word to refer to whoever the Hearer is, we have a word to refer to whoever the Speaker is: *I* (or *me*). *I* gets its meaning from the utterance context just as *you* does. *You* and *I* are examples of personal pronouns words which refer directly to participants in the utterance context. All languages apparently have personal pronouns, a quite striking universal property, though languages differ greatly in the details, as we will see later on.

Person

Exercise 2.5.3

What information is conveyed by the word *her* in the sentence *I love her*? Does (or could) the meaning of the word change with the utterance context as it does for *I*?

The personal pronouns *I* and *you* are alike in a number of ways; they differ with respect to their person, that is, which utterance participant they refer to. Conventionally we call reference to the Speaker **first person** and reference to the Hearer **second person**. Note that reference that includes the Speaker as well as other people is considered first person; thus *we* is a first person pronoun. The commonality between *I* and *we* is not something that is reflected in form in English, but we should not be surprised to find it elsewhere. In Japanese, there are many words for *I*, but each of these words has a corresponding form ending in *-tachi* that means 'we': for example, *watashi*, *boku*, *ore* 'I'; *watashitachi*, *bokutachi*, *oretachi* 'we'.

Reference that does not include the Speaker but includes at least the Hearer is considered second person; thus the Southern English pronoun *y'all* and the Spanish pronoun *ustedes* are second person. Notice that many English dialects use the same pronoun, *you*, for both reference to a single Hearer and to multiple Hearers or to the Hearer and other people. First and second person seem to be universal categories in languages; they appear not only in the form of pronouns, as we'll see in Chapter 7.

What about reference to things (or abstractions) that are (or include) neither the Speaker nor the Hearer? Here things get a little more complicated. Clearly any reference which is not first person or second person belongs to this category, which is known as **third person**. So in the following sentence, the expressions in bold are third person.

1. Clark told **his sister** about **the movie**.

This sentence makes reference to three different things, and none of these either is or includes the Speaker or Hearer. None of these references seems to be deictic either since their meanings do not seem to depend on the utterance context, on who says the sentence to whom and on when and where it is said. (This is not quite true — the interpretation of *the movie* in this sentence does depend on the utterance context — but we will not worry about this aspect of deixis.)

Just as we have first and second person pronouns which say nothing more about the referent than that it is (or includes) the Speaker or Hearer, we have third person pronouns which say little more about the referent than that it does *not* include the Speaker or

Hearer. Third person pronouns in English include *she*, *he*, *it*, and *that*. Notice that these pronouns do provide a little more information about the referent than that it is third person (neither Speaker nor Hearer). They also tell us something about the *gender* of the referent. *He* refers to something that is not the Speaker or Hearer and is perceived as male. (Note that the referent doesn't have to be male in a biological sense. In informal English we often use *he* in informal English to refer to animals without actually checking first to make sure we have their gender right. Another way to look at this usage is to think of the male pronouns as the **default**, the form we use for animals when we don't know the gender.) Similarly, *she* refers to something that is not the Speaker or Hearer and is perceived as female, and *it* and *that* refer to something which is perceived as non-human.

Pronouns as Speaker-oriented Shortcuts

Note how third person pronouns act as shortcuts; in this sense they are Speaker-oriented. To refer to something, a Speaker can just say "it" and not bother coming up with the name of the thing or a common noun for its category. Of course the burden is then on the Hearer to figure out which non-human thing the Speaker is referring to. This implies two things about language learning. First, people have to learn how to interpret such pronouns. This is not trivial, and it has proven to be one of the most difficult language behaviors to get computers to do. Second, people have to learn in what situations it is appropriate to use such pronouns. Consider the following sentence uttered at the beginning of a conversation.

2. *Did you find it?*

This sentence sounds silly unless the Speaker somehow knows that whatever *it* refers to is on the Hearer's mind and that nothing else the Hearer might have been looking for is.

Let's summarize what we've given our Lexies in this section. Personal pronouns don't actually allow them to refer to any new things in the world that they couldn't already refer to; they already had proper nouns or common nouns for this purpose. Instead personal pronouns give them a new way to refer, using the roles of the utterance context directly. As we have seen, utterance contexts have more than just a Speaker and a Hearer, and we can expect languages to have deictic words that refer to the other roles as well. Words such as *here* and *now* do exactly that.

Learning Deixis

Exercise 2.5.4

If young children learning English treat the word *you* like a proper noun instead of a person pronoun, what kinds of mistakes will they make?

Pronouns May be Easy to Produce, But in the Beginning They're Hard to Figure Out.

As we have seen, young children must learn to refer using words that point directly to individuals (proper nouns), words that point to categories of things (common nouns), and words that point to deictic roles. Deixis seems to be the last of these to emerge. Early on children often treat first and second person pronouns as though they were proper nouns. So when the Speaker uses the word *you* to refer to the Hearer (the child), the child may also use the word *you* to refer to herself, who is now the Speaker. Perhaps deixis is hard for a child because it requires switching perspective. The child can't simply imitate the adult usage because the roles of Speaker and Hearer switch when this happens, and deictic words like *I* and *you* change their referents. In some sense the child has to understand a word like *you* from the perspective of the Speaker, realizing that that person then fills the *you* role when she becomes the Speaker.

Third person pronouns are difficult for children in a related way. As we have seen, their appropriate use depends on the ability to know whether the Hearer can interpret them, that is, in some sense on the ability to put oneself in the position of the Hearer. Because this is apparently difficult for children, they often produce sentences like the following in a situation where the Hearer would have no way of figuring out who *he* is.

3. *He hit me.*

This page titled [2.5: Deixis and Person](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.6: Lexical Differences Among Languages

Some Reasons Languages Differ Lexically

So far we have endowed our Lexies with an amazing capacity, one that to date has only been found among human beings. Over the generations, they can now invent a very large store of labels for individuals and categories of things in the world (even categories of things not in the world). And, equally important, they can pass on this store of labels to their children.

Now let's imagine various tribes of Lexies in different parts of the world with no contact with each other. Each tribe will experience a different environment, containing its own potentially unique set of animals and plants and its own climate and geology. Each tribe will invent words for the things in its environment that matter to it, and we will naturally expect to find words for different things in each tribe. Modern languages also differ from each other in this way. Amharic has a word for hippopotamus because hippopotamuses are found in Ethiopia, but Inuktitut does not because hippopotamuses are not found (normally) in northern Canada.

We can also expect the cultures of the different tribes of Lexies to differ. This will result in several differences in their store of words. First, certain naturally occurring things will become more important. A tribe that makes pots out of clay will want a word for clay; another tribe may not bother. Second, as culture develops, there will be more and more cultural artifacts, that is, objects produced by the members of the culture. Naturally the tribe will want words for these as well, and if they are not producing them, they will not have such words. Finally, culture results in abstractions, concepts that do not represent (physical) things in the world at all: political units, social relationships, rituals, laws, and unseen forces. These will vary a great deal in their details from tribe to tribe, and we can expect these differences to be reflected in the words that each tribe comes up with.

Culture and Nouns

Modern languages also differ from each other in these ways. Amharic has the word *agelgil* meaning a leather-covered basket that Ethiopians used traditionally to carry prepared food when they traveled. Other languages don't have a word for this concept. English now has the word *nerd* to refer to a particular kind of person who is fascinated with technology and lacking in social skills. This is a relatively new concept, specific to certain cultures, and there is probably no word for it in most languages.

Differences Within and Among Languages

Exercise 2.6.1

Languages such as English, Spanish, Mandarin Chinese, and Japanese have many specialized terms for computers and their use, whereas many other languages, such as Tzeltal and Inuktitut, do not. Does this represent some kind of fundamental limitation of these languages?

Finally we can also expect the store of words to vary among the individuals *within* each tribe. As culture progresses, experts emerge, people who specialize in agriculture or pottery or music or religion. Each of these groups will invent words that are not known to everyone in the tribe. Modern languages also have this property. A carpenter knows what a hasp is; I have no idea. I know what a morpheme is because I'm a linguist, but I don't expect most English speakers to know this.

This brings up an important distinction, that between the words that a language has and the words that an individual speaker of the language knows. Because some speakers of languages such as Mandarin Chinese, English, Spanish, and Japanese have traveled all over the world and studied the physical environments as well as the cultures they have found, these languages have words for concepts such as hippopotamus and polygamy, concepts that are not part of the everyday life of speakers of these languages. Thus it is almost certainly true that Mandarin Chinese, English, Spanish, and Japanese have more words than Amharic, Tzeltal, Lingala, and Inuktitut. But this fact is of little interest to linguists and other language scientists, who, if you remember, are concerned with what individual people know about their language (and sometimes other languages) and how they use this knowledge. There is no evidence that individual speakers of English or Japanese know any more words than individual speakers of Amharic or Tzeltal.

Where New Words Come From

Furthermore, if a language is lacking a word for a particular concept, it is a simple matter for the speakers of the language to add a new word when they become familiar with the concept. One way for this to happen is through semantic extension of an existing

word; we saw this earlier with *mouse* in English. Another way is to create a new word out of combinations of old words or pieces of old words; we will see how this works in Chapter 5 and Chapter 8. A third, very common, way is to simply borrow the word from another language. Thus English speakers borrowed the word *algebra* from Arabic; Japanese speakers borrowed their word for 'bread', *pan*, from Portuguese; Amharic speakers borrowed their word for 'automobile', *mekina*, from Italian; and Lingala speakers borrowed their word for 'chair', *kiti*, from Swahili.

Lexical Domains: Personal Pronouns

What are the differences between the personal pronouns *you* and *you guys*? (There are at least two differences.)

More interesting than isolated differences in the words that are available in different languages is how the concepts within a particular *domain* are conveyed in different languages. We'll consider two examples here, personal pronouns and nouns for kinship relations; we'll look at others later on when we discuss words for relations.

A complete set of personal pronouns in my dialect of English includes the following: *I, me, you, she, her, he, him, it, we, us, you guys, they, them*. Note that I'm writing *you guys* as two words, but in most important ways it behaves like one word. For our present purposes, we can ignore the following group: *me, her, him, us, them*; we're not really ready to discuss how they differ from the others. Among the ones that are left, let's consider how they differ from each other. We have already seen how they differ with respect to person: *I* and *we* are first person; *you* and *you guys* are second person; *she, he, it, and they* are third person. We can view person as a **dimension**, a kind of scale along which concepts can vary. Each concept that varies along the dimension has a **value** for that dimension. The person dimension has only three possible values, first, second, and third, and each personal pronoun has one of these values.

Person is not just a conceptual dimension; it is a **semantic dimension** because the different values are reflected in different linguistic forms. That is, like words, semantic dimensions have both form and meaning. When we speak of "person", we may be talking about form, for example, the difference between the word forms *I* and *you*, about meaning, for example, the difference between Speaker and Hearer, or about the association between form and meaning.

But person alone is not enough to account for all of the differences among the pronouns. It does not distinguish *I* from *we*, for example. These two words differ on another semantic dimension, **number**. *I* is **singular**: it refers to an individual. *We* is **plural**: it refers to more than one individual. What values are possible on the number dimension? Of course languages have words for all of the different numbers, but within the personal pronouns, there seem to be only the following possibilities: singular, dual (two individuals), trial (three individuals), and plural (unspecified multiple individuals). Of these trial is very rare, and, among our set of nine languages, dual is used only in Inuktitut. Thus Inuktitut has three first person pronouns, *uvanga* 'I', *uvaguk* 'we (two people)', *uvagut* 'we (more than two people)'.

Given the two dimensions of person and number, we can divide up the English personal pronouns as shown in the table below. The third person pronouns fall into the singular group of three, *she, he, and it*, and the single plural pronoun *they*. The second person is more complicated. In relatively formal speech and writing, we use *you* for both singular and plural, but informally, at least in my dialect, we may also use *you guys* for the plural. (Note that other English dialects have other second person plural pronouns, *you all/y'all, yunz*, etc.) Thus we need to include both *you* and *you guys* in the plural column.

	singular	plural
1 person	I	we
2 person	you	you, guys
3 person	she, he, it	they

Clearly we need more dimensions to distinguish the words since two of the cells in our table contain more than one word. Among the third person singular pronouns, the remaining difference has to do with **gender**, whether the referent is being viewed as male, female, or neither. Instead of *male* and *female*, I will use the conventional linguistic terms **masculine** and **feminine** to emphasize that we are dealing with linguistic categories rather than biological categories in the world, and for the third value I will use **neuter**. Thus there are three possible values on the gender dimension for English, and three seems to be all that is needed for other languages, though some languages have a dimension similar to gender that has many more values.

That leaves the distinction between *you* and *you guys* in the plural. As we have already seen, this is related to **formality**, another semantic dimension and a very complicated one. I will have little to say about it here, except that it is related to the larger context (not just the utterance context) and to the relationship between the Speaker and Hearer. For example, language is likely to be relatively formal in the context of a public speech or when people talk to their employers. For now, let's assume that the formality dimension has only two values, informal and formal. The table below shows the breakdown of the English personal pronouns along the four dimensions of person, number, gender, and formality.

	singular			plural	
1 person	I			we	
2 person	you			formal	informal
				you	you guys
3 person	feminine	masculine	neutral	they	
	she	he	it		

Gaps in Pronoun Systems

Notice that there seem to be gaps in the English system. There is a word for third person singular feminine, but no word for second person singular feminine, and formality is only relevant for second person plural. Because there is no masculine or feminine *you* in English, we can say that *you* is unspecified for the gender dimension. As we will see many times in the book, languages tend to be **systematic** — if they make a distinction somewhere, they tend to make that distinction elsewhere — but they are not always so. English personal pronouns *are* systematic in one important way: the distinction between first, second, and third person is maintained in both singular and plural. But they are not in other ways, as we have just seen.

You will probably not be surprised to learn that there is nothing special about the English system; other languages organize things somewhat differently, though it seems that person and number are relevant for all languages. Here is the set of Amharic personal pronouns.

singular				plural
1 person	iné			innya
2 person	plain		respect	innante
	feminine	masculine	irswō	
	anči	ante		
3 person	plainplain		respect	innessu
	feminine	masculine	issaccew	
	isswa	issu		

Notice that Amharic fills some of the apparent gaps that English has; for example, there is both a masculine and a feminine second person singular pronoun, while English only makes the gender distinction in third person. But Amharic is unsystematic in some ways too; while gender is relevant for singular pronouns, it is not for plural pronouns, and, as in English, it doesn't enter into first person at all. Notice also that there is a new dimension, **respect**, that is relevant for Amharic pronouns, at least in second and third person singular. Respect is similar to formality, but it relates specifically to the attitude that the Speaker wants to convey toward the referent, that is, the Hearer in the case of second person and another person in the case of third person. In Amharic, there are two values for this dimension, plain and respectful. Finally, notice that while English has three values for gender, Amharic has only two, masculine and feminine. This means that one or the other of these must make do to refer to things that are neither male nor female. Many languages have only two genders, and each of these languages has its own way of determining which gender is appropriate for things that don't have "natural" gender.

We have seen only two examples of personal pronoun systems. Other languages have quite different systems, some making use of dimensions that are not relevant for English or Amharic, some ignoring dimensions that matter for English and Amharic. For example, in many languages, including Tzeltal and Inuktitut, gender plays no role at all in the personal pronoun systems: there is no distinction like that between *he* and *she*. It is not clear why pronoun systems vary the way they do. For example, it would be wrong to assume that Tzeltal pronouns lack gender because Tzeltal speakers are less conscious of gender in the world or that children learning Tzeltal become less sensitive to gender differences than children learning English or Amharic or Spanish. At least there is no evidence for these kinds of relationships. The relationship between language and thought has been most often studied in the context of grammar, and since we are looking at personal pronouns, we are getting pretty close to grammar, but we will save this topic for later.

Lexical Domains: Kinship Terms

What do the meanings of the words *father* and *uncle* have in common? What sort of dimension would you need to distinguish the meanings of these words?

Now let's look at the words we use to refer to kinship relations. We won't consider all of the words in a given language, just some of the basic ones. Let's start by taking two similar words and trying to figure out what dimension distinguishes their meanings, say *brother* and *sister*. This is easy since we've already been discussing this dimension; it's gender.

But gender won't help us with the distinction between *daughter* and *mother* since both are female. For these words we have to consider their relationship to the person who provides the reference point for the relationship, what cultural anthropologists (the experts on this topic) call **ego**. In both cases, there is a direct relationship (what anthropologists call **lineal**), but in one case the relationship goes in one direction (back into the past); in the other, it goes in the opposite direction (forward into the future). Let's call this dimension "vertical separation from ego". We can use positive and negative numbers to represent values on this dimension. In the case of *mother*, the separation is -1 (one generation back); in the case of *daughter*, it is +1 (one generation forward).

But these two dimensions won't suffice to distinguish all basic English kinship terms. What about *mother* and *aunt*? Both are female, and both are separated by -1 from ego. What distinguishes these two relations is the closeness of the relationship to ego. For *mother*, the person is in a lineal relation to ego. For *aunt*, we need to go back another generation, to ego's grandparents, to find a common ancestor. We will call this dimension "horizontal distance from ego" and represent it again with a number (but no sign). For *mother*, we will say the distance is 0; for *aunt* (and *cousin* and *niece*), it is 1. Here is a list of some English kinship terms with their values on the three dimensions. If a cell is left blank, the dimension is unspecified for that term.

	vertical	horizontal	gender
mother	-1	0	feminine
daughter	+1	0	feminine
sister	0	1	feminine
aunt	-1	1	feminine
parent	-1	0	
grandchild	+2	0	
niece	+1	1	feminine
cousin	0	2	

Not All Languages have "Aunts" and "Uncles"

Now let's look at some of the terms that Lingala speakers use for kinship terms. Some of these are just like English, but others require different dimensions than are required for English. Lingala speakers use different words for siblings that are older or younger than ego and for aunts and uncles that are older or younger than their parents, but they don't normally distinguish siblings or aunts and uncles by gender. We'll refer to this as the "relative age" dimension. Lingala speakers also distinguish maternal and paternal aunts and uncles; we'll call this the "parent path" dimension. Finally, Lingala speakers use the same words for grandparents

and grandchildren; that is, at least some of the time they are concerned only with vertical distance, not vertical direction (earlier or later). The table below shows values on the kinship dimensions for some Lingala kinship terms.

	vertical	horizontal	gender	parent path	age
<i>mama</i> 'mother'	-1	0	feminine	maternal	
<i>tata</i> 'father'	-1	0	masculine	paternal	
<i>nkoko</i> 'grandparent/ grandchild'	2 (+/-)	0			
<i>nkulutu</i> 'older sibling'	0	1			older
<i>leki</i> 'younger sibling'	0	1			younger
<i>mama-nkulutu</i> 'older sibling of mother'	-1	1		maternal	older
<i>tata-leki</i> 'younger sibling of father'	-1	1		paternal	younger

Differences in kinship terms are more likely to be related to culture than differences in personal pronouns. That is, when a single term (such as Lingala *nkulutu* 'older sibling') groups different relatives together, we might expect that in the culture where the language is spoken, those relatives are treated similarly by ego. (I don't know whether this is the case for Lingala speakers, however.) Words refer to categories, after all, and categories are a way in which people group the things in the world. Children growing up in a particular culture are learning the cultural concepts and the words simultaneously. Their experience with the culture should help them learn the words referring to cultural concepts, and their exposure to the words should help them learn the concepts. But little is actually known about how this sort of interaction works. In the next section we'll consider the learning of the meanings of apparently simpler nouns, those referring to physical objects. Even here we'll discover that there is considerable disagreement on how babies manage to master the words.

This page titled [2.6: Lexical Differences Among Languages](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.7: Learning Meaning

Why Language Learning Seems Hard

Exercise 2.7.1

Say a child is presented with a single example of an apple along with the word *apple*. What would it take for the child to be able to correctly apply the word in the future to other apples and not incorrectly apply it to, say, pears or strawberries or mushrooms?

One of the striking, and very powerful, features of human language is that it is learned. Within the domain of words alone, the fact that language is learned allows people to continually come up with new words for new concepts (by extension or combination of existing words, by borrowing, or by inventing). When words are combined to form sentences, even greater flexibility is possible, as we'll say later on in the book.

In and of itself, learning is not so impressive; after all, some birds learn their songs. Language learning seems amazing for two reasons.

1. What is learned, a human language, is very complex.
2. The information that children are provided about the language seems insufficient.

We'll be examining the complexity of language throughout this book. For now, I'll mention just one aspect of language you should already be familiar with. Language is applicable to a potentially infinite range of situations. That is, the learner has to **generalize**, to go beyond the examples that are seen during learning and apply language to *novel* situations. Once you know a word like *rock*, for example, you have the ability to use it to refer to any rock you might encounter and to understand it when someone else does this. Conversely, you know when *not* to use the word *rock*; you would not apply it to some sand, for example. And when you hear someone else using the word and there are several possible referents around, you know that the speaker is not referring to a tree or a puddle. The problem is that when you're learning the meaning of the word, you can't possibly be shown all rocks and told that they belong to the semantic category rock and shown everything else and told that it doesn't belong to the category.

Learning What Is and What Isn't a Rock

To see why this would matter, let's consider a very simple "language", consisting of just one word, *rock*. As a language learner, you have to figure out the meaning of each word. (You also have to figure out how to pronounce each word, but I'll save that kind of complexity for the next chapter.) Let's simplify further by assuming that the information you receive about the language consists of presentations of pairs of objects and words with no distractions of any kind (other objects that could be possible referents for the word, other words that could refer to the object). After lots of these presentations, you should begin to have an idea of what a rock is; that is, you would know how to use the word *rock* when you see an object that is very similar to the rocks you've been presented. But what about a potential referent that is *not* so similar, for example, a rock that is much larger than the ones you've seen so far, or a clump of soil? How would you know to apply *rock* to the first of these and not to the second? In particular, without ever being told that a clump of soil or some sand is not a rock, what would prevent you from using the word to refer to these?

So it seems that to learn word meanings, you'd need two kinds of examples, examples illustrating where particular words apply (*rock* in the presence of a rock) and examples illustrating where particular words do not apply (**not** *rock* in the presence of a clump of soil). The first kind of information is called **positive evidence**, the second kind **negative evidence**. Even with all of this help, you might often find yourself stuck if you were faced with an object quite unlike any of the ones you'd seen (though to be fair people sometimes do have trouble deciding which word applies in a given situation).

Distractions: Other Objects and Other Words in the Context

So the learning task itself looks very challenging. What about the information that learners actually get about the words in the target language? First, it obviously does not consist of simple pairings of objects and words. Objects do not present themselves in isolation from the rest of the world. For a given presentation of a noun, there will usually be other objects and masses around in addition to the intended referent of the noun. Furthermore, as we'll see beginning in Chapter 5, words don't just refer to things; they also refer to properties of things or relations between things. Say a child has just heard the word *tiger*. Even if there is a tiger

present, and the child is able to pick out the tiger in the scene, how does she know the word refers to that animal and not, say, the trees around the tiger, the ground beneath the tiger, the sky overhead, the tiger's legs, the relation between the tiger and the ground, the relation between the tiger's legs and the tiger's body, or any number of other noticeable aspects of the current situation? This problem, described in 1960 by the philosopher of language W. V. O. Quine in the context of translation rather than language learning, is sometimes called **Quine's problem**.

Furthermore, nouns usually do not occur in isolation; they occur together with other words. The child has to segment the noun out from the stream of sounds and figure out which aspects of the situation the different words refer to. If the child already knows all of the other words, this may not be too difficult, but the utterance may contain more than one unfamiliar word.

Finally, although we've seen that negative evidence seems crucial for learning, adults apparently do not provide children with negative evidence, at least not directly. That is, they do not say things like "this is *not* a rock" very often. Of course they may correct the mistakes that children make ("no, that's not a rock; it's dirt"), but children often seem to ignore the corrections, especially if they have made themselves understood in spite of their error.

So How is Language Learned?

These arguments about the complexity of language and the seeming lack of information available to children have led many researchers to look for constraints on how language is learned. The idea is that if the child somehow "knew" that only certain things were possible in human language, this would make learning simpler because the set of possibilities would be smaller. For example, if children knew from the outset that all languages have a category of words (common nouns) that are used to refer to categories of things in the world, then they could focus the learning process by "looking for" nouns in what they hear and not making misguided hypotheses about what the words are for. Words *could* work a very different way from the way they work in human language. For example, instead of referring to all instances of a category, a noun could refer (in a sort of metonymic way) to things that go together. So hearing *cat* for one cat, a learner could assume that the word is used for that cat and the things that go with it, its food, its owner, the place where it sleeps. Of course nouns *don't* work this way, and the point of the constraint would be to prevent learners from even considering that they do. We'll see other examples of possible constraints on language learning later on.

If we can agree on what some of these constraints are, how would the child get them? One common kind of proposal is that the constraints are **innate**, that is, that they are in the child's genes and not learned at all. The learner "knows" the constraints on what is possible in human language just because the learner is a human being. Another possibility is that the constraints come through experience in the first few months of life and are already in place when the relevant aspect of language is learned. However, this could only apply to constraints that could be learned by an infant through experience with the world. Many people don't distinguish these two positions because they both make the strong claim that constraints on how language is learned are already in place when the process begins in earnest. For the learning of words and their meanings, this means that the constraints are available by the child is about nine months old.

Innate Constraints vs. Statistical Learning

This view of language learning as constrained by universal principles that are either innate or learned very early on is quite popular. But it is not the only possibility. An alternative position gives a central role to the child's experience with language and the world. I'll refer to this as the "empiricist position". The empiricists believe that the linguistic information children receive is richer than it appears at first glance, that it contains many sorts of **regularities**, that is, recurring patterns. In addition to the regularities that are the conventions of language, how words are pronounced and combined in meaningful ways, there are regularities in the way people *present* language to children. For example, people tend to look at an object they are referring to if it is present, and this can help a child figure out what the word refers to.

This account of language learning requires that children be good **statistical learners**, that they be very sensitive to the regularities in the input around them, and empiricists have made the case this is so. If children can learn this way, then they can not only pick up on the regularities that are present in the environment but can also compensate for the lack of negative evidence by noticing what tends *not* to occur with what as well as what tends to occur with what and by learning about the consequences of their own mistakes. Finally, empiricists argue that children are predisposed (probably innately) to be social creatures, to notice what other people are attending to and to be interested in what they are doing. Being interested in certain things and not in others constrains the space of things that the learner will attend to or will guess that language is about. For all of these reasons, the empiricists hold

that people do not start out with innate knowledge of language but rather pick up what they know about language (or at least most of what they know) through experience. Because experience is so important in their view, some empiricists also argue that children learning different sorts of languages may behave somewhat differently, at least early on.

Many linguistic textbooks dismiss the empiricist position, giving the impression that the issue is solved, that language (though obviously not particular languages) is basically innate. But this is not at all the case; if anything, the controversy is more heated now than it was ten or twenty years ago. As someone squarely on the empiricist side, I will try to show how this position makes sense, here and in other chapters, but I will *not* be claiming that we have all the answers. The question of how language is learned is still one of the great outstanding questions facing science.

The Shape Bias

Exercise 2.7.2

How would the knowledge that nouns tend to refer to categories defined partly by shape help the child learning the meaning of the word *apple*?

In this subsection, we'll look at one example of a possible constraint on learning that seems to be reflected in the behavior of learners, and I'll suggest an alternative explanation for the behavior that doesn't require a specific constraint.

Many of the early words of children learning English and many other languages are nouns, both proper nouns (*Daddy*) and common nouns for concrete things, especially solid objects (*dog*, *cup*, *book*). In learning these kinds of nouns, it has been shown that children tend to generalize on the basis of shape, rather than material, color, or texture, for example. This tendency is called the **shape bias**. For example, consider the following experiment, which could be performed with a child of two-and-a-half years, say. The experiment begins with a training phase (above the line in the figure below) in which the child is shown an unfamiliar object and hears it labeled with a new word, such as *dax*. Next, in the test phase (below the line in the figure), the child is shown a set of objects, each of which matches the original object on one or more dimensions, and asked to find the one that the word best applies to ("show me the *dax*"). In the figure, the first test object matches the training object on shape, the second matches on color, and the third matches on texture pattern. In other words, the child is being asked to generalize on the basis of one example to another. In experiments like this, children tend to pick the object that matches the original object in shape, in the figure, the first object.



The shape bias is just a tendency; children do not exhibit the bias for (non-solid) masses, and they exhibit less of a shape bias for objects that they believe are animals, for example, things that appear to have eyes.

How the Shape Bias Can Help in Noun Learning

The shape bias can help children in the learning of noun meanings because it restricts the possible semantic categories to those in which shape, rather than color or texture, is a relevant dimension. If a child learning the word *flute* sees an example of the word, say a silver concert flute, the child can later extend the word to other objects that are similar in shape to the original example but avoid extending it to other objects that are similar only on other dimensions, for example, to a silver teapot. But where does the shape bias come from? One view is that it precedes word learning; it is either innate or it is learned on the basis of the child's early experience with objects.

But researchers led by psychologist Linda B. Smith have shown that there is a simpler account and one that agrees better with children's behavior. There are two details of their behavior that are relevant. First, the shape bias seems only to apply in linguistic tasks, that is, when children are labeling objects. When they are asked to group objects on the basis of their similarity, for example,

they don't necessarily base their groupings on the shapes of the objects. This implies that the bias is not necessarily a general cognitive bias but only a bias that is relevant for language and that it could not be learned on the basis of pre-linguistic experience. Second, even for language, the shape bias does not appear until children have learned fifty or so words. The implication is that the shape bias is learned on the basis of the language that the children are exposed to. The research has shown that most of the early nouns children learn refer to categories that *are* defined by shape, categories such as cup and horse and apple. The idea is that as children learn more and more words of this type, they make the generalization that *shape matters* for nouns and then go on to use this generalization to help them in learning further nouns; shape is what they pay attention to when they are learning new nouns for objects. In other words, for this small aspect of language learning, at least, if children are good statistical learners, special constraints are not required.

For a more in-depth discussion of the shape bias and related issues, see this 2000 paper by Smith and cognitive scientist Eliana Colunga.

This page titled [2.7: Learning Meaning](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.8: Problems

1.8.1 Word Senses and Taxonomies, Metaphor and Metonymy

In these sections, we saw examples of four ways in which the meanings of words can be shifted or extended or mislearned (in the case of children or second-language learners): generalization, specialization, metaphor (based on similarity), metonymy (based on close association, belonging).

For each of the following examples, say which of these kinds of shifts or extensions is involved, explaining your answer. The extended/shifted word is in **boldface** in each case.

Example 2.8.1:

(The speaker is an IU student discussing a basketball game; he does not play on the team.)

*You should have seen **us** beat **Purdue** last night.*

Answer

Both are metonymy. There is an strong three-way association between a school, its athletic teams, and its students, so the speaker can refer to Purdue's team as "Purdue" and his own team as "us". The student and the team "belong" to the university.

1. *They're performing **Shakespeare** tonight.*
2. A baby uses the word **car** for toy cars, but not real cars.
3. *What a **jock**. Is there any sport he's not involved in?*
4. *Don't pay any attention to what she's saying. It's just **bullshit**.*
5. The word **tool** was originally used only for physical devices such as hammers and axes. Now its meaning includes more abstract devices such as software to help people design machines or produce art.
6. The English words **beef**, **pork**, and **mutton** were borrowed from French, where they referred to animals rather than meat: 'ox', 'pig', 'sheep'.
7. The English word **infant** derives ultimately from a Latin word meaning 'a very young child who has not yet learned to speak'. Of course we now use it for all very young children, whether they can speak or not.
8. Originally the English word **vaccine** was used only to refer to a weak form of the smallpox virus used to prevent smallpox. Later it was extended to include any substance that works in the same way against a disease.
9. The English word **igloo**, usually referring to a house made of blocks of snow, comes from a word in Eskimo languages meaning simply 'house'.
10. The English word **clock** comes from a French word meaning 'bell'. Before there were clocks, it was the bells in church towers that people in Europe relied on for knowing the time of day. Later those bells were replaced by clocks located on the same church towers.

2.8.2 Differences Between Languages

In this section, we saw that in the domain of personal pronouns, languages may make use of different dimensions to distinguish the different forms from each other. Below are the Tzeltal personal pronouns (San Jerónimo dialect). Decide which dimensions matter for Tzeltal personal pronouns, what the possible values are for each dimension, and what the values on each dimension are for each pronoun. Remember that a dimension value can be unspecified for a given pronoun (for example, *you* in English is unspecified for gender). Remember also that each word must have a unique combination of values; otherwise you have not distinguished all of the words from one another. Of course you don't have to know how these words are pronounced, but when writing them, note that the apostrophe is a separate letter in the Tzeltal alphabet.

- *jo'on* 'I'
- *ja'at* 'you (one person)'
- *ja'* 'he, she, it'
- *jo'otik* 'we = I and you and possibly others'

- *jo'otkotik* 'we = I and others but **not** you'
- *ja'ex* 'you (more than one person)'
- *ja'ik* 'they'

	Person	Number
jo'on	1	Singular
ja'at	2	Singular
ja'	3	Singular
jo'otik	1, 2	Plural
jo'otkotik	1	Plural
ja'ex	2	Plural
ja'ik	3	Plural

	Speaker	Hearer	Number
jo'on	Yes	No	Singular
ja'at	No	Yes	Singular
ja'	No	No	Singular
jo'otik	Yes	Yes	Plural
jo'otkotik	Yes	No	Plural
ja'ex	No	Yes	Plural
ja'ik	No	No	Plural

	Person	Exclusiveness	Number
jo'on	1		Singular
ja'at	2		Singular
ja'	3		Singular
jo'otik	1	No	Plural
jo'otkotik	1	Yes	Plural
ja'ex	2		Plural
ja'ik	3		Plural

This page titled [2.8: Problems](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

3: Word Forms - Units

- [3.0: Introduction](#)
- [3.1: Phonemes](#)
- [3.2: Iconicity](#)
- [3.3: Vowels \(Part 1\)](#)
- [3.4: Vowels \(Part 2\)](#)
- [3.5: English Consonants](#)
- [3.6: Consonants in Other Languages](#)
- [3.7: Syllables](#)
- [3.8: Problems](#)

This page titled [3: Word Forms - Units](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.0: Introduction

In the last chapter, we focused on word meanings, saying nothing at all about word forms, that is, how words are produced and what they sound or look like. In this chapter, we'll focus on form, specifically the units that make up words and how those units are combined. We'll see that, just as meaning is organized around the way that people divide the continuous world into a set of categories, form is organized around the way that people divide the continuous space of possible sounds (or gestures for sign language) into a set of categories. We'll also meet dimensions and values again, this time dimensions relating to how the units of form are produced by the Speaker. We'll also see that while different languages have quite different sets of categories, different patterns for combining them into words, and even different dimensions, there are also obvious constraints on what is possible in word forms in human languages. This, and the next chapter, will be the only ones in the book dedicated almost exclusively to form. Linguistic form, however, is always in service of meaning, and we will see that the function of the basic form categories is to distinguish words, that is, patterns with different meanings, from one another.

This page titled [3.0: Introduction](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.1: Phonemes

Vowel Categories

Exercise 3.1.1

Each word in a language needs its own form. From the perspective of the Hearer, what should be generally true for these forms? From the perspective of the Speaker, what should be true?

We saw in the last chapter how there may be an infinite range of possible things in the world, even an infinite range of perceivable things, and that people cope with this by categorizing what they experience, by grouping the very large set of possible things into a relatively small set of categories. We saw also that the categories that people have for things are learned, that people seem to be very good at learning categories for what's around them, in particular for the categories that they label with words. There are good reasons for these categories. People can learn a particular response or set of responses for a category and then apply it to all instances of that category. They don't have to start over with each new thing they face. For language this means that Speakers can use a known word for a novel situation and that Hearers can understand known words when applied to novel situations.

Now let's consider the other end of words, their form. Linguistic forms require two **channels**, one for production, one for perception. The production channel could theoretically be any means we have of creating distinguishable signals. Probably because of constraints related to the human body and how it is controlled by the human nervous system, the only known human production channels are speech, signing, and writing. The perception channel makes use of a sensory system capable of distinguishing the signals produced on the production channel. This sensory system is audition (hearing) for speech and vision for signing and writing.

The Simplest Sounds We Can Make

Let's return to our imaginary tribe of Lexies, who are enjoying the communicative advantages brought by the invention of nouns. We've had nothing to say about how these words are produced. Our Lexies could have started with signing, using the hands, arms, upper body, and face, and we'll return to this possibility later. But for now let's assume they started with speech, using the vocal tract, the lips, tongue, oral (mouth) cavity, nasal (nose) cavity, throat, larynx (voice box), and lungs. In the beginning the idea is relatively simple, consisting of the following actions.

1. Opening the oral cavity some amount
2. Placing the tongue in some position within the oral cavity that allows air to pass through
3. Tightening the vocal cords in the larynx
4. Expelling air from the lungs to cause the tightened vocal cords to vibrate

By varying 1 and 2, the Lexies get different sounds. We'll call these sounds vowels. Here are some examples of the vowels they get.

- vowel 1
- vowel 2
- vowel 3
- vowel 4

Playing around with this mechanism, the Lexies realize that there are very many different sounds, maybe even an infinite number of sounds, that they can produce. Here is what the playing around sounds like.

- playing with vowels

Vowel Categories and Word Form Categories

But what the Lexies require is *not* an infinite number of word forms; they require exactly one word form for each word. Just as they group the very large number of possible things in the world into a relatively small set of categories, they need to group the very large number of possible word forms into a small set of *categories*, one for each word. Let's assume that early on in the history of their developing language, they have only four words, two proper nouns naming important individuals in the tribe and two common nouns naming important categories of things in their world, tiger and apple.

Given their vowel mechanism, they could easily develop four separate categories for the four words. Let's say the four categories are centered on vowels 1-4 above; that is, these vowels are the prototypes for the categories. But recall what it means to be a category. Each actual pronunciation of one of the four words is an instance of one of the four word form categories. No two of these instances are likely to ever be exactly the same. So even if I know how to produce the four word forms, there will be some variation. For example, the following might be some possible instances of the word meaning apple, which is centered on vowel 1 above.

- three instances of the word for apple

The differences between the vowels are very real differences, both in terms of how they are produced and what they sound like, but these differences would not *matter* to the Lexies because the meaning in each case is the same, apple. On the other hand, the difference between any of the above vowels and the following vowel *would* matter because this is an instance of the word form that means tiger.

- the word for tiger

Differences that Matter and Differences that Don't

The vowel in the word for tiger **contrasts** with the vowels in the three instances of the word for apple because it makes a difference in the meaning of the word. The vowels in the three instances of the word for apple do not contrast with each other, even though they do differ from each other, because they do not make a difference in the meaning of the word.

Given the very simple language at this stage, let's consider what the members of the tribe need to store in long-term memory. To say the words, they need to remember, for each one, roughly how to place their tongue and their jaw, and to *recognize* the words (in order to be able to understand them), they need to remember, for each one, roughly what it sounds like. I won't go into the acoustics of vowels here; all you need to know is that there are ways in which it corresponds roughly to the placement of the tongue. While this is quite a lot of information to remember for each vowel, there are only four of them, so it does not tax long-term memory too severely.

Syllable and Consonant Categories

Exercise 3.1.2

What would the disadvantage be of a language whose word forms consisted only of single vowels?

Now let's see what might happen when the Lexies find the need for more words as they find the need to refer to more and more individuals and categories. They might get to fifteen or twenty words using the approach we've described so far, but at this point they'd begin to run into problems as the categories become more and more similar. From the Speaker's perspective, it will be more and more difficult to produce the sounds accurately enough so that they don't overlap. More importantly, from the Hearer's perspective, it will be more and more difficult to distinguish the sounds as they become closer and closer to one another acoustically.

Clearly the Lexies need a way to come up with more easily distinguishable words. Through experimenting more with their vocal tracts, they discover that they can produce a much larger set of sounds by *moving* their lips, tongue, and/or jaw *during* the production of the word. Each of these new sets of sounds is produced by performing the following sequence of actions.

1. Closing the vocal tract at some point using the lips or tongue.
2. Beginning the vibration of the vocal cords.
3. Opening the vocal tract by releasing the contact that was made by the lips or tongue.
4. Moving the tongue and jaw to the position of one of the vowels as the vocal cords continue to vibrate.

I'll call the new sounds that result from movement of the lips, tongue, and/or jaw, along with the simpler ones that don't involve movement, that is, that consist just of vowels, syllables. Here is a syllable that results when the contact at the beginning is made with the lips.

- lip contact + vowel

Here are three syllables with the contact made by the tip of the tongue, all with the same vowel.

- tongue top contact + vowel

Here are three syllables with the contact made by the middle and back of the tongue, also with the same vowel.

- tongue body contact + vowel

The important point about these different ways of making the contact is that they result in qualitatively different sounds to hearers. Thus using syllables greatly increases the number of distinguishable words the Lexies have access to. But note that, as with vowels, there is a very large number of possibilities for how the lips and tongue are placed for the contact. Since the Lexies only need a different syllable for each word, what is called for again are *categories*. Each syllable representing a word is a separate category. To remember a word, then, a Lexie would need to remember how the syllable for that word sounds and how to produce the contact at the beginning and the vowel part at the end. Again this is a lot to remember for each syllable, but if there aren't too many, it might be manageable.

Syllable Categories Consisting of Consonant Categories and Vowel Categories

Let's say the Lexies make use of five easily distinguished vowels, and five easily distinguished initial contacts, which we will call consonants. In all that gives them 25 (5×5) possible syllables. To remember all of these, they'd have to store 25 separate sets of production instructions and acoustic data in long-term memory. But there is a more efficient way; they could treat each of the five contact possibilities as categories and each of the five vowels as categories — only ten things to remember. Then to produce or recognize a syllable, they would only need to combine the information for the consonant and the vowel. And to remember a word form, they would only need to remember what combination of consonant and vowel was used.

Unfortunately it's not quite that simple, especially on the perception (acoustic) end, because the way a consonant sounds depends a lot on the vowel that follows it. This means that it is probably impossible for people to remember the acoustic properties of consonants independently from vowels. And this is why there are people who argue that the basic categories making up word forms are syllables, rather than consonants and vowels.

Phonemes

Exercise 3.1.3

Why do you think the writing systems of some languages, such as English, Spanish, Lingala, and Tzeltal, use separate characters for the consonants and vowels rather than separate characters for each syllable? (Note: the writing systems of many languages, for example, Amharic, Japanese, Chinese, and Inuktitut, do use separate characters for each syllable.)

Let's recap where the Lexies are. Their words, whose meanings we learned about in the last section, need forms. Each word form needs to be stored in long-term memory; they need to be able to remember how to produce it and what it sounds like. Each word form is a category because it represents a set of possible forms, a very large, possibly infinite, set. Word forms consist of syllables, each representing for the Speaker either a particular open position of the vocal tract (that is, a vowel) or a closing of the vocal tract followed by an opening, ending in a particular position (that is, a consonant followed by a vowel). The most efficient way for the Lexies to remember all of the possible syllable categories is to remember the consonant categories and the vowel categories.

Thus the consonant and vowel categories make up a kind of alphabet to be used in making new words. Given a concept that the Lexies would like to refer to, all they need to do is find a new (previously unused) combination of a consonant and a vowel to make the new word form. The elements of this basic alphabet for linguistic form are called phonemes; each vowel or consonant category is a phoneme of the language. The phoneme is one of the most important concepts in modern linguistics, and we will spend a lot of time getting a handle on it. For now, the important points to keep in mind are the following.

Variability Within Languages and Between Languages

1. Each phoneme is a category; that is, it represents a cluster of possible consonant or vowel instances, centered on a prototype.
2. The way the space of possible sounds (consonants and vowels) is divided up into phonemes is to some extent arbitrary. That is, it can be expected to vary from language to language; the phonemes that a given language has are conventions. What is important is that the phonemes are distinctive enough to be distinguished by hearers. Thus one tribe of Lexies might end up with five vowel phonemes, another tribe might also end up with five vowel phonemes but centered on different sounds from the first tribe, and a third tribe might end up with eight vowel phonemes.

3. Speakers can produce differences that contrast, that is, that represent different phonemes and make a difference in meaning, and differences that do not contrast, that is, that represent different ways of producing the same phoneme and do not change the meaning.
4. Phonemes are *not* letters. Letters are the basic elements of a writing system, which may or may not have been designed to represent the phonemes of a language. But the phonemes of a language such as English are represented only very imperfectly by the English alphabet, as we will see, and other languages, such as Mandarin Chinese, have writing systems that do not even pretend to represent phonemes. Equally importantly, most languages are not written at all.

The word forms in all modern languages can be described in terms of a small set of phonemes, though the set is different for each language (and may vary between dialects of the same language, as we will see for English dialects). The number of phonemes varies from about 13 to about 150; English has about 40, depending on how we count and which dialect we are considering.

Extending the Lexicon

More Vowels, More Consonants, More Syllables

With only 25 or so different syllables to choose from, the Lexies will only be able to make use of a very small set of words. There are several ways in which they can extend the set of possible word forms. One is to produce more possible vowels by varying something other than the position of the tongue and opening of the jaw. For example, they could round their lips during the production of the vowel and get a different effect, or they could move their tongue, jaw, and/or lips during the production of the vowel to get a combination of vowel sounds. Another way to extend the number of possible word forms is to produce more possible consonants by varying something other than how the contact is made at the beginning of the syllable. For example, they could make the contact incomplete, allowing some air to pass through the gap that is made, or they could delay the beginning of the vibration of their vocal cords until after the contact is released. A third possibility is to produce more complicated syllables, for example, by allowing the vocal tract to close at the end as well as the beginning of the syllable. A final possibility is to combine syllables to make longer word forms. If there are only 25 possible syllables, there are 625 (25×25) possible two-syllable words.

All four of these ways of coming up with more word forms are used in modern languages. We will be discussing them all at length later in this chapter.

Visual/Spatial Phonemes

Exercise 3.1.4

There seems to be nothing in sign languages corresponding to vowels and consonants, though there are syllables. What dimensions do you think would distinguish different sign syllables from one another?

Sign Languages also have Categories of Form.

One tribe of Lexies takes a different approach from the others. Instead of using their vocal tracts to produce words, they use their hands and arms. Rather than beginning like the speaking Lexies with static vowel-like patterns to represent the first words, these signing Lexies realize that hand/arm *movements* would be easier to see than static hand/arm configurations. As with speaking, however, it is clear that there are very many possible movements, and they realize they need movement *categories*. So they settle on a small set of such categories and use these in combination to make word forms.

Like spoken languages, modern sign languages also have phonemes as the units that words (signs) are built up from. As with the signing Lexies, in modern sign languages, the phonemes consist of movements of the fingers, hands, or arms. Corresponding to dimensions such as the position of the tongue or lips for spoken language are the particular configuration of the hands, for example, which fingers are extended; the position of the hands with the respect to the body; and the direction of movement of the hands, fingers, and arms. Also, as with spoken languages, we can expect the specific set of phonemes to vary from one sign language to another. However, research on the phonemes of sign languages is still relatively new, so there is not yet as much agreement on their properties as there is for spoken languages. In the rest of this chapter, we will be focusing mainly on the phonemes of spoken languages and how they are combined to form syllables and words.

This page titled [3.1: Phonemes](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.


3.2: Iconicity

Iconicity and Arbitrariness in Spoken Words

If you had to invent a word for monkey, what word form would you choose to make it as memorable as possible to other speakers? What about a word for axe?

We have discussed two aspects of words, their meanings and their forms (and in the rest of this chapter, there will be lots more to learn about forms), but we have not yet considered the third aspect, the connection between form and meaning. How do people settle on which forms to use with which meanings?

Recall that the speaking Lexies begin with a small set of vowel phonemes for their first word forms and that these are used to refer to a set of individuals and categories. One category is tiger. Which vowel would be a good choice for this category? The Lexies more often hear than see the tigers in their vicinity, so they are familiar with the sound the animals make. While they are not capable of a very good imitation of the roaring of the tigers, it is reasonably clear to them which of their vowels are further and which are closer to this sound. So they choose this vowel to refer to this category. The figure below diagrams the relationship. The form of the word is indicated by *aaa*. In addition to the form-meaning relation joining *aaa* and tiger, there is a similarity relation joining the word form to the tiger's sound, as well as the relation between a tiger and its sound.



Apples don't make sounds of course, but a sound is produced when we eat them. A vowel by itself is a very poor imitation of this sound, however, and the Lexies have to wait till they have syllables that have vocal tract closures at the beginning and end to produce something passable, for example, this sound.

So far we've seen that the imitation of sounds is a possible *motivation* for the choice of a form for a word. But it's not the only possible motivation. We'll see others later on in this section and other sections. For now, the important point is to distinguish two different kinds of form-meaning connections, those that are motivated, for example, based on a resemblance between the sound of the word and a sound associated with the meaning, and those that are not motivated, where the association between a form and its meaning is apparently random. The first type of connection is referred to as **iconic**, the second as **arbitrary**. For most categories in the world, there will be no obvious motivation for the spoken word form. For example, what sound would suggest the category house or the category sand? In such cases spoken languages have no alternative but to arbitrarily associate forms with meanings. The forms of the English words *house* and *sand* don't suggest their meanings at all, and this is probably true for most, if not all, other spoken languages.

Consider what the advantages of iconicity might be. Someone learning a new word has to learn three things: the meaning, the form, and the association between the meaning and the form. We have discussed the learning of word meaning, and, in the next chapter, we'll discuss the learning of form. It seems clear from what is known about long-term memory that the more the learner has to go on, the easier it is to remember an association. Say a Speaker (and learner) of our Lexies' language is thinking of the category tiger and trying to remember how to refer to it. If the word-meaning association is arbitrary, the Speaker has nothing more than the concept itself to go on. But if (as shown in the figure above), there is an additional similarity relation connecting the meaning (or an associate of the meaning, the tiger's sound) and the form, the Speaker has an extra path to get to the word form. In fact it is not known if iconicity actually helps Speakers (and learners) in this way; ongoing research should give us some clues.

In modern spoken languages, iconicity is not very common among the nouns used for things. In English, we have a number of words imitating the sounds of animals, for example, *bow-wow*, *roar*, *cluck*, and *chirp*, and for various natural and artificial events, for example, *splash*, *rumble*, and *bang*, but we don't usually use these words as nouns to refer to the things that make the sounds. In Japanese it is common in speech directed at children to use animal sounds to refer to the animals that make the sounds. For example, *wan-wan* may mean 'dog', and *moo-moo* may mean 'cow'.

Why Dogs Sound Different in Japanese and English

These examples of animal sounds bring up the obvious differences between languages in their imitations of what are the same sounds in nature, the sounds that dogs or cows make, for example. These differences result in part from the different phonemes used in the languages; thus a word form like *crash* is not possible in Japanese, as we will see later because Japanese syllables cannot begin with sequences like *kr* or end with a consonant like *sh*. Another reason for these differences is that the imitations are only very imperfect (people are not very good at imitating sounds from nature), and a great deal of variation is possible in how they're achieved. Thus the vowel Japanese speakers use in their word for the sound of a cow is different from the vowel English speakers use in their word. When a form is iconic, the implication is not that a person with no knowledge of the language could predict the form given the meaning; the form is still a convention. The point is only that the form bears *some* non-arbitrary relation to the meaning.

There is another potential difference between languages with respect to iconicity. It seems that some languages are more iconic than others. Japanese, Amharic, and Lingala, for example, all have entire categories of words, called **expressives** numbering in the hundreds or thousands that seem to have some sort of iconicity as a defining feature. The iconicity is often not based on the imitation of a sound from nature but may relate more to how the position or the movement of the organs of speech directly corresponds to something in nature. For example, the Japanese word *pyon-pyon*, roughly 'jumping movement', seems to suggest its meaning through the "jumping" of the lips. Investigation of these aspects of spoken languages is still very inconclusive and controversial, however.

Iconicity in Signs

Exercise 3.2.1

If you had to invent a *sign* for monkey, what form would you choose to make it as memorable as possible to other signers? What about a sign for axe?

The fact that the most form-to-meaning relations for words in spoken languages are arbitrary rather than iconic has led many linguists to believe that iconicity is only a marginal phenomenon. An examination of sign languages may lead us to a different conclusion.

Let's return to the Lexie tribe that decided to go with visual/spatial word forms rather than spoken word forms. For tiger there is now no possibility of imitating the sound of the animal. But a tiger has more visual properties than acoustic properties. They could choose to represent the stripes, the gait, the tail, or the whiskers of the tiger. And for categories that have no obvious sound associated with them such as tree, lake, and sand, the signing Lexies can resort to representing the shape of the thing, the way it is used by people, or the way it responds when it is manipulated in some way. The fact that signs are executed in space means simply that sign language offers more possibilities for iconicity than spoken language does.

Sign languages are More Iconic than Spoken Languages.

It is not surprising then that in modern sign languages, the lexicons are full of iconicity. Here are some examples from American Sign Language.

In the case of the sign for tiger, the movements of the hands convey something of the appearance of the members of the category, namely the tiger's stripes. The sign for apple is more abstract; the signer's cheek represents the fruit, and the moved finger represents the stem of the fruit being twisted off, as when one is going to eat the fruit. The sign for key represents the action of using the key to lock or unlock something.

The ways in which these signs are iconic should remind you of what you learned about metaphor in the last chapter. The source domain in this case is the real world in which things have shapes, move about, and get used by people. The target domain is the signing space in which signers move their hands and arms. Metaphor, and iconicity, is possible because there are mappings between the two domains. Note that signs are also metonymic; the sign for tiger represents a part of the animal rather than the whole animal. One linguist who focuses on iconicity and metaphor in sign languages is Sarah Taub.

In sum, iconicity makes sense since it facilitates memory for the association between form and meaning, so it should not surprise us that language will sometimes capitalize on iconicity when this is possible. There are many more possibilities when the medium is

visual/spatial than vocal/auditory, and the great majority of spoken word forms are probably not iconic. We will see more examples of iconicity in both spoken and signed language later in the book, both for words and grammar.

This page titled [3.2: Iconicity](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.3: Vowels (Part 1)

The goals of linguists, if you remember, are to describe what people know about their language and to figure out how languages are similar and different. For word forms, specifically for phonemes, this means that we must describe both how speakers and hearers distinguish phonemes *within* a given language and how individual phonemes and systems of phonemes differ *between* languages. To satisfy both of these goals, we will be looking for ways to describe the variation between speech sounds. Just as we saw that categories of things could be described in terms of values along different dimensions, we will be looking at dimensions of sound, dimensions that allow us to make distinctions between phonemes within and between languages.

Phonetic Symbols

Exercise 3.3.1

If we want to write down how words in different languages are pronounced, why not just use the letters of English to do this? Could we use the letters of English to write down the pronunciation of English words? What if we want to distinguish the pronunciation of words in different English dialects?

Before we discuss the specific phonemes of languages, we need to decide how to represent the sounds. Let's briefly consider English. Like other spoken languages, English has a set of vowel and consonant phonemes that its speakers use to make the words of their language. Like other written languages, English also has a way of graphically representing spoken word forms. The English writing system is an example of an **alphabetic** writing system in which phonemes are represented by characters or combinations of characters. But, for various historical reasons, the English writing system does this very imperfectly. Consider the words *way*, *weigh*, *wait*, and *wake*. These words share the same vowel phoneme, but it is spelled in four different ways. That is, a single phoneme may be represented using different letters or combinations of letters. Now consider the words *bother*, *brother*, *border*, and *voter*. These words share the letter *o*, but it represents four different vowels, each a different phoneme. That is, a single letter may represent multiple phonemes. We can conclude two things from this.

- We must be careful not to confuse sounds with letters; the letter *ois* *not* a vowel, though it is used to represent vowels.
- We cannot rely on English spelling when we are concerned with the pronunciation of English words.

Because we will need a way to represent the phonemes of English and other languages unambiguously, we must rely on a set of symbols for this that are *not* used quite like the alphabets of any alphabetic writing systems. Symbols representing the basic sounds, or **phones**, of spoken languages, are called phonetic symbols. Linguists use a set of phonetic symbols called the **International Phonetic Alphabet (IPA)**. The symbols in the IPA are based on the characters in the Roman alphabet, which is also the basis for the writing systems of many languages, including English, Spanish, Lingala, and Tzeltal. The IPA is overseen by the International Phonetics Association. Here is a full chart of the IPA symbols that you can click on to hear some of the sounds. Note that you probably won't understand this table until after you've studied this section, the next one on English consonants, and the one after that on consonants in other languages. I will use a subset of the IPA symbols in this book. (To view many of these symbols, you will need Unicode support in your browser. Go to this appendix to see if your browser has this support.)

Symbols for Phonemes vs. Symbols for Phonetic Details and Language Differences

I will use phonetic symbols in two ways. Much of the time they will be used to distinguish the phonemes within a language. The important thing will be to make sure each phoneme has a unique symbol; which symbol we use is not as important. When the symbols are used in this way, they will be enclosed by slashes. As an example, my pronunciation of the word *phonemes* is /'fonimz/. I will also be using symbols in another way: to represent the actual sound that is produced in a given situation, rather than a category of sounds (a phoneme, that is), and to represent sound differences between different dialects and different languages. For this purpose we will have to be more careful about which symbols are used for which phones. When the symbols are used in this way, they will be enclosed in square brackets. As an example, my (more detailed) pronunciation of the word *phonemes* is ['foʊnɪmz]. It is important to note that the symbols are at best only an approximation to the actual phone that is used in any given situation because the space of possible phones is enormous, possibly infinite, while we only have a relatively small, finite set of phonetic symbols.


Vowel Features

Exercise 3.3.2

Notice what you do with your mouth when you pronounce the names of the letters "a", "e", "i", and "o". What features of your mouth seem to distinguish these vowels from one another?

A complete account of the sounds of a spoken language (or of spoken language in general) would have to make reference both to the way sounds are produced, articulation, and the way they are perceived, their **auditory** properties. In this book we will not have time to go into the auditory properties of speech in any detail, but we cannot neglect them completely. Differences in articulation are obviously pointless if they aren't reflected in the way the results sound to hearers. In fact, two quite different ways of articulating a sound can sometimes produce the same auditory effect. From the perspective of a hearer, those would have to belong to the same phoneme.

In order to understand how vowels and consonants work, we need to know a little bit about the physical apparatus that is used to produce them. The figure below shows a side view of the vocal tract, with labels for some of the parts that we'll be discussing in this section and the next sections.




The heart of the system is the larynx, located in the throat under the Adam's apple; the larynx contains the vocal cords, two pieces of flesh that can be loosened or tightened. When air from the lungs passes through the larynx, it may be allowed to pass unhindered, as when we are breathing out. Or the vocal cords may be tightened and brought together so the air causes them to vibrate. This vibration, or voicing, is what distinguishes normal speech from whispering and certain speech sounds from others. Here is a slow-motion movie, produced at the UCLA Phonetics Laboratory, of the vocal cords vibrating during speech.

How the Vocal Tract is Superior to a Trumpet

Given an outward airstream and vibrating vocal cords, we have a device that is a little like a brass musical instrument. The vocal cords are like the vibrating lips of the musician, and the region of the vocal tract between the larynx and the opening of the mouth (the yellow region in the figure), the oral cavity, is like the body of the horn. What makes the vocal tract far more versatile than a trumpet or a trombone, however, is that speakers can change the shape of the instrument as they are playing it, producing a great variety of sounds.

What distinguishes vowels from each other auditorily is the precise shape and volume of the oral cavity, and the main organ involved in adjusting the shape and volume is the tongue. Given the constraints on the shape of the vocal tract and the way the tongue is manipulated, there are three extreme positions the tongue can take that lead to vowel sounds. If the body of the tongue is pushed forward and toward the roof of the mouth, we get this vowel, something like the vowel in the word *beat* and written with the symbol [i]. If the body of the tongue is lowered while the back of the tongue is pushed toward the roof of the mouth, we get this vowel, something like the vowel in the word *boot* and written with the symbol [u]. Note that in both cases the tongue cannot approach the roof of the mouth too closely, or we get the sound of friction near the contact, that is, a consonant rather than a vowel sound. Finally, if the body and back of the tongue are both moved away from the roof of the mouth, at the same time narrowing the region that is between the larynx and the mouth (the "pharynx") we get this vowel, something like the vowel in the word *hot* (as pronounced by most Americans) and written in this book with the symbol [ɑ]. (The IPA symbol for this sound is the variant of *a* in

the Roman alphabet that resembles this Greek letter.) The configurations of the vocal tract for these three vowels are shown in the figures below.



How can we turn this informal description of what is going on in the vocal tract into a more compact description that makes it clear how vowels differ from one another? First, [i] and [u] share one property; they are both associated with a relatively narrow gap between the tongue and the roof of the mouth (and a relatively wide pharynx). They differ in this way from [ɑ], which has the widest possible gap between the body of the tongue and the roof of the mouth (and a relatively narrow pharynx). Just as we spoke of dimensions along which categories of meaning varied, we will speak here of dimensions along which these sound categories vary. The dimension that distinguishes [i] and [u] from [ɑ] we'll call height. So far we have seen two different values on this dimension, one for [i] and [u] and another for [ɑ].

Another, equivalent way to talk about the difference between these vowels is in terms of **features**; [i] and [u] have the feature high, while [ɑ] has the feature low. A feature always corresponds to some value on a dimension. For example, with the domain of personal pronouns we saw that the gender dimension had possible values of masculine, feminine, and neuter; each of these can also be seen as feature. All of the things having a certain feature can also be viewed as a category; thus there is the category of masculine pronouns and the category of high vowels.


Turning a Huge Number of Mouth Configurations into a Small Set of Categories

Dimensions can be **discrete**, that is, with a finite set of possible values, or **continuous**, that is, with an open-ended and possibly infinite set of possible values. Gender is a discrete dimension — there are only three possible values — while vowel height is a continuous dimension because we can theoretically put the highest part of the tongue at any point between the lowest and highest possible points. The fact that height is a continuous dimension means that the values high and low identify *approximate* points along the dimension and that a given vowel is only *relatively* high or low. But note that the fact that a dimension is continuous does not imply that people can distinguish all of the possible values; they cannot. For example, speakers can probably reliably produce and hearers can probably reliably distinguish no more than about five different vowel heights. Note also that what is a continuous dimension in the world may be treated as discrete within language (or elsewhere in cognition). Thus within a given language, as we'll see soon, only a small number of height distinctions are made.

The dimension of height distinguishes [i] and [u] from [ɑ], but it does not distinguish [u] from [i]. Looking at the diagrams of the vocal tract again, we can see that these two vowels differ mainly in terms of where the narrowest gap between the tongue and the roof of the mouth is. For [i] it is as far forward as it is possible for the tongue body to move and still produce a vowel sound; for [u] it is as far backward as it is possible for the tongue body to move and still produce a vowel sound. This dimension is called backness; again we have seen (so far) two possible values on this dimension: front, as for [i], and back, as for [u]. But note that, like height, backness is a continuous dimension, so front and back are labels for approximate regions along the dimension. For [ɑ], the closest point of approach between the tongue body and the roof of the mouth is near the back of the mouth, so this is also considered a back vowel. Summarizing, the two articulatory dimensions of height and backness represent the location in the mouth of the closest point of contact between the tongue body and the roof of the mouth.

It turns out that these two articulatory dimensions correspond roughly to two fundamental auditory dimensions. That is, as a speaker changes the height of a vowel, the way that vowel sounds to a hearer changes along one auditory dimension, and as a speaker changes the backness of a vowel, the way that the vowel sounds to a hearer changes along another auditory dimension. However, some recent research has shown that the correspondence between the two articulatory and the two auditory dimensions is only approximate. For vowels other than the extreme ones we've been discussing, there are often multiple articulatory ways of achieving the same auditory effect; precisely adjusting the position of the closest approach between the tongue body and the roof of the mouth is only one of these ways. Thus in a sense it is a bit of an over-simplification to characterize vowels in terms of their precise height and backness. However, height and backness still seem to provide the simplest articulatory description of vowels, and I will follow this traditional approach in what follows.


If every vowel has a height value and a backness value, then we can visualize all of the possible vowels in a two-dimensional space known as vowel space. Given the constraints of the mouth and tongue, it turns out that vowel space has roughly the shape of a trapezoid; this is shown in the figure below, along with some of the terms used to refer to different regions within vowel space. Recall that these terms are only for convenience; there are really very many, maybe even an infinite number of, possible positions within this space.




Vowels of Spanish

Spanish has Five Vowels, None Quite like English.

We'll start with a relatively simple vowel system, that of Spanish. Spanish has five vowel phonemes, represented in the Spanish alphabet by the letters *a*, *e*, *i*, *o*, and *u*. The same characters are conventionally used to represent the vowels themselves. Listen to the vowels in these five Spanish words, and try to pronounce them yourself (here and elsewhere in this book, we will focus on Western Hemisphere, rather than European (Castilian), Spanish): *piso*, *puso*, *peso*, *poso*, *paso*. You may not be aware of the position of your tongue with respect to the roof of your mouth, but you should be aware of how the position of your tongue changes as you change from one vowel to another and how similar particular pairs of vowels are to each other in terms of tongue position. The figure below shows the approximate positions of the Spanish vowel phonemes in the vowel trapezoid. The positions are approximate because each phoneme is a category with a fairly wide range of possible pronunciations, depending in particular on which consonants appear before or after the vowel. You can click on each vowel in the figure to hear its pronunciation.



We can describe each of the Spanish vowels along the two dimensions of height and backness. For Spanish we need three values of height and two values of backness to do this. There is a further dimension that by itself does not distinguish any of the vowels, but it will be useful for us in discussing how languages and dialects differ from one another. This is the degree of lip **rounding** that accompanies the vowel. In Spanish /u/ and /o/ are the only vowels accompanied by significant lip rounding. For /i/ and /e/ the lips take the opposite position: they are spread. I will treat lip rounding and spreading as a single dimension. The figure below shows roughly where each of the Spanish vowels falls along the three dimensions of height, backness, and rounding. Each vowel is represented by three small circles of the same color, one for each dimension.



Another way to depict the differences is in a table that gives a label (feature name) for the value of each vowel on each dimension. I will use the symbol "Ø" to indicate that a phone has no value or a default value for a dimension.

	/i/	/e/	/u/	/o/	/a/
Height	high	mid	high	mid	low
Backness	front	front	back	back	back
Rounding	spread	spread	rounded	rounded	Ø

From the table, we can see that height is a **contrastive dimension** for Spanish vowels. That is, if we change nothing but the height of a vowel phoneme, we get a different phoneme: lowering the height of /i/ gives us /e/; raising the height of /o/ gives us /u/. The situation is a little more complicated for the other two dimensions because they tend to change together. Changing only the backness of a Spanish vowel phoneme does not give us a Spanish vowel at all. If we make /u/ a front vowel, we also have to make it spread to get /i/. If we make /e/ a back vowel, we also have to make it rounded to get /o/. If we make /a/ a front vowel, we also have to make it mid to get /e/. The same holds if we attempt to change only the rounding of a Spanish vowel. To distinguish all of the Spanish vowels, we actually need only *either* backness or rounding, but not both; that is, all we'd need to know about a Spanish vowel would be its values on two dimensions to know which vowel it is (verify this for yourself from the table). So we can call backness contrastive or rounding contrastive, or we can call the combination of backness and rounding contrastive. The point is that the system is **redundant**; it provides more information than is actually needed to make the distinctions.

Redundancy Can Come in Handy in Case the Hearer Misses Something.

Redundancy is common in language. We will see further examples with English vowels and in the grammar of various languages. Redundancy is a Hearer-oriented feature of language. When two or more different dimensions agree with one another, the Hearer can still extract the meaning even if they fail to note what is going on one of the dimensions. So with Spanish vowels, if the Hearer detects that a vowel is high and back but misses the fact that it is rounded, they can still know which vowel it is (/u/). The same would be true if they only detected that the vowel was high and rounded, but not that it was back.

Vowels of Japanese

The vowels of Standard Japanese are quite similar to those of Spanish. There are five phonemes centered on roughly the same positions in the vowel trapezoid as for Spanish. There are two differences. First, the high back vowel, /u/ in Spanish, is not normally rounded in Japanese; in Japanese it sounds like this. When we need to make it clear that it's an unrounded high back vowel we are discussing, we use the symbol [ɯ], but /u/ is conventionally used to represent this Japanese vowel phoneme.

Unlike in English, in Japanese Long Vowels are Just like Short Vowels, Only Longer.

A second difference between Spanish and Japanese is more significant. In Spanish, vowels tend not to vary too much in length, and how much they vary depends on the dialect. In any case, if the length of a vowel is changed in a word, we get the same word. So if a Spanish speaker pronounces the word *peso* with an extra-long /e/, the difference will probably be noticeable to a hearer, who might find it a bit odd but would not interpret it as a different word. In other words, vowel **length** is *not contrastive* in Spanish. Japanese differs in this regard. If we take a Japanese word with a short vowel and lengthen that vowel, we get potentially a different word. For example, *su* means 'nest', *suu*, with the same vowel, roughly twice as long, means 'number'; *koke* means 'moss', *kokee* means 'solid', *kooke* means 'succession'. (Note that *ee* and *oo* here represent long versions of *e* and *o*, not the sounds these letters would represent in English.) That is, vowel length is contrastive in Japanese. Note that this does not mean that changing the length of any vowel in a Japanese word must yield a different Japanese word, only that it potentially does. Thus *mise* 'store' is a word, but *miise* is not; *miira* 'mummy' is word, but *mira* is not.

There are two ways to analyze Japanese long vowels. We could consider each to be a separate phoneme, making ten vowel phonemes all together. Alternatively we could consider each long vowel to be two short vowels in succession. For our purposes, the choice doesn't really matter, and I will go with the second alternative (which is the usual way to treat Japanese vowels).

This page titled [3.3: Vowels \(Part 1\)](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.4: Vowels (Part 2)


Vowels of General American English

Exercise 3.4.1

Listen to the vowels in the words *bit*, *bet*, *bat*, *but*, *beat*, *bait*, *boot*, *boat*, *bought*, *bite*, and *bout*. At least how many distinct vowels does English have?

The first thing to note about the English vowels is that they are a great deal more complicated than the Spanish or Japanese vowels. This complexity has several consequences. First, there is considerable disagreement among linguists on the details of how to describe English vowels, and even some of the brief comments I have to make about them below may be disputed by one or another linguist. Second, there are considerable differences between dialects, and this makes an overall description of English vowels difficult, if not impossible. Third, the English vowel system (within a given dialect) is not as stable as a simpler system like that of Spanish, and it is more likely to change over time.

In this section, we will consider only the vowels of General American English (GA), specifically the variety that I speak. (If you are not a speaker of GA, some of your vowels may differ from those we'll be describing. Many of these differences are described in the section on English accents.) Let's begin with the front vowels. These are the vowels in the following words, each followed by the symbol I will use to represent the vowel: *beat* (/i/), *bit* (/ɪ/), *bait* (/e/), *bet* (/ɛ/), *bat* (/æ/). The approximate positions of these vowels in vowel space are shown in the figure below. Again you can click on a symbol to hear the vowel.



You can see that five vowels are squeezed into the front part of the vowel space, the same number of vowels that Spanish has all together. If the differences in height (and minor differences in backness) were the only dimensions distinguishing these vowels, this would put quite a burden on the hearer, who after all has to figure out that *beat* and *bit* are different words. It would also be a burden on the speaker, who would have to be very careful in placing the tongue for each vowel.

How English "Long" and "Short" Vowels Differ

So not surprisingly, these vowels are also distinguished from each other along other dimensions. We can put the highest four vowels in two groups, one consisting of the those traditionally called "long", /i/ and /e/, the other consisting of those traditionally called "short", /ɪ/ and /ɛ/. (The fifth vowel, /æ/, tends to group with the "short" vowels.) There are at least three differences between these two groups. First, the "long" vowels really do tend to be longer than the short vowels. Second, the "short" vowels tend to be "pure"; that is, the position of the tongue does not change during the production of the vowels. The "long" vowels, on the other hand, involve some movement, in both cases toward the upper left corner of vowel space. This is more noticeable for /e/ than it is for /i/ (but in other dialects, it is quite striking for /i/ too). Third, though you probably can't feel this yourself, for the "long" vowels, the base ("root") of the tongue is pushed forward. Finally, the muscles in the tongue are tenser than for the short vowels. Because values on the last two dimensions tend to go together, they are often considered to be a single dimension called [tenseness](#). The two extremes on this dimension are referred to as "tense" and "lax".


So we again have redundancy, as many as five different dimensions distinguishing /i/ from /ɪ/ and /e/ from /ɛ/. It appears that speakers do not always make use of all of these dimensions to achieve the differences that hearers hear in them.

Now let's look at the central and back vowels phonemes of GA. These are the vowels in the following words, with the symbols for each after them: *boot* (/u/), *put* (/ʊ/), *boat* (/o/), *bought* (/ɔ/), *pot* (/ɑ/), *but* (/ʌ/). (Note: many GA speakers do not make a distinction between /ɔ/ and /ɑ/; more on this in the section on English accents.)

English Vowels are So Complicated that Linguists Don't Agree on How to Treat Some of Them.

There is an additional central vowel, symbolized by /ə/, that appears in two quite distinct environments. (The IPA symbol for this sound is an upside-down e.) First, it is the most common vowel in unstressed syllables in English, that is, syllables that are weakly articulated. An example is the first phone in the word *about* and the last phone in the word *comma*. Some people consider this to be the same phoneme as /ʌ/; the main reason for treating it separately is that it behaves differently from /ʌ/ in the history of the language as well as in a number of modern dialects. Second, it appears followed by the consonant /r/ (in stressed and unstressed syllables). This combination is worth mentioning because speakers normally combine the two phones into one vowel sound. Examples are the beginning of the word *earth* and the end of the word *actor*. Again there is disagreement about how to treat the phone; some people would prefer to treat /ər/ as a single phoneme in the language.


The approximate positions of the back and central vowels in vowel space are shown in the figure below. Again you can click on a symbol to hear the vowel.



Again we find a large number of vowels squeezed into a region of the height-backness space, and, not surprisingly, the vowels differ on other dimensions. Like /i/ and /e/, /u/ and /o/ belong to the set of "long" vowels. Both share the features of the other long vowels: length, tendency to "move" as they are pronounced, and tenseness. /ʊ/, /ɑ/, and /ʌ/ are clearly "short" vowels, sharing features with the front "short" vowels /ɪ/, /ɛ/, and /æ/. /ə/ is "short" in its unstressed form, but in the /ər/ combination, it resembles the "long" vowels. /ɔ/ can also be relatively long, although it is usually lax, so it doesn't fit neatly into the division between "long" and "short" vowels.

How Two Vowels Can Behave like One

There are three other "long" English vowels, all of them involving significant movement during their production. Vowels of this sort, in which the position of the tongue changes a lot during their production, are called diphthongs. The symbols for diphthongs indicate both the beginning and the end positions. Diphthongs are different from pairs of vowels in that one of the two parts is pronounced more forcefully than the other; for these English diphthongs the more forceful part is the first. The English diphthongs are exemplified in the words *bite*, *bout*, and *boy*; they are represented by the symbols /ay/, /aw/, and /ɔy/. (Note that the symbol at the beginning of the first two has not been used elsewhere for English, though it was used for Spanish and Japanese. It represents a low vowel that is further front than /ɑ/; we'll meet it again later when we discuss English accents.) The symbols [ɪ] and [ʊ] are used for the second members of the diphthongs; these correspond closely to the vowels [ɪ] and [ʊ]. We will meet these symbols again in the section on English consonants. The use of these consonant symbols helps to show that the part of the diphthong that is "stronger" is the first element. The figure below shows the paths that each of the diphthongs traces in vowel space. You can click on the beginning symbol to hear the vowel it represents.



You may be wondering why the diphthongs /ay/, /aw/, and /ɔy/ are considered individual phonemes when each seems to consist of more than one phone. Why can't we just treat them as sequences of two vowel phonemes? The main reason is that the initial elements of these diphthongs behave quite differently from any of the simple vowels. As we'll see later when talk about differences between English accents and phonological change, the pronunciation of these diphthongs in a particular dialect and at a particular point in the history of the language seems to have nothing directly to do with the pronunciation of any of the simple vowels. For example, changes in /æ/ and /ɑ/ over time may be unrelated to changes in /ay/ and /aw/. Thus these diphthongs appear to be inseparable units for English speakers and hearers.

General American has one more diphthong that is *not* usually considered to be an individual phoneme. This is the part of the word *few* following the initial "f" sound. Unlike the three diphthongs discussed above, in this one it is the second element that is "stronger". I will write this diphthong as /yu/. In modern English we can treat /yu/ as two separate phonemes because the second element tends to behave like the vowel /u/ in most dialects; when /u/ changes its pronunciation over time, /yu/ changes as well. The /y/ also appears to be a separate element in that its pattern of occurrence depends on the consonant that precedes it. Consider the following words: *resume*, *lewd*, *news*, *Tuesday*, *few*, *beautiful*, *music*, *cute*. In some accents, all of these words have /yu/; in others, /yu/ is used in the last six; in others (for example, GA), /yu/ is used only in the last four; in others, /yu/ never occurs.

Note that we have already discussed six dimensions for describing vowels: height, backness, length, rounding, tenseness, and whether the vowel height and backness change during the production of the vowel. I'll call this last dimension **diphthongization**. For GA, all we need to specify for this dimension is whether there is movement in the direction of a high front position ([i]), in the direction of a high back position ([u]), or not at all. (There are other possibilities in other English dialects, as we will see later.) Now we can describe all of the English vowels in terms of their values on these dimensions. The table below does this (for all of the dimensions except length). The point here is not to memorize any of this — there is, after all, considerable disagreement about the details — the point is to see how articulatory dimensions (features) allow us to begin to capture what speakers and listeners seem to know about the vowel phonemes of a language.

	/i/	/ɪ/	/e/	/ɛ/	/æ/
Height	hi	hi	hi mid	low mid	low
Backness	front	front	front	front	front
Rounding	spread	spread	spread	spread	spread
Tenseness	tense	lax	tense	lax	lax
Diphthongization	→ i	∅	→ i	∅	∅

	/u/	/ʊ/	/o/	/ɔ/	/ɑ/	/ɒ/	/ə/
Height	hi	hi	hi mid	low mid	low	hi mid	low mid
Backness	back	back	back	back	back	center	center
Rounding	round	round	round	round	∅	∅	∅
Tenseness	tense	lax	tense	lax	lax	lax	
Diphthongization	→ u	∅	→ u	∅	∅	∅	∅

	/ay/	/aw/	/ɔy/
Height	low → hi	low → hi	mid → hi
Backness	center → front	center → back	back → front
Rounding	∅	∅ → round	round → ∅
Tenseness	tense	tense	tense
Diphthongization	→ i	→ u	→ i

Here are examples of each vowel, again as pronounced in my accent. You may notice that the vowels do not sound exactly the same in different words. Since each phoneme is a *category*, we should expect there to be different variants in the same way that not all apples look alike. I'll have more to say about this variation within phonemes in the section on phonetic contexts.

/i/	<i>bead, need, happy</i>
/ɪ/	<i>bid, near, insist</i>
/e/	<i>bait, bay</i>
/ɛ/	<i>bed, head, berry</i>
/æ/	<i>bad, ban, bat</i>
/ʊ/	<i>booed, new, tune</i>
/ʊ/	<i>book, put, poor</i>
/o/	<i>boat, wrote, old</i>
/ɔ/	<i>bought, saw, lost, born</i>
/ɑ/	<i>body, father, bar</i>
/ʌ/	<i>bud, none</i>
/ə/	<i>about, sofa, further, listen, convince</i>
/ay/	<i>ride, write, buy</i>
/aw/	<i>bound, how</i>
/ɔy/	<i>boy, boil</i>

This page titled [3.4: Vowels \(Part 2\)](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.5: English Consonants

We have seen how the vowels of languages can be described in terms of values on a small number of dimensions, or equivalently, features. In this section we will see that the same holds true for consonants. Languages vary a great deal with respect to how many vowel and consonant phonemes they have, but all languages seem to have more consonants than vowel phonemes. Not surprisingly, more consonant dimensions than vowel dimensions are contrastive for languages. We'll first look at the dimensions and the values on those dimensions that are relevant for English. In the next section, we'll see how the consonants of some other languages differ from those of English.

Stops

Exercise 3.5.1

With your tongue touching your upper teeth, pronounce a syllable like *dah*. Then do the same thing with the tip of your tongue touching a position a little further back. Continue moving your tongue back a little at a time, and see how many distinguishable consonants you can produce, some like English, some not. Now try the same thing with the body of your tongue (you can put the tip of your tongue behind your lower teeth to keep it out of the way). The furthest forward you can put it is probably close to the position for the consonant at the beginning of the word *church*.


Let's start with the simple consonants that the speaking Lexies developed. Recall that these are produced with a complete closure of the vocal tract, blocking the passage of air. Consonants made by completely closing the vocal tract are called **stops**. As we discussed informally, different consonants can be produced by varying the place where the closure occurs. This consonant dimension is known as **place of articulation**; we will see later that the place of articulation is also relevant when there isn't a complete closure of the vocal tract. In one sense place of articulation is really six different dimensions because it involves the independent movement of six separate parts of the vocal tract. Each of these structures is called an **articulator**. The articulators relevant for place of articulation are the lips, the tongue tip, the tongue body, the tongue root, the pharynx (the region behind and below the oral cavity), and the **glottis** (the gap between the vocal cords). Each of these, except the glottis, is indicated in this figure from the last section.

You Get Different Stops by Making the Contact at Different Places in the Vocal Tract.


Each language makes use of several places of articulation, usually between three and six, to distinguish its consonant phonemes. In this section, we'll only consider those places that are relevant for English.

There are two possible places of articulation involving the lips as articulators. For **bilabial** place of articulation, the lips are brought together (or for non-stops, as we'll see later, close together). The first and last consonants in the word *bib* are bilabial stops. The symbol for this consonant is /b/, so the pronunciation of *bib* is written /bɪb/. A further possible position of the lips is contact between the lips and upper teeth; this is not used for English stops, though it is used for other English consonants. It is discussed below.

With one part or another of the tongue as articulator, there is a continuous range of possible places for contact with the roof of the mouth, beginning with the upper teeth and extending back to the uvula at the back of the mouth. All languages apparently make use of at least two positions within this range. For English stops, two positions are relevant. One of these is contact between the tip of the tongue and the ridge that is just behind the upper teeth, the **alveolar ridge**. This is referred to as **alveolar** place of articulation. It is a feature of the first and last consonants in the word *did*. The symbol for these stops is /d/, so the pronunciation of the word *did* is written /dɪd/. The configuration of the vocal tract for the pronunciation of /d/ is shown in the figure below.



There is a further possibility for contact between the tongue and the roof of the mouth that is used in most languages. The back of the tongue body contacts the region near the back of the roof of the mouth, near the structure called the **velum**, which I'll have more to say about below. This is called **velar** place of articulation. It is a feature of the first and last consonants in the word *gag*. The symbol for these stops is /g/, so the pronunciation of *gag* is written /gæg/. The position of the vocal tract for the pronunciation of /g/ is shown in the figure below.



So English distinguishes the stops /b, d, g/ along the dimension of place of articulation. Other places of articulation are utilized for other English consonants that are not stops and for consonants in other languages.

Now notice what these three stops share. Like all stops, they are produced with complete contact between the articulators. In addition, all of them are accompanied by voicing, that is, vibration of the vocal cords, during the contact. Thus when they appear at the beginning of a word, the voicing starts before the contact is released (and continues through the following vowel), and when they appear at the end of a word, the voicing continues after the contact is made. (Note that the voicing can't start too much before the beginning of a word like *bib* or stop too much after the end of a word like *bib* because the air being passed through the vocal cords can't escape, and pressure builds up quickly behind the point of contact, preventing more air from being expelled from the lungs. You can verify this for yourself by trying to pronounce a long /b/, /d/, or /g/ sound without releasing the contact between the articulators.)

You Get Different Stops by Varying When You Start or Stop Voicing.

But this is not the only possibility for how the voicing and the beginning and end of the contact can be timed. When the stop consonant comes at the beginning of the word, we get a different effect when the voicing begins *after* the release of the contact; listen to the difference between the words *bay* and *pay*. Similarly at the end of the word, the effect is different if the voicing ends before the contact or roughly at the same time as the contact; compare *add* and *at*. The dimension that distinguishes these pairs of words from each other is called **voicing**. For the moment we will consider only two values for this dimension — **voiced** and **voiceless** — but, as we will see later, voicing is actually more complicated than this.

Just as English has voiced stops at the bilabial, alveolar, and velar places of articulation, it also has voiceless stops at these places. The voiceless bilabial stop is illustrated at the beginning and end of the word *pep*. It is symbolized with /p/, so the pronunciation of *pep* is written /pep/. The voiceless alveolar stop is illustrated at the beginning and end of the word *tot*. It is symbolized with /t/, so the pronunciation of *tot* is written /tat/. The voiceless velar stop is illustrated at the beginning and end of the word *kick*. It is symbolized with /k/, so the pronunciation of *kick* is written /kɪk/. There is also a voiceless stop with its place of articulation at the glottis; this is referred to as **glottal** place of articulation. For a **glottal stop**, the vocal cords are brought together, blocking the airstream as for other stops, and then released suddenly. The glottal stop may appear at the beginning of English words that begin with a vowel, and it appears in the middle of the word *uh-oh*. It is not normally considered an English phoneme, however, because it is not used to create new English words or to distinguish English words from one another. The symbol for a glottal stop in this book is /ʔ/ (the IPA symbol is like a question mark with no dot at the bottom), so the pronunciation of *uh-oh* is written /ʔ^ʔo/.

Fricatives and Affricates

Exercise 3.5.2

Produce the syllables /ba/, /da/, /ga/. Now do the same thing but instead of making a complete closure for the stops at the beginnings of the syllables, leave a little gap between the articulators, and see what consonants result. Do the same thing for /pa/, /ta/, /ka/. (Note that the resulting consonants should sound like English only for /da/ and /ta/.)

So far all of the consonants we have looked at have involved a complete closure of the vocal tract, blocking air from passing out. But this is not the only way to make consonants. In fact we need a new dimension for the various possibilities (really a whole

cluster of dimensions); this is called **manner of articulation**. One crucial variable within manner of articulation is the distance between the articulators. For stops, the closure is complete, but there are two further possibilities. One, discussed in this subsection, involves a narrow, but not complete, closure that allows air to pass through the aperture but with accompanying noise. The other, discussed in the next subsection, involves an opening that is wide enough for the air to pass through unimpeded.

Stops and Fricatives are Different Manners of Articulation.

Consider what happens when you bite your lower lip with your upper teeth and then blow air out. Unless you're biting too hard, some of the air can pass between your teeth and lip, creating a sound like that at the beginning and end of the word *fife*. A phone made like this, with an incomplete or approximate closure that permits air to pass through and produces a noisy sound due to the resulting turbulence, is called a **fricative**. The fricative at the beginning and end of the word *fife* is voiceless because the fricative sound is not accompanied by voicing. That is, the voicing starts after the vocal tract is opened up for the vowel and stops just before the closure made again at the end of the word. The place of articulation for this consonant is one we didn't encounter for English stops; it is the second of the possible places associated with the lips (in addition to bilabial place of articulation). It is called **labiodental** place of articulation. The symbol for the voiceless labiodental fricative is /f/, so the pronunciation of the word *fife* is written /fayf/. English also has a phoneme that is the same as /f/, but voiced. This is the sound at the beginning and end of the word *verve*. It is symbolized by /v/, so the pronunciation of *verve* is written /vərv/.

English has a pair of fricatives at another place of articulation where there are no English stops. Try putting your tongue between your teeth or against the back of your upper teeth and then expelling air from your mouth. Again if the contact is not too tight, some air should pass between your tongue and your teeth, generating turbulence that results in the consonant that appears at the beginning of the word *thing* and at the end of the word *both*. The place of articulation for this consonant is called **dental** place of articulation. The symbol for the voiceless dental fricative is /θ/, so the pronunciation of the word *both* is written /boθ/. English also has the corresponding voiced phoneme; it is the initial consonant in the word *this* and the final consonant in the word *bathe*. It is symbolized by /ð/ in this book, so the pronunciation of *bathe* is written /beð/.

Although English has no bilabial fricatives, it does have alveolar fricatives. When the tongue is allowed to approach but not quite come in contact with the alveolar ridge, we get the consonants in the word *sauce* if it is not accompanied by voice and the consonants in the word *zoos* if it is accompanied by voice. The symbols for these alveolar fricatives are /s/ and /z/, so the pronunciation of *sauce* is written /sɔs/, and the pronunciation of *zoos* is written /zuz/.

Most English Fricatives are Produced at Different Places of Articulation than English Stops.

Somewhat behind the alveolar ridge, it is possible to bring part of the body of the tongue near the roof of the mouth and produce voiceless and voiced fricatives that are distinguishable from /s/ and /z/. The voiceless fricative appears at the beginning and end of the word *shush*. It is symbolized by /ʃ/, so the pronunciation of *shush* is written /ʃʌʃ/. The voiced fricative at this place of articulation is a somewhat marginal phoneme in English, and it does not normally appear at the beginnings of words. It is the consonant in the middle of the word *Asia*. The symbol for this consonant in this book is /ʒ/ (somewhat like the IPA symbol), so the pronunciation of *Asia* is written /eʒə/. /ʃ/ and /ʒ/ are produced at what is called the **postalveolar** place of articulation.

English does not have velar fricatives, but it does have a voiceless glottal fricative, produced by making the glottis narrow enough for a breathy sound to be created. This is the consonant at the beginning of the word *hot*; this phoneme does not occur at the end of English words. It is symbolized by /h/, so the pronunciation of *hot* is written /hət/.

We have seen that stops involve complete closure, and fricatives involve approximate closure. It is also possible to combine these two by beginning with a complete (stop) closure and ending with an approximate (fricative) closure. Such phones are called **affricates**. English has two of them, voiced and voiceless affricates produced at the postalveolar place of articulation. The voiceless postalveolar affricate is the first and last consonant in the word *church*; it is symbolized by /tʃ/ in this book, so the pronunciation of *church* is written /tʃɜrʃ/. The voiced postalveolar affricate is the first and last consonant in the word *judge*; it is symbolized by /dʒ/ in this book, so the pronunciation of *judge* is written /dʒʌdʒ/. Notice how /c/ is similar to a /t/ followed by a /ʃ/ and how /j/ is similar to a /d/ followed by a /ʒ/; in fact, an alternate way to write these affricates is /tʃ/ and /dʒ/.

Sonorants


Exercise 3.5.3

Produce the syllables /ba/, /da/, /ga/. Now do the same thing but instead of making a complete closure for the stops at the beginnings of the syllables, leave a little gap between the articulators, and see what consonants result. Do the same thing for /pa/, /ta/, /ka/. (Note that the resulting consonants should sound like English only for /da/ and /ta/.)

Pronounce the syllable /ba/ while holding your nose. Now try the same thing, replacing the /b/ with an "m" sound (as in *mama*). What can you conclude about the difference between /b/ and the sound of "m"? Do the same thing with /d/ and the sound of "n". Also try to pronounce the word *sing* while holding your nose, and notice what happens to the final consonant (written with the letter combination "ng").

English Nasal and Lateral Consonants Have a Complete Stop-like Contact in One Place, But the Air Escapes Somewhere Else.

One way to produce a sonorant consonant is to completely close the oral cavity, just as for a stop, but to open up the **nasal cavity**, the empty region behind the nostrils and above the oral cavity. This is achieved by lowering the **velum**, the flap at the back of the roof of the mouth. The nasal cavity and velum are shown in the figure below, in which the vocal tract is configured for the production of the sound at the beginning and end of the word *mom*. Such phones are called **nasal consonants** (as we will see later on, there are also nasal vowels). For nasal consonants, the air is allowed to pass through the nasal cavity, but it also resonates in the oral cavity, and the place of articulation (within the oral cavity) distinguishes different nasal consonants from one another. English has three nasal consonant phonemes, at the bilabial, alveolar, and velar places of articulation. The bilabial nasal is the one at the beginning and end of the word *mom*; it is symbolized by /m/, so the pronunciation of *mom* is written /mɒm/. The alveolar nasal is the one at the beginning and end of the word *none*; it is symbolized by /n/, so the pronunciation of *none* is written /nʌn/. The velar nasal is the one at the end of the word *sing* (this phoneme does not appear at the beginning of words in English); it is symbolized in this book by /ŋ/, which is close to the IPA symbol, so the pronunciation of *sing* is written /sɪŋ/.



The other way to produce a sonorant is to leave an opening in the oral cavity that is wide enough so that there is none of the noise that characterizes fricatives. These consonants are called **approximants** because the approach of the articulators is only approximate. One way to achieve this is to make a complete contact as for a stop consonant but release the air at one or both sides of the tongue. Such a sound is called a **lateral approximant**. English has one lateral approximant phoneme, with the contact at the alveolar place of articulation. This is the sound at the beginning and end of the word *lull*; it is symbolized with /l/, so the pronunciation of the word *lull* is written /lʌl/.

A further possibility is for no closure of the oral cavity at all. English has three such consonants. One is produced with the tip of the tongue curled somewhat back and approaching the roof of the mouth behind the alveolar ridge. This is the sound at the beginning and end of the word *rear* (as pronounced by most North Americans). I will use the symbol /r/ for this consonant (though the more precise IPA symbol is [ɹ]; [r] is often used for an alveolar trill or tap, described in the next section). Sounds produced with the tongue curved in this way are called **retroflex**; we can treat this as a particular place of articulation (though not everyone does).

The two other approximants are produced similarly to high vowels, except that the articulators are usually not brought as close together as they would be for vowels. One of these consonants approaches the vowels /u/ and /ʊ/. It is the consonant found at the beginning of the word *we* and is symbolized by /w/, so the pronunciation of *we* is written /wi/. Note that this phoneme has two simultaneous places of articulation: velar, because the back of the tongue approaches the velar region, just as for /u/, and bilabial, because the lips are rounded and brought close together, as for /u/. The other English approximant resembles the vowels /i/ and /ɪ/. It is the consonant found at the beginning of the word *you* and is symbolized by /j/, so the pronunciation of *you* is written /yu/. (Note that in IPA, this consonant is symbolized by /j/.) The place of articulation for this consonant is one we haven't seen yet for

any consonants; as for the vowel /i/, the middle of the tongue approaches the region in the middle of the roof of the mouth. This is referred to as the **palatal** place of articulation.

It should be clear by now that the distinction between vowels and consonants is really a matter of degree. The consonants that are the least vowel-like are stops, which involve a complete closure of the vocal tract and cannot be pronounced continuously. A little more like vowels are fricatives, which can be pronounced continuously but which still have the characteristic fricative noise resulting from the narrow opening in the vocal tract. Closest to vowels are sonorants. All of these can be pronounced continuously. In fact, the English sonorants /m/, /n/, /ŋ/, /r/, and /l/ can all be pronounced as separate syllables by themselves, in which case they behave something like vowels. We have already seen how the combination /ər/, as in *burn*, is sometimes treated as a separate English vowel phoneme. We can treat it as a retroflex vowel or a syllabic retroflex consonant; it really doesn't matter for our purposes. We find the syllabic variant of /n/ in the usual pronunciation of the word *button*; unless the word is pronounced very carefully, there is no /ə/ between the /t/ and the /n/. The same is true for the /l/ in *saddle*. To indicate that a sonorant is syllabic, a short vertical line is added under the consonant symbol, as in [sæd̩l̩].

Not All Phones are Unambiguously Vowels or Consonants.

The phones /w/ and /y/ are the closest of all to vowels. Each has a place and manner of articulation very similar to a high vowel. The main difference between these phones and "true vowels" is that, unlike the other sonorants, they cannot be syllabic; they always require a vowel before or after them to create a syllable. So possible English syllables include [kway], [pyus], and [yaw], but not [kw], [py], or [y]. Because of their similarity to "true vowels", [w] and [y] are sometimes called **semivowels**. For our purposes (and maybe for anyone's purposes) there will be no point to arguing about whether semivowels are vowels or consonants. When I am emphasizing their consonant properties, I will speak of them as consonants; when I am emphasizing their vowel properties, I will speak of them as vowels.

The tables below are a summary of the English consonants in terms of three dimensions, place of articulation (POA), manner of articulation (MOA), and voicing, but recall that place and manner of articulation are more precisely viewed as multiple dimensions.

	/p/	/b/	/m/	/f/	/v/	/θ/	/ð/
POA	bil	bil	bil	lbdn	lbdn	dent	dent
MOA	stop	stop	nas	fric	fric	fric	fric
Voice	vcls	vcd	vcd	vcls	vcd	vcls	vcd

	/t/	/d/	/s/	/z/	/n/	/l/
POA	alv	alv	alv	alv	alv	alv
MOA	stop	stop	fric	fric	nas	lat
Voice	vcls	vcd	vcls	vcd	vcd	vcd

	/ç/	/j/	/ʃ/	/ʒ/	/r/	/y/
POA	alvpl	alvpl	alvpl	alvpl	rflex	pal
MOA	affr	affr	fric	fric	approx	approx
Voice	vcls	vcd	vcls	vcd	vcd	vcd

	/k/	/g/	/ŋ/	/w/	/h/	(/ʔ/)
POA	vel	vel	vel	vel+bil	glot	glot
MOA	stop	stop	nas	approx	fric	stop
Voice	vcls	vcd	vcd	vcd	vcls	vcls

The table below gives examples for each English consonant phoneme. As with vowels, each phoneme may vary somewhat from word to word.

/p/	<i>pin, spin, lap</i>
/b/	<i>bin, lab</i>
/m/	<i>man, ham</i>
/f/	<i>fin, if</i>
/v/	<i>vine, live</i>
/θ/	<i>thin, both, ether</i>
/ð/	<i>this, bathe, either</i>
/t/	<i>talked, stone, lit,</i>
/d/	<i>den, lid, hugged</i>
/s/	<i>sin, kiss, lips</i>
/z/	<i>zoo, easy, lose, eggs</i>
/n/	<i>pin, manner, listen</i>
/l/	<i>lip, sell, castle</i>
/tʃ/	<i>church, nature</i>
/dʒ/	<i>gene, jar, gradual</i>
/ʃ/	<i>shin, mission, nation, fish, machine</i>
/ʒ/	<i>leisure, garage (for some speakers)</i>
/r/	<i>rip, narrow, year</i>
/j/	<i>year, cute /kyut/</i>
/k/	<i>kin, call, lick, chemical</i>
/g/	<i>get, anger, leg</i>
/ŋ/	<i>sing, anger, anchor</i>
/w/	<i>witch, which, reward</i>
/h/	<i>hip</i>
/ʔ/	<i>uh-oh /ʔ^ʔo/</i>

In this section we have seen how the consonants of one language, English, can be described in terms of three dimensions. These three dimensions do most of the work in distinguishing the consonants of the world's languages, but, as we will see in the next section, others are also required. Furthermore, as you can see from the tables above, there are gaps in English, combinations of

features on the three dimensions for which there is no English phoneme. For example, English has no velar fricative, either voiced or voiceless. To some extent, this is just an accident of the history of English, and other languages will fill those gaps (and have gaps of their own elsewhere). The next section gives examples.

This page titled [3.5: English Consonants](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.6: Consonants in Other Languages

Listen to this phrase in Amharic, meaning 'he finished carefully'. See what consonant sounds you can pick out that do not occur in English.

▶ 0:00 / 0:01 ———— 🔊 ⋮

Recall how vowel phonemes in different languages differ from each other. One possibility is that one of the vowel dimensions may be organized differently. For example, the backness dimension has two contrastive values (front and back) in Spanish and Japanese but three contrastive values (front, central, and back) in English (and Amharic). A second possibility is that there is a gap in one system that is filled in the other. For example, Spanish and Japanese have low vowels and front vowels, but they have no low front vowel (/æ/), whereas English does have such a vowel. A third possibility is that one language may use a dimension contrastively which is not used contrastively at all in other languages. Thus the dimension of tenseness distinguishes English vowels from one another (/i/ from /ɪ/ and /u/ from /ʊ/), while this dimension is irrelevant for Spanish and Japanese vowels.

Organization of dimensions

As with vowels, a language may make fewer distinctions on a given dimension than other languages make. Consider Lingala, which, like English, has bilabial, alveolar, and velar stops, nasals, and fricatives, but, unlike English, Spanish, Japanese, Amharic, and Tzeltal, makes no use of the postalveolar place of articulation. That is, Lingala has no phonemes like /č/, /ǰ/, /š/, and /ž/.


Another possibility is that two languages make the same number of distinctions along a dimension but not the same distinctions. Consider place of articulation for stops, affricates, and nasals. English stops, affricates, and nasals (other than the marginal glottal stop) appear at four places of articulation: bilabial (/p/, /b/, /m/), alveolar (/t/, /d/, /n/), postalveolar (/č/, /ǰ/), and velar (/k/, /g/, /ŋ/). Spanish and Japanese also have stops and affricates at four different positions, and three of these are roughly the same as for English, but alveolar is replaced by dental place of articulation, that is, with the tongue tip against the upper teeth rather than against the alveolar ridge. (Recall that English has fricatives at this place of articulation (/θ/, /ð/), but no stops or nasals.) When we are concerned only about the phonemes within a language, we can use the same symbols that we use for the English alveolar phonemes — /t/, /d/, /n/ — for the dental phonemes in these languages because it is not to make sure each phone is kept distinct from every other. However, when it is important to make it clear that the place of articulation is dental rather than alveolar, I will use the IPA symbols [t̪], [d̪], [n̪]. See if you can hear the difference between the alveolar and dental places of articulation in the following syllables: [tata], [t̪at̪a], [dada], [d̪ad̪a].

Now consider voicing. Recall that English consonants are either voiced, with voicing during the production of the consonant, or voiceless, with voicing beginning after or ending before (or simultaneously with) the consonant. Spanish also has voiced and voiceless consonants, but it differs in the details. Listen to your pronunciation of the word *pie*. The lips are brought together for the /p/. Next the lips are opened with a kind of explosive puff of air (which you can feel if you put your hand in front of your mouth). Then the vocal cords begin to vibrate and the vowel /ay/ is produced. Now listen to the pronunciation of the [Spanish word](#) *pai*, a word borrowed from English with the meaning 'pie': /pai/. The consonant at the beginning sounds something like English /p/, but the release of the lip closure and the beginning of voicing happen almost simultaneously for Spanish, and there is no puff of air. We call the English voiceless stop in *pie* [aspirated](#), and when we need to distinguish it from consonants like those in Spanish, we use [h] following the consonant symbol, for example, [p^h]. So if we want to show the detailed pronunciations of the English and Spanish words, we would write them [p^hay] and [pay] respectively.

As in English, the Spanish /p/ is distinct from its voiced counterpart, /b/, as in the word *vaya*, pronounced /baya/. That is, Spanish and English both make a two-way distinction in voicing. The pattern that holds for /p/ and /b/ in the two languages also holds for the other stops and affricates. So English /t/ is aspirated, while Spanish /t/ is not; English /k/ is aspirated, while Spanish /k/ is not. In brief, then, both English and Spanish have voiced and voiceless stops, but for Spanish voiceless stops there is no lag between the release and the voicing as there is for English voiceless stops.

We can see that the voicing dimension is really a continuous dimension with many different possibilities for the relative timing of the release of the consonant closure and the voicing. When we think of voicing in this way, the dimension is sometimes called

voice onset time. Voice onset time is illustrated for three bilabial stops in the figure below. The top line shows the closure (single line) and opening (double line) of the lips. Each of the three other lines shows when voicing begins relative to the opening of the lips (the dashed vertical line) for three different stops, [p^h], as in English *pie*; [p], as in Spanish *pie*; and [b], as in English *buy* and Spanish *vaya*.



The figure shows only three of many possible times for voicing to begin. But both English and Spanish have exactly two categories along this continuum. This means that English and Spanish hearers perceive *discontinuity* where there is continuity. An English hearer would perceive some of the cases as /p/ and some as /b/. The differences between the different /p/ cases and the differences between the different /b/ cases might not be perceived at all; the /p/s and the /b/s would tend to sound the same. At the same time, the differences between the /p/s and the /b/s would be exaggerated; they would tend to sound more different than they actually are. This phenomenon is referred to as **categorical perception**. Both Spanish and English hearers experience categorical perception for voice onset time, but the line dividing their categories is in different places. As we will see in the section on phonetic contexts, the situation in English is somewhat more complicated than what we've seen so far; English speakers actually produce a range of voiceless stops that include ones like the Spanish voiceless stops.

You should not be surprised to know that other languages divide up the voice onset time dimension differently from English or Spanish. In Mandarin Chinese there is also a two-way distinction for stops (and affricates), but the distinction is between voiceless aspirated stops (like English [p^h]) and voiceless unaspirated stops (like Spanish [p]). In problems later on, we will see how other languages treat voice onset time.

What about manner of articulation, the third major dimension distinguishing consonants? While all languages make use of different manners of articulation, some make use of more possibilities. Recall that to some extent, manner of articulation can be seen as a collection of possible ways of configuring the vocal tract to produce sounds. The possibilities we have seen are stops, fricatives, affricates, nasals, and approximants, including lateral approximants.

There are two other possibilities that involve bringing the articulators into contact, as for stops and affricates. For stops and affricates the articulators are brought together and held there till the release. A different approach is for one articulator to quickly tap against the other but not remain in contact with it. This is known as a **tap**. This is the manner of articulation used for the second consonant in the Spanish word *pero* 'but'. For this consonant, the tip of the tongue strikes the alveolar ridge quickly but does not remain there. A similar sound is also used in Japanese and Amharic. When we need to distinguish this tap from other *r*-like sounds, I will use the symbol [ɾ] for this purpose.

The other possibility is for one articulator to be brought quickly in contact with the other several times in succession. This is known as a **trill**. This is the manner of articulation used for the second consonant in the Spanish word *perro* 'dog'. The place of articulation is the same as for [ɾ], but in this case the tongue strikes the alveolar ridge several times. A similar sound is also used in Amharic; it appears in the Amharic phrase referred to in the box at the beginning of this section. The usual IPA symbol for the alveolar trill is [r], but we can also use a double [r] for this purpose to distinguish it from a tap when this is necessary. With this notation, the pronunciation of the Spanish word *perro* is written /perro/.

Gaps

Now consider how some languages fill the gaps of other languages. Notice that in the alveolar place of articulation, English has stops, fricatives, and a nasal. In the bilabial place of articulation, it has stops and a nasal but no fricatives, though there are fricatives in the nearby labiodental place of articulation. In the velar place of articulation, English has stops and a nasal but no fricatives in this or any nearby position. Notice also that while English has bilabial, alveolar, and velar nasals, it has no nasal phoneme in the postalveolar or palatal places of articulation.

These last two gaps are filled in Spanish. Spanish has a voiceless velar fricative, the consonant in the middle of the word *México*. This sound is symbolized by /x/, so the pronunciation of the word *México* is written /mexiko/. Spanish also has a palatal nasal, the consonant in the middle of the word *año* 'year'. This consonant is symbolized in this book by /ɲ/, so the pronunciation of the word *año* is written /aɲo/. But Spanish has gaps of its own. There is no velar nasal phoneme, a gap which is filled by the English phoneme /ŋ/. And while Spanish has a postalveolar affricate (/tʃ/), it has no postalveolar fricative (/ʃ/), unlike English, Japanese, Amharic, and Tzeltal.

New dimensions

Just as for vowels, one language may make use of a consonant dimension which other languages do not. Two of the languages on our list, Amharic and Tzeltal, have an alternate way of producing stops and affricates. This mode of production involves a buildup of pressure behind the point of contact and an explosive release accompanied by a glottal stop; such consonants are referred to as **ejective** consonants. I will use a following apostrophe to indicate these sounds and will refer to this dimension as **glottalization**. (Note that we have to consider this a new dimension rather than just a new value for manner of articulation because it is possible with *different* manners of articulation; there are glottalized stops, fricatives, affricates, and sonorants in different languages.)

Listen to the following contrasts, and see if you can hear the difference between ejectives and plain, non-glottalized stops and affricates: [papa], [p'ap'a], [tata], [t'at'a], [kaka], [k'ak'a], [čača], [č'ač'a]. Glottalization is used contrastively in Amharic and Tzeltal. That is, the sounds /t/, /č/, and /k/ are used, like /t/, /č/, and /k/, to make distinct words in these languages. Three of these sounds appear in the Amharic phrase mentioned in the box at the beginning of this section, which it may be worth transcribing now: /bet'InIkk'ak'e c'errese/. To show that glottalization is contrastive in Amharic, we can cite pairs like /kok/ 'peach' vs. /k'ok/ 'partridge' and /tɪl/ 'worm' vs. /t'ɪl/ 'quarrel'.

This section has not been a complete survey of possible consonants in human languages, or even in the ten languages discussed in this book. There are other places of articulation, other manners of articulation, and even other dimensions. The point has been to show how consonant systems differ along the basic dimensions of place of articulation and manner of articulation and what all languages share: a small set of phonemes produced with different kinds of constrictions within the vocal tract. In the next section we'll see how these consonant phonemes are combined with the vowel phonemes we discussed earlier to form syllables and how languages resemble and differ from each other in how this is done.

This page titled [3.6: Consonants in Other Languages](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.7: Syllables

We have seen how each spoken language has a set of consonant and vowel categories that are used by its speakers and hearers to distinguish the words of the language. The consonants and vowels in turn are combined into larger units, syllables. Syllables are distinguished from one another in terms of the consonants and vowels that they consist of. But syllables can also be distinguished from one another in other ways, and some of these ways are very commonly used contrastively, that is, to distinguish words from each other. We will look at some of these "suprasegmental" features of language in this section. Languages also differ in terms of how consonants and vowels can be combined into syllables, the "phonotactics" of the language, and we will also look at this property of languages in this section.

Suprasegmentals

Note

Let's go back to our Lexies in an early stage of their word development. They have vowels and consonants, and their word forms consist of one or two syllables. Consider the possible word form /bago/. Without changing the vowels and consonants, how would it be possible to make this pair of syllables into more than one distinguishable word?

We've discussed vowels and consonants, and in the [section on phonemes](#) we looked briefly at how they are combined to form syllables such as /pa/, /bi/, and /ne/. We've also discussed the dimensions that distinguish different vowels from each other and the dimensions that distinguish different consonants from each other. Now we consider the dimensions and features that distinguish syllables from each other, independently of the consonants and vowels in them. Since consonants and vowels are sometimes referred to as "segments", these dimensions and features are referred to as **suprasegmentals**, that is, 'above the segments'.

One property that clearly characterizes a syllable and that could distinguish one syllable from another is **loudness**. A particular one-syllable word could be spoken more loudly than other words, or a two-syllable word could have one syllable spoken more loudly than the other. In this case we would be concerned with **relative** rather than **absolute** loudness. That is, we would only care that the first syllable of a two-syllable word is **louder** than the second, not that the first syllable has a particular loudness.

Another property of syllables is their length (though this may amount to the same thing as **vowel length**). One syllable in a word may be held for a longer time than the other(s). Again what seems to matter for language is relative, rather than absolute, length.

Finally syllables may differ from one another in their **pitch**, that is, the dimension that distinguishes musical notes from one another. Once again, what will we care about is relative pitch; if absolute pitch mattered, as it does in music, women, men, and young children would be unable to achieve the same effects. One syllable can have a higher pitch than another. A syllable can also be characterized by a particular pitch movement, say, rising or falling, rather than a level pitch. Note that movement is a separate dimension from overall relative pitch; a pitch fall could start and end relatively high or relatively low.

The main question that should concern us, because our focus in this chapter is what distinguishes word forms from one another, is whether any of these suprasegmental dimensions is used contrastively. Let's start with English. Consider the two instances of *permit* in the following sentence.

1. Without a **permit**, they wouldn't **permit** me to participate.

Both of these words would be transcribed with the same set of consonant and vowel phonemes: /pərmIt/. But they are pronounced differently. There is more "effort" expended on the first than the second syllable in the first word and on the second than the first syllable in the second word. The actual difference may involve loudness, length, and pitch: the first syllable in the first *permit* is probably louder and longer than the second, and the first syllable probably involves a fall from a relatively high to a low pitch while the second syllable is more or less level and low. The reverse is probably true for the second *permit*.

This suprasegmental dimension of English is called **stress**. Because there are words like the two *permits* in English, we can see that this dimension is contrastive in English. English may have as many as three different values (levels) of stress within a word. I will symbolize them with '/' before a syllable with high ("primary") stress, /, before a syllable with medium ("secondary") stress, and nothing before a syllable with weak stress. Thus the two words we have been discussing would be written /'pərmIt/ and /pə'rmit/, and the word *constitution* would be written /,kɒnstə'tuʃən/. Other English examples in which stress alone distinguishes words are

torment (/ˈtɔrmɛnt/ and /tɔrˈmɛnt/) and *survey* (/ˈsɜrve/ and /sɜrˈve/). Spanish also has contrastive stress. For example, the words *canto* 'I sing' and *cantó* 'he sang' differ only in stress: /ˈkanto/ and /kanˈto/ respectively.

Now let's look at how pitch alone behaves in some languages, for example, Lingala. Consider the following words, in which /ˈ/ over a character indicates a relatively high pitch and no mark over a character indicates a relatively low pitch.

2. /moto/ 'person', /motó/ 'head'
3. /ebóló/ 'piece of cloth', /eboló/ 'skull', /ebólo/ 'group'
4. /moluka/ 'fishing', /molúka/ 'canoe trip', /molúká/ 'river'

Clearly pitch alone is enough to distinguish words in Lingala. That is, pitch is used contrastively in this language. This use of pitch is called **tone**. In a tone language such as Lingala, Mandarin Chinese, or one of the thousands of other tone languages of Africa, Asia, or the Americas, each syllable has an associated tone, that is, a pitch level or movement. Each tone language has a small set of tone categories, or **tonemes**, which are used to distinguish words in the language just as phonemes are (and, as we'll see later, in languages like Lingala also to distinguish grammatical forms). In Lingala the basic tonemes are high and low tone; there are also somewhat marginal rising and falling tones. Note that in a tone language, it is *relative* pitch that matters. "High tone" means high relative to the pitch of the speaker's voice and to the pitch of the rest of the utterance in which the syllable occurs, not a particular pitch or range of pitches.

Japanese also uses pitch alone to distinguish words, but the system works somewhat differently from that in a tone language like Lingala. In Japanese, for words of a given number of syllables there are a small number of possible pitch patterns. Rather than specify the pitch of every syllable for a given word, we just need to specify (and the learner needs to remember) which of the pitch patterns is used for that word. A language like this is called a **pitch accent** language. The following examples illustrate the three possible patterns for a two-syllable noun followed by the word *wa*, which indicates that the meaning of that noun is the "topic" of the sentence. In the transcription, pitch is indicated by the height of the syllables, with low pitch at the level of the line.

5. [haši wa] 'edge TOPIC'
6. [haši wa] 'bridge TOPIC'
7. [haši wa] 'chopsticks TOPIC'

Because these three phrases are distinguished by pitch alone, we can see that pitch is used contrastively in Japanese, as it is in Lingala.

The suprasegmental dimensions of pitch, loudness, and length also play a somewhat different role in languages. Consider the following English sentences, in which the word in boldface is emphasized.

8. *Lois married Clark.*
9. *Lois **married** Clark.*
10. **Lois** married Clark.
11. *Lois married **Clark**?*
12. *Lois **married** Clark?*
13. □ ***Lois** married Clark?*

Notice how suprasegmentals (loudness, length, and pitch) are used to emphasize different words in the sentences and to indicate whether the sentence is a statement or a question. These uses of suprasegmentals are referred to as **intonation**. All human languages appear to use intonation.

Phonotactics

Note

Consider the following made-up words, each written both how it might be spelled in English and with phonetic symbols.

14. *glooce* /glus/
15. *verm* /vərm/
16. *binzle* /bɪnzəl/
17. *fkotch* /fkɑc/

18. *sreep* /srip/
19. *noo* /nʊ/
20. *taheh* /tʌɛ/
21. *lingg* /lɪŋg/

Do all of these seem like possible English words to you? If some don't, what about them seems to be impossible in English?

As we have seen, each spoken language has an "alphabet" of form categories — consonant and vowel phonemes — which are combined to form the syllables that make up words. But languages differ not only in the particular vowel and consonant phonemes they have. They also differ with respect to how the vowels and consonants may be combined to form syllables.

Let's start with simple English syllables consisting of a consonant followed by a vowel; I'll abbreviate this as "CV". First, can any consonant appear in the "C" position? Taking the vowel as the constant /o/, certainly all of the following are possible syllables in English: /po/, /bo/, /mo/, /vo/, /to/, /co/, /ʃo/, /ko/, /lo/, /ro/, /wo/, /ho/. But what about /ŋo/? A complete search of the English lexicon reveals that there are no English words that have syllables beginning with the phoneme /ŋ/. Although other nasal consonants (/m/ and /n/) and other velar consonants (/k/ and /g/) can appear at the beginnings of syllables, English seems to constrain syllables to not begin with the phoneme /ŋ/.

What about the vowels in a CV syllable? Let's be more specific and assume that the syllable is stressed and comes at the end of an English word. Keeping the consonant as the constant /b/, all of the following seem possible: /bi/, /be/, /bu/, /bo/, /bɔ/, /bay/, /baw/, /bɔy/. (For speakers who do not make the distinction between /ɔ/ and /ɑ/, /bɑ/ would also be possible.) But what about the following: /bɪ/, /bɛ/, /bæ/, /bʊ/, /bʌ/, /bɑ/ (for speakers who distinguish /ɑ/ and /ɔ/)? None of these syllables seems possible. Again there is apparently a sort of prohibition on the kinds of phonemes that can appear in English syllables. In this case, the most efficient way to state the prohibition is to say that *English forbids lax vowels, other than /ɔ/, from appearing at the ends of syllables* (at least stressed syllables at the end of words). Note that /ɔ/ presents a problem for the generalization; this is one of the ways in which this vowel does not quite fit into the lax/tense, short/long distinction.

Thus English has constraints on the structure of syllables. Such constraints are referred to as **phonotactics**. It's beyond our goals to go into English phonotactics in detail, but let's investigate a bit further what the bounds are on English syllables.

What about syllables with more than one consonant at the beginning? In general, clusters of consonants not separated by vowels are more difficult for speakers to produce than consonants that are separated by vowels. This is because the articulators must move from one consonant position to another without opening up in between (because the opening would be realized as a vowel). And the difficulty of particular combinations varies considerably. Thus we should expect more constraints on what is possible in clusters than for single consonants. An examination of the English lexicon reveals that the following consonant clusters can appear at the beginnings of General American English syllables (my accent) if we count the semivowels /w/ and /y/ as consonants.

22. /tw/, /dw/, /kw/, /gw/
23. /by/, /py/, /my/, /fy/, /vy/, /ky/, /hy/
24. /pl/, /bl/, /fl/, /kl/, /gl/, /sl/, /ʃl/
25. /pr/, /br/, /fr/, /θr/, /tr/, /dr/, /kr/, /gr/
26. /sp/, /st/, /sk/, /sm/, /sn/, /ʃp/
27. /spl/, /spr/, /str/, /skl/, /skr/

We can see some patterns in what is possible. /s/ seems to be special. If we leave it out, we see that all of the clusters end in a sonorant consonant, /w/, /y/, /l/, or /r/. Clusters of three consonants must consist of /s/ followed by a voiceless stop followed by either /l/ or /r/. In fact, for this and other reasons, /l/ and /r/ are often treated as forming a category in their own right.

English has a range of more detailed constraints when it comes to which vowels can occur before which consonants. Consider syllables ending in a vowel and a consonant. Some syllable-final consonants, for example, /t/, permit any English vowel before them. But before /r/, the possibilities are quite restricted. In my accent, only the following vowels are possible before /r/: /ɪ/ (*pier*), /ɛ/ (*pair*), /ʊ/ (*poor*), /ɔ/ (*pour*), /ə/ (*per*). In other words, none of the tense vowels may appear before /r/. Another way to see this is as the **neutralization** of the lax-tense distinction before /r/; that is, the distinction between tense and lax vowels has disappeared in the context before an /r/. Evidence for this is that the vowel in the word *pour* is actually somewhere between the usual /ɔ/ and the

usual /o/ in this accent, and we could actually represent it with either symbol when we are just representing the phonemes of the dialect.

To some extent, the constraints on English syllable clusters seem to be related to what is easy to do. A cluster such as /mk/ or /lpr/, not possible in English, is quite difficult to produce. But the constraints also seem somewhat arbitrary. For example, there is no reason to believe that /ʃk/, which is not possible in English, is any more difficult than /sk/, which is possible. And /bw/, which does not occur, seems no more difficult to produce than /tw/, which does.

If we examine the phonotactics of other languages, we see these same basic properties. In addition, we see that the *degree* of complexity which is permitted for syllable structure is specific to the language (and varies considerably between languages). Let's look at [Japanese](#). A Japanese syllable can begin with at most one consonant; no consonant clusters are permitted. A Japanese syllable can end with a vowel, or /n/, or, if the syllable is not at the end of a word, with the consonant that begins the next syllable (but only if that consonant is /p/, /t/, /k/, /s/, /ʃ/, or /c/). Thus the first group below includes possible Japanese words, and the second group includes impossible Japanese words.

28. /e/, /se/, /te/, /ten/, /kantan/, /henkai/, /nattoo/, /makka/

29. /nat/, /mak/, /bum/, /nas/, /ste/

We can see that Japanese draws the line between what is phonotactically possible and what is not in a very different place than English does.

Spanish fits somewhere between English and Japanese. Spanish permits at most one consonant at the end of a syllable, and this consonant can only be one of the following: /d, s, n, l/. (If you know some Spanish and think that /m/ can appear at the end of a syllable, as in *comprender*, you're right in a way. Later we'll see that this "m" can be seen as a kind of variant of /n/.) Spanish does permit consonant clusters at the beginnings of words but no clusters beginning with /s/. It does permit two-consonant clusters ending in /l/ and /r/, however, much as English does. And, because Spanish has a number of diphthongs beginning with /w/ and /y/, Spanish syllables permit *more* clusters at the beginning of syllables ending in /w/ and /y/ than English does, for example, /fw/ and /sy/. In fact, Spanish permits some three-element clusters consisting of two consonants followed by /w/ or /y/, for example, /prw/, as in the word *prueba* /prweba/ 'test'.

Not surprisingly, there are languages which are *more* extreme than English in terms of the complexity they permit in syllables, though none of these is among the eight other spoken languages discussed in this book. Among these language with more complex syllables are familiar languages like Russian. To take perhaps the most extreme of all, in the Canadian Indian language Nuxalk, a word may consist of as many as four consonants and no vowel, for example, /sk'st/.

We have seen that each language has its own idea about what counts as a good syllable; that is, each language has a syllable *category*. Though we will not have time to go into it, it turns out that syllables are grouped together into higher-level units and that languages also differ in the ways this can be done. Some of the constraints make good sense from the perspective of the Hearer. It turns out that for a hearer, it is easier to distinguish consonants at the beginnings than at the ends of syllables. Thus it should not be surprising that more different consonants are possible at the beginnings than the ends of syllables in many languages; making distinctions that are hard to hear would not serve any function.

What about the constraints that distinguish one language from another? Why do languages seem so different when it comes to phonotactics? These differences seem to be related again to the Hearer-oriented pressure to make syllables more distinctive so that they are easier to distinguish along with the opposing Speaker-oriented pressure to make words easy to pronounce. The more different consonants and consonant clusters are possible at the beginnings of syllables, the more distinct syllables are, but, at the same time, the more difficult syllables become to produce. Different languages have sorted out the conflict in different ways.

So how do languages with relatively constrained syllable structure deal with the need for word forms to be distinct? One strategy is tone; more different syllables are possible if each syllable has an associated tone as well as a sequence of consonants and vowels. Mandarin Chinese is a language with relatively simple phonotactics and hence relatively few possible syllable types, but it compensates this by having four separate tones.

Another strategy is words consisting of more than one syllable. Compare these Japanese and English nouns referring to basic body parts: *atama*, *head*; *kokoro*, *heart*; *karada*, *body*; *ashi*, *leg/foot*; *te*, *hand*; *hana*, *nose*; *mimi*, *ear*; *mune*, *chest*; *koshi*, *hip*. Only one of these Japanese words consists of one syllable, while only one of the English words consists of more than one syllable. Japanese

has fewer distinct syllable types than English, so it compensates by making words out of longer sequences of syllables than English does.

This page titled [3.7: Syllables](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.8: Problems


In transcribing the phones in these problems, you can replace the symbols that are not found on a keyboard with uppercase letters. (These conventions agree mostly with emerging internet conventions for representing IPA in email and newsgroups.)

- $\varepsilon \rightarrow E$
- $\alpha \rightarrow A$
- $\mathfrak{O} \rightarrow O$
- $\theta \rightarrow T$
- $\delta \rightarrow D$
- $\mathfrak{S} \rightarrow S$
- $\mathfrak{Z} \rightarrow Z$
- $\mathfrak{C} \rightarrow C$
- $\mathfrak{J} \rightarrow J$
- $\eta \rightarrow N$

For "æ", use "&".

Vowels

1. The chart below shows the vowel phonemes of Amharic in vowel space. Rounded vowels are circled; spread vowels are underlined. Note that some of the symbols do not refer to exactly the same sounds that they refer to for English.



Describe each vowel in terms of its values on the relevant vowel dimensions. There should be three dimensions.

2. For each of the following English words, write the symbol for the vowel in the word. Do not use a dictionary to help you.

- a. *mine*
- b. *main*
- c. *mean*
- d. *moon*
- e. *moan*
- f. *men*
- g. *man*
- h. *love*
- i. *weird*
- j. *weigh*
- k. *pants*
- l. *soft*
- m. *put*
- n. *owe*
- o. *far*
- p. *loud*
- q. *point*
- r. *head*
- s. *niece*

English consonants

1. For each of the following sets of English phonemes, say which features (dimension values) they share. For item j, you will have to come up with a feature that is not mentioned in the book; think of what the articulation of these phones shares, and describe it in a short sentence.

Example: /p, b/ — bilabial, stop

- a. /ɛ, æ, ɪ, i, e/
 - b. /ʊ, ɪ, u, ɪ/
 - c. /z, d, n, t, s/
 - d. /w, y/
 - e. /n, ŋ, m/
 - f. /v, z, ʒ/
 - g. /g, ŋ/
 - h. /s, z/
 - i. /t, k, p/
 - j. /p, m, f, b, v/
2. Transcribe each of the following English words, using the phonetic symbols in the book. Do not use a dictionary to help you.
 - a. *creature*
 - b. *anger*
 - c. *hanger*
 - d. *luxury*
 - e. *pushed*
 - f. *physics*
 - g. *badges*
 - h. *thistle*
 - i. *rhythm*
 - j. *choir*

Syllables

1. For each of the following, say whether it is a possible word in English (whether it obeys English phonotactics). If it is not possible, say why in the most general terms possible.

Example: /læbg/ No, because English syllables cannot end with a /bg/ cluster.

- a. /psu/
- b. /plʊ/
- c. /nawz/
- d. /wrayn/
- e. /brith/
- f. /bɪŋf/
- g. /sprʊks/
- h. /'kræbəl/
- i. /sædb/
- j. /'pyuzi/

This page titled [3.8: Problems](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

4: Word Forms - Processes

- [4.0: Introduction to Word Forms](#)
- [4.1: Phonetic Contexts](#)
- [4.2: Assimilation](#)
- [4.3: Distribution of Phones](#)
- [4.4: Learning Phonology](#)
- [4.5: English Accents](#)
- [4.6: Phonological Change](#)
- [4.7: Phonology in the Wild](#)
- [4.8: Problems](#)

This page titled [4: Word Forms - Processes](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.0: Introduction to Word Forms

In the last chapter, we looked at the basic units that make up the forms of words in spoken language: syllables and the consonants and vowels that combine to form syllables. In this chapter, we will look at various ways in which these units may change. First, a given phoneme may be pronounced differently depending on the phonemes that immediately precede and follow it. Though processes like this seem to originate in making it easier for Speakers to produce sequences of phones, they become conventional. That is, different languages make use of different processes of this sort. These kinds of processes also depend on the formality of language; in casual speech, there are often additional simplifications to the forms. Second, the units of linguistic form obviously change through the course of language learning. A beginning first or second language learner does a poor job of producing and recognizing the units of the target language but gets closer to the capability of a native Speaker/Hearer as learning progresses. Finally, the units of linguistic form change more slowly throughout the entire community of Speakers/Hearers. That is, the phonological conventions that define the forms of any language are not constant. In this chapter, we will also review much of what we've covered concerning linguistic form in the context of different English accents.

This page titled [4.0: Introduction to Word Forms](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.1: Phonetic Contexts

Lexical and Phonological Knowledge

Exercise 4.1.1

I've been claiming that the sounds of a human language are clustered into categories called phonemes. As a speaker of English, then, what do you have to know to be able to pronounce an English word, say, the word *fun*? Hint: some of what you know is about just that word and some of what you know is more general.

The main point of what I said about the sounds of spoken language in the last chapter was to show what phonemes are. Let's review this concept.

1. They are specific to particular spoken languages (or dialects actually); that is, learning a language means in part learning the phonemes of the language.
2. They are categories, with all of the familiar properties that categories have.
 - a. Phonemes are realized as individual sounds (phones), instances of the categories, that may differ from one another but are centered on a prototypical member. That is, a phoneme is not really a phone, but an abstraction over a collection of phones.
 - b. Phonemes divide a continuous space (one or more continuous dimensions) into a small set of types.
 - c. Individual phones are categorized by hearers as belonging to one or another of the phonemes of the hearers' language. Phones that are closer to the prototypical member are easier to categorize. Phones that are within a phoneme are perceived as closer to each other than phones that belong to different phonemes, even when the real difference is the same (categorical perception).
3. Phonemes provide a kind of "alphabet" in terms of which speakers and hearers remember the pronunciations of words.

Let's focus on the last point for now. I tried to argue in the section on phonemes that phonemes provide a more efficient way of remembering a large number of words than the alternative. How might this memory system work? We'll concentrate on what the Speaker has to remember. Let's take the English word *fun*, which we assume is stored in the Speaker's (and the Hearer's) memory as the sequence of phonemes that we have been writing as /fʌn/. (Of course the claim is not that anything like these symbols is written in the person's brain, only that there is some representation of the three phoneme categories and their order.)

Remembering How to Pronounce an English Word Involves Specific Knowledge About that Word and General Knowledge About the Pronunciation of English.


Now how would the Speaker use this representation in memory to actually pronounce the word *fun*? Another sort of knowledge is needed: knowledge about how each phoneme is realized in terms of articulation. So for /f/ the Speaker would need to know to bring the lower lip in contact with the upper teeth and pass air through the opening without vibration of the vocal cords. And for /ʌ/ the Speaker would need to know to open the mouth a certain extent and raise the body of the tongue somewhat while causing the vocal cords to vibrate. Finally for /n/ the Speaker would need to know to bring the tongue tip in contact with the alveolar ridge, to lower the velum, and again to cause the vocal cords to vibrate.

I'll refer to this knowledge about how phonemes are produced as **realization rules**. This is our first use of the term **rule**, which you'll meet a lot later in this book. It refers to general knowledge about what to do in a particular kind of situation. Notice how this differs from the way the word *rule* is used outside of social science; realization rules are learned but not taught, and speakers are not conscious of the rules they know. And calling the knowledge a "rule" should not make us think that it takes the form of some sort of explicit statement in the brain of the Speaker. In fact little is known about the precise form that rules take; this is currently an area of great controversy in cognitive science.


Just as the Speaker needs to know how to produce each phoneme, the Hearer (who of course is also a Speaker) needs comparable knowledge for how to perceive each phoneme, knowledge about what the phonemes sound like. However, since I have had little to say about the acoustic or auditory properties of phonemes, I'm not in a position to spell out in more detail what this knowledge is.

The two kinds of knowledge, knowledge about the form of a particular word in terms of a sequence of phonemes and knowledge about how particular phonemes are articulated or recognized, differ in one very important way. Knowledge about the form of

particular words must be memorized for each word; none of this is general knowledge. This knowledge belongs in the lexicon, the storehouse of knowledge about particular words. On the other hand, knowledge about how phonemes are produced or perceived is general; it applies to all words containing the phonemes. This knowledge is part of the Speaker's and Hearer's **phonology**, that is, general knowledge about the form that words can take in the language. The figure below illustrates these two types of knowledge.



The fact that phonological knowledge is general means that it applies to other words as well. For example, the word *laugh* /læf/ also contains the phoneme /f/, so the realization rule for /f/ applies to this word as well. And the word *no* /no/ also contains the phoneme /n/, so the realization rule for /n/ applies here too. These examples are illustrated in the figure below. The arrows below the phonemes indicate that each phoneme gets spelled out as a set of articulatory actions and an auditory pattern. The arrows all go in both directions because the knowledge has to be usable by both speakers (down in the figure) and hearers (up in the figure).



The same would hold true for other languages, except that the phonemes would be different (so the knowledge about how they are produced and perceived would be different), and of course the words would also be different.

Realization rules may also specify how combinations of phonemes are pronounced. One confusing area in English is the behavior of the sonorants /m, n, ŋ, l, r/ in unstressed syllables, in words like *prizm* (/m/), *happen* (/n/), *incredible* (/ŋ, l/), *under* (/r/). One possibility is to see these in each case as a vowel, /ə/ or /ɪ/, followed by one of the sonorant consonants. So *prizm* is /'prɪzəm/, and *incredible* is /ɪŋ'krɛdəbəl/. But, as we saw in the section on English sonorants, the unstressed vowel is sometimes not pronounced; instead the sonorant is syllabic, behaving almost like a vowel. Sometimes this is the only natural pronunciation; sometimes it depends on the speed and informality of the situation. So a relatively casual pronunciation of *incredible* is [ɪŋ'krɛdəbəl] (recall that a short line under a consonant symbol indicates a syllabic consonant). Since this knowledge about how to *realize* combinations like /ə/ is general knowledge about English, we can put it in the realization rules in the phonological component.

Phonemes in Context

Exercise 4.1.2

Listen to the sound represented by the letter "t" in the word *put* in the following sentences.

1. *Put something on.*
2. *Put me down.*
3. *Put it on the table.*
4. *Put this on.*
5. *Put your shirt on.*

If you listen carefully, you may hear as many as five different "t" sounds. What does this mean? Should we assume that English speakers have five different phonemes in place of the single /t/ that we discussed in the section on English consonants? Keep in mind what phonemes are and what function they serve in language.

But the picture is not this simple. Remember point 2a about phonemes above: the actual instances of a given phoneme will differ from one another. So one /f/ will not necessarily be the same as another /f/. But are the differences just random? Let's consider /t/,

where the variations are quite striking. Take the word *at*, which in terms of the phonemes that have been proposed would be represented in the lexicon as /æt/. Given what we've said so far about English /t/, the realization rule for this phoneme would have to include specifications that the speaker place the tip of the tongue against the alveolar ridge and produce the sound without vibration of the vocal cords. (The specification is intentionally somewhat vague about how the consonant is released so that it can apply to /t/ at both the beginnings and the ends of syllables.)

English /t/ has a Number of Different Pronunciations

This works fine for the /t/ in *at* when the word appears before a word beginning with /f/ (*at four o'clock*) or /s/ (*at six o'clock*), say. But when it appears before a word beginning with /θ/ (*at three o'clock*), and we say it in a natural way, we see that the tongue tip is not against the alveolar ridge, but against the teeth, as for a dental stop ([t̪]). And when it appears before a vowel (*at eight o'clock*), the consonant, as pronounced by North Americans anyway, is voiced and articulated as a tap rather than a stop. In fact it is very similar to the Spanish /ɾ/, that is, the phone that is more accurately represented in our notation by [ɾ]. So we see that not only does the final consonant in *at* take different forms, there is a regularity to the forms it takes; the pronunciation depends on what phoneme follows that consonant.

How are we to deal with this kind of variability in our theory of how the pronunciation of words is represented? One possibility would be multiple representations of the word *at* in the lexicon: /æt/, /æt̪/, /æɾ/. Since we would now be using /t̪/ and /ɾ/ to represent words in the lexicon, these would have to be seen as English phonemes, in addition to /t/. But note that this wouldn't be enough; a list of different pronunciations of the word would have to say something about when each one was appropriate, for example, *use /æɾ/ when the next word begins with a vowel*. And of course the phonological component of memory would need to specify how /t̪/ and /ɾ/ are produced and perceived, as well as /t/.

But this would be a strange way for language to work. Why should speakers be forced to remember more than one form for a word? The whole point of phonemes (at least that is what I'm trying to argue) is to make words easy to remember. Recall also that the lexicon is supposed to be for knowledge about specific words, *not* general knowledge about how words are pronounced. If the information about different ways to pronounce *at* applies to more than just the pronunciation of *at*, then it does not belong in the lexicon; it belongs in the phonological component. If we examine a lot of English words, we discover that many of them (for example, *put* and *let*, as illustrated in the box above) end in consonants that behave exactly like the consonant at the end of *at*. In each case the different pronunciations of that consonant depend on the first phoneme in the next word.

The English Phonological Component Tells How to Pronounce /t/ in Different Contexts

So now consider a second alternative. The lexicon records only one pronunciation of *at*, /æt/, and the phonological component specifies how the /t/ is to be pronounced (or perceived). The difference from the simple picture described in the last subsection is that this specification has to refer to what comes after the /t/. So it would include something like the following: *if /t/ is followed by a dental fricative, place the tongue against the upper teeth*, and a corresponding statement for each of the other variant pronunciations. Now the realization rule for /t/ is really a set of rules, one for each type of following phone and the corresponding realization of /t/. The different realizations of a phoneme are called its **allophones**. We'll come back later to a more detailed statement of what the realization rules and allophones for English /t/ are. For now, there are these important points to note.

- An allophone is always an allophone *of* some phoneme; it makes no sense to say simply that a particular phone is an allophone in a language. So in English we've seen that [t], [t̪], and [ɾ] are all allophones of the English phoneme /t/.
- This sort of knowledge about how a phoneme is realized as different allophones is general knowledge that speakers have about English, not about particular English words. It is in fact part of the knowledge of what the different English phonemes are.
- The different ways in which a phoneme such as /t/ is pronounced, that is, the allophones of /t/, are similar. They differ at most in one or two features from each other. In fact /t/ is a rather extreme example for English; a phoneme such as /f/ varies relatively little.
- The realization of the phoneme pronounced as one or another allophone usually depends on what other sounds are near the phoneme in question, that is, on the phonetic **context** of the phoneme. In the section on assimilation, we'll learn more about how contexts determine allophones, and in the section on distribution, we'll learn about how you can use the different contexts that phones appear in to decide whether they're allophones of a single phoneme or separate phonemes.
- Allophones are *not* categories. This means that speakers and hearers do not need to categorize the different allophones of English /t/. The only category they need among these variants is /t/. It also means that hearers may not even notice the

difference between the different allophones of a phoneme. As far as they are concerned, they all belong to the same category. This fact is often reflected in writing systems, which almost never assign different characters to different allophones of a phoneme. So English uses a single letter "t" for all of the allophones of the phoneme /t/.

Let's look at a Spanish example of the same phenomenon. Consider the word *de* 'of, from'. If we listen to this word spoken in isolation or at the beginning of a sentence, for example in the expression *de nada* 'don't mention it', we would conclude that its form consists of the two Spanish phonemes /d/ and /e/. As in the English example above, the actual pronunciation of the word would require knowledge about how /d/ and /e/ are articulated, that is, realization rules for these phonemes. This is a different language, so the realization rules would not be the same as for the English phonemes that we write with /d/ and /e/. Thus Spanish /d/ is a dental, not an alveolar, stop ([d̪]), but the basic principle still applies: phonological knowledge is general; it applies to many words.

But consider now how the word *de* sounds when it appears in a phrase following a vowel, for example, in *un vaso de leche* 'a glass of milk'. Here the consonant in *de* involves no complete closure between the tongue and upper teeth, so this is not a stop. Some people refer to this sound as a fricative, but since the closure is usually not close enough to allow for any of the turbulence associated with fricatives, it's better seen as an approximant. However, since the difference won't matter for our purposes, I'll use the same symbol that was used for the English voiced dental fricative, [ð], for this sound.

Like English /t/, Spanish /d/ Varies in a Predictable Way

As in the example of English /t/, we need to ask whether this variability in the initial consonant is a special property of the word *de* or whether it's true of Spanish words in general. A brief examination of words containing the voiced dental stop ([d̪]) and the voiced dental approximant ([ð]) in Spanish reveals that this is a general feature of Spanish phonology. So the word *de* has one form in the lexicon, /de/, and the phonological component spells out how /d/ is realized in terms of specific articulations in different phonetic contexts, that is, as different allophones. Here is rough statement of how Spanish /d/ is pronounced.

- When /d/ appears at the beginning of a word following a break, or when it appears following /n/ or /l/, it is pronounced as the voiced dental stop [d̪].
- Otherwise /d/ is pronounced as the voiced dental approximant [ð].

The Spanish example should make two more points clear. First, the relevant context for determining how a phoneme is pronounced can be before as well as after the phoneme in question. Second, the sort of general knowledge about how phonemes are pronounced is specific to particular languages; it is general within the language, but it is not general enough to apply to all languages. We can see this by looking at the English phonemes /t/, /d/, and /ð/, and the Spanish phonemes /d/ and /r/. The English phoneme /t/ can be realized as [t̪], but in Spanish, /t/ (dental, not alveolar as in English) and /r/ (pronounced as [r̪]) are separate phonemes. Spanish /t/ is never voiced and never pronounced as a tap. The Spanish phoneme /d/ can be realized as [d̪], but in English /d/ (alveolar, not dental as in Spanish) and /r/ (always a dental fricative) are separate phonemes. English /d/ is never pronounced as a fricative or approximant.

However, as we will see in the next section, the rules that specify how phonemes in a given language are to be produced or perceived in different contexts are not completely arbitrary. There are good reasons for all of the rules, and, although almost none of them are universal, we can expect many of them to turn up in multiple languages.

This page titled [4.1: Phonetic Contexts](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.2: Assimilation


Anticipatory Assimilation

Exercise 4.2.1

The syllable *un-* means roughly 'not' in a word like *unhappy*, in which it is pronounced [ʌn]. Say the following words beginning with this syllable rapidly, and listen to how the nasal consonant in the syllable is pronounced: *unbelievable*, *unpretentious*, *unkind*, *ungrateful*. Assuming they did not sound like [n], what pattern do you notice in how the pronunciation of the nasal consonant depends on the following consonant?

Why would a phoneme have multiple realizations (allophones) in different contexts? Let's start with a simple example, the realization of /t/ as a dental, rather than an alveolar, stop, that is, as [t̪]. We saw this happening when /t/ comes before a dental fricative, for example, the first /t/ in *at the top*. Compare your pronunciation of this phrase with what it would be like if you pronounced the /t/ as an alveolar stop. (I'm assuming the alveolar pronunciation is the prototypical articulation for /t/ because it's the most common place of articulation for this consonant.) You would have to slide your tongue forward from the alveolar ridge to the upper teeth as you go from the /t/ to the dental fricative /ð/ in *the*. It is simply *easier* to put the tongue behind the teeth for both the stop and the fricative.

This is illustrated in the figure below. The lines labeled "tongue tip contact" and "voicing" represent the relative timing of the articulatory actions that are required for the production of the sequence /æ/, /t/, /ð/ at the beginning of the phrase *at the top*. The two lines at the top illustrate the positions of the tongue for the two possible ways of pronouncing the sequence. The top line shows the movements for the normal pronunciation with a dental stop, [t̪]. The second line shows the movements that would be necessary if the /t/ were alveolar instead. There is an extra movement of the tongue tip to the alveolar ridge before it moves to the front teeth. That is, pronouncing the /t/ this way involves more movements.




If the first stop in *at the top* is really an example of the phoneme /t/, we need an account for why this stop isn't produced in the prototypical way in this word. In this case the speaker *anticipates* the place of articulation (dental) of the following fricative. For English /t/ this appears to be quite general. That is, if we examine a lot of English words, looking for voiceless alveolar stops, we'll see that, in unaffected speech anyway, they don't occur right before dental fricatives. Instead they're replaced in that context by voiceless dental stops. The generalization holds not only for cases where the fricative following the /t/ comes in a separate word, as in *at the top*, but also when both phones are in the same word. One example is in the word *eighth*. Note that in this case the /t/ is not reflected in the spelling, but it is there, at least in my accent. Another example is the /t/ in *width* (spelled "d").

How do we represent words like *eighth* and *width* in the lexicon? The stops in these words are a little different from the stop in *at* because they are *always* pronounced as dental stops (in natural speech anyway). But it's easy to see why this is so: they are always followed by a dental fricative. So, just as with *at*, we will represent these words in the lexicon using /t/, and a realization rule in the phonological component will specify that the /t/ is realized as a dental rather than an alveolar stop. So the phonemic representation

of *eighth* (in my accent) is /eɪt̪θ/, and a more detailed phonetic representation showing the place of articulation of the /t/ (and the diphthongization of the /e/) would be [eɪt̪θ]. Likewise, the phonemic representation of *width* (in my accent) is /wɪt̪θ/, and the detailed phonetic representation is [wɪt̪θ].


This is illustrated in the figure below. The prototypical pronunciation of /t/ is shown in bold and designated the "default", that is, the allophone that is used *unless* there is some reason to use another allophone. The arrow connecting the phoneme to the default allophone is dashed in the figure to indicate that it isn't the one used for this word.



But what we see with /t/ is even more general in English. Consider the /n/ in *on the top* and the /d/ in *hide the money*. In both cases, the natural articulation is dental, rather than alveolar. If we take the prototypical articulation of /t/, /d/, and /n/ to be alveolar, we can see what is going on here as a change or process: the prototypical articulation is modified when the phoneme is realized in a particular context. So the realization rules for English can include the general rule:

- *Alveolar stops and nasals become dental when they precede a dental fricative.*

The figure below illustrates the rule for /t/ only. The label on the arrow connecting the phoneme to its dental allophone means that that allophone is appropriate when the phoneme occurs before a dental consonant. The "_" represents the position of the phoneme, that is, directly before the dental consonant.



Assimilation Applying to One Phoneme Often Generalizes to Other, Similar Phonemes

This rule makes sense because it makes articulation easier; one phone, the alveolar stop or nasal, *agrees with* another, the dental fricative. A process in which one phone comes to agree with one or more others in its context is called **assimilation**. Assimilation is a Speaker-oriented process because it makes articulation easier. But notice that the change from an alveolar to a dental consonant should not interfere seriously with comprehension because the resulting sounds are quite similar to the original ones and because English has no dental stop or dental nasal phonemes that could be confused with the sounds that result. In simplifying things for the Speaker, a Speaker-oriented process should not make things too difficult for the Hearer.


Let's look at an example with vowels. Listen carefully to the vowels in the following words as you produce them: *tad*, *tan*; *sag*, *sang*; *jab*, *jam*.

The vowels in the first word in each pair are probably not quite the same as those in the second word in each pair. The vowels in *tan*, *sang*, and *jam* are normally *nasalized*; that is, the velum is lowered during the production of these vowels, allowing air to pass through the nasal cavity as well as the oral cavity. But this makes sense because it is what will be required for the following nasal consonant (/n/, /ŋ/, /m/) in each case. The speaker *anticipates* the nasal articulation of the consonant during the production of the vowel. It is of course possible to keep the velum up during the vowel and then simultaneously make the oral closure and lower the velum for the consonant, but it is apparently easier to get the velum lowering out of the way during the relatively long vowel production. This avoids the need to perfectly coordinate the lowering of the velum with the contact in the oral cavity (bilabial for /m/, alveolar for /n/, velar for /ŋ/).

Note that the nasalization of the vowel in *tan*, *sang*, and *jam* is completely *predictable* from the vowel's context, specifically, the following nasal consonant. This means that there is no need to record the vowel's nasalization in the lexicon; this is a general

property of the phoneme /æ/. So *jam* is /jæm/ in the lexicon, and a realization rule specifies that the /æ/ should be nasal because a nasal consonant follows.


This rule is illustrated in the figure below. [̃] above the vowel symbol is used to indicate nasalization. As before there is a label on one of the realization arrows indicating the context for the rule, in this case, before a nasal consonant.



Not surprisingly the same holds in English for other vowels, though how much the velum is lowered depends on the particular vowel. We can include something like the following in the phonological component for English:

- *Vowels tend to be nasalized when they precede nasal consonants.*

We can illustrate this more general rule as in the figure below. Here "V" means any English vowel.



This means that all English vowels have a nasalized (at least to some degree) allophone, which occurs when the vowels precede a nasal consonant. As before, this Speaker-oriented process should not interfere too seriously with comprehension; Hearers should have no difficulty recognizing the phonemes for the nasalized vowels.

Perseverative Assimilation


Exercise 4.2.2

One rapid, informal pronunciation of the word *something* is ['sʌmʔŋ]. (Recall that [ŋ] indicates a syllabic nasal). Why do you think the final nasal, [ŋ] in the careful pronunciation in many accents, has become [m]?

Each Language Has its Own Assimilation Rules

In the examples we've seen so far, the assimilation is **anticipatory**; a phoneme changes its pronunciation in order to agree with a *following* phone on some dimension. Assimilation can work in the other direction as well. Let's look at Spanish /d/ again. Recall that this has an approximant allophone, [ð], when it follows any phone other than /n/ or /l/. For Spanish the only other possible previous phones are the others that can end a syllable: vowels, semivowels (/w/, /y/), /r/, /s/, or another /d/ (realized as the approximant). None of these is a stop; that is, there is no complete and sustained closure of the oral cavity (for /r/, a tap, there is a rapid closure followed by a release). Following any of these phones, pronouncing the /d/ as a stop would interrupt the open oral cavity. The approximant [ð] preserves this opening, and in this sense it is an example of assimilation: the consonant takes on one of the features of the previous phone (in the case of /r/ it is the open state of the oral cavity *following* the tap itself). In this case the assimilation is **perseverative**; a feature of one phone "perseveres" during a following phone.

The figure below illustrates this Spanish rule, though just for the case where /d/ follows a vowel. In this case, the "_" symbol follows the relevant context ("vowel") because it is what precedes /d/ that determines which allophone to use.



But why wouldn't this happen when the /d/ follows /n/ or /l/? Note that both of these Spanish consonants *do* involve contact of the tip of the tongue with the teeth, precisely the place where the stop [d] has its contact. Thus in these cases pronouncing the /d/ as the

approximant would involve *releasing* the contact that has already been made for the previous consonant. So it makes some sense that /d/ does not change its manner of articulation in these contexts.

As with the English assimilation of alveolar to dental place of articulation, this process in Spanish applies more generally than to just /d/. Let's look at what happens to words beginning with voiced bilabial stops and voiced velar stops. I'll use the symbols [β] and [ɣ] to represent voiced bilabial and velar approximants.

1. *blanco, caballo blanco*
2. *gordo, caballo gordo*

In these examples, [β] occurs when the phone appears at the beginning of a word following a pause, and [ɣ] occurs when the phone follows a vowel. Similarly, for [g] and [ɣ], the stop occurs following a pause and the approximant occurs following a vowel. This is exactly what we saw with /d/, so it appears that Spanish also has /b/ and /g/ phonemes, each with (at least) two allophones and that a general rule applies to change the default stop manner of articulation to approximant manner of articulation when the phoneme follows a vowel.

However, unlike for /d/, /b/ and /g/ also become approximants following /l/, as in these examples:

3. *el blanco*
4. *el gordo*

Why would /d/ remain a stop following /l/ while /b/ and /g/ become approximants? Recall that the behavior of /d/ following /l/ made sense because the tongue makes contact with the teeth for /l/ as it does for the dental stop [t̪]. But the articulator for [b] is the lips and for [g] the back of the tongue (making contact with the velar region of the roof of the mouth). Since neither of these gaps is closed during the production of /l/, leaving them open for a following /b/ or /g/ would represent the same sort of perseverative assimilation we see when these consonants follow vowels.

That leaves the situation with /b/ and /g/ following nasal consonants, which is a bit more complicated and more interesting. I will leave it for one of the problems for this section.

Let's look at one more English example, another of the allophones of English /t/, which illustrates both perseverative and anticipatory assimilation. But before we consider /t/, let's look at one of the allophones of English /d/. Listen to what happens to the [d] in *do* when it is preceded by a vowel and not stressed.

1. *How do we go?*
2. *Who do you like?*
3. *Why do I think that?*

When *do* is pronounced in isolation, it starts with a stop, but when it follows a vowel and is not stressed, the [d] becomes a tap. The difference is not very great since there is still a contact at the alveolar ridge (as well as voicing), but the contact in the case of the tap is made by a quick gesture of the tongue tip and the contact is very brief. If we look at a lot of English words, we see that this is generally true for North American English /d/. For example, /d/ is pronounced like a tap in words like *rider*, *muddy*, and *needed*. The realization rule is something like the following:


- When /d/ begins an unstressed syllable, follows a vowel, and precedes a vowel (including /əɪ/ but excluding [ŋ] and [ɹ]), that is, the variants of /əɪn/ and /əɪl/ without the vowel), it is realized as a tap ([ɾ]).

Why might this be? Though we probably cannot call it assimilation because the oral cavity is open before and after the consonant and closed as the consonant is produced, this does appear to be a Speaker-oriented process. To pronounce the /d/ like a stop, the tongue tip must make contact with the alveolar ridge and remain there for some time. For the tap articulation, it must only make brief contact with the alveolar ridge. Thus in general the tap appears to be easier to execute in this environment.

Both English and Spanish Have Examples of Consonants Becoming More Vowel-like Following Vowels

Now let's consider /t/ in the same context. We already saw in the section on contexts that /t/ may also be realized as a voiced tap. We saw this with the word *at*, but we can also see it in the middle of words such as *butter*, *city*, and *Italy*. If the default (prototypical) pronunciation of /t/ is as a voiceless stop, there are two changes here: the /t/ becomes a tap and it becomes voiced as well. We have just seen (for /d/) how pronouncing a stop as a tap when it comes between two vowels may simplify articulation for the speaker. The voicing of the /t/ is an example of assimilation. Both the phone before and the phone after the /t/ in these words are

voiced, so allowing the voicing to continue through the articulation of the /t/ simplifies matters somewhat for the speaker. Because the assimilation points in both directions, it is both perseverative and anticipatory. The figure below illustrates this process. It shows the third allophone of /t/ that we have seen, [R], and an indication of the context where it is used. The "_" appears between two vowels ("V"), the second of which is unstressed.



While assimilation makes good sense from the perspective of the Speaker, we can't use it in general to predict how phonemes in different languages will behave. First, particular assimilation rules that operate in one language or dialect may not operate in another. For example, English does not have a rule like the Spanish rule that causes voiced stops to become approximants (although something like this may happen in English in very rapid, casual speech), and in most English accents /t/ is not voiced when it comes between vowels as it is in North American accents. In other words, there is an arbitrary aspect to assimilation; which kinds of assimilation apply to a particular language and dialect must be learned. Second, some of the allophonic variation that is exhibited by phonemes seems to have nothing to do with assimilation. We will see examples of this type next.

Non-assimilative Allophonic Variation

Let's go back once more to the most complicated consonant in English, /t/. Compare the /t/ in *till* with the /t/ in *still*. As described when we discussed consonant voicing and voice onset time, the /t/ in *till* is aspirated. That is, there is a significant lag between the release of the contact and the beginning of voicing, and you can feel a puff of air being expelled during this lag. To show this detailed pronunciation, we can write [t^h]. But note that this is not the way the /t/ in *still* is pronounced. For this sound the release and the beginning of voicing roughly coincide.

But why not treat these two kinds of *t*-sounds as two different phonemes? Recall once again that different phonemes are used to distinguish words from one another. Can these two kinds of English *t*-sounds distinguish words from each other? If we could have the aspirated *t* ([t^h]) in a word like *still*, then we could have one word pronounced [stɪl] (with unaspirated *t*) and another word pronounced [st^hɪl] (with aspirated *t*). But the aspirated sound never appears in this position, and if it did (and this is important), English hearers would still hear the word as *still*. The same would be true if we used the unaspirated sound in the word *till*. English hearers would still hear the word as *till*, or perhaps as *dill*, but not as some other word with neither /t/ nor /d/ at the beginning.

So it seems that the distinction between [t] and [t^h] cannot be used in English to distinguish one word from another. Since these two sounds are very similar to each other and they cannot be used to distinguish words from each other, we conclude that they belong to the same phoneme, that they are both allophones of /t/. But notice that the realization of /t/ as (unaspirated) [t] or (aspirated) [t^h] is a regular phenomenon, something that English speakers know how to do. This means that we need realization rules in the phonological component telling when [t] is appropriate and when [t^h] is appropriate. Here's a stab at what they would look like.

- When /t/ appears at the beginning of a stressed syllable, it is aspirated. (The details of its pronunciation depend on what follows it.)
- Otherwise /t/ is unaspirated. (The details of its pronunciation depend on what precedes and what follows it.)

But we can be even more general than this. Consider the words *pot* and *spot*. The /p/ sounds in the words show exactly the same alternatives as the /t/ sounds in *till* and *still*. And the same is true for the /k/ sounds in the words *car* and *scar*. If we examine a lot of words containing /p/ and /k/, we see that the alternation is just as general as it is for /t/. So the realization rule can apply to all three phonemes:

- When a voiceless stop appears at the beginning of a stressed syllable, it is aspirated.
- Otherwise a voiceless stop is unaspirated.

Speakers of a Language Know How and When to Produce the Allophones of a Phoneme, But They May Not be Aware of Any of the Differences

It is usually not surprising for English speakers to learn that the [t^h] in *till* and the [t] in *still* belong to the same category. In fact they probably were not even aware that there were these two different sounds, and some English speakers cannot hear the

difference even after it is described to them. The difference between these two allophones of /t/ in one sense *doesn't matter* for English in the way that the difference between /t/ and /d/ does matter. English hearers do not have to identify a sound as being [t] or [t^h] because what they care about is what words they are hearing, and the difference between these two sounds never matters for the words. However, they do have to identify a sound as being /t/ or /d/ because it can make a difference, say, between *till* and *dill* or between *bat* and *bad*. Learning English apparently includes learning to emphasize the differences between sounds such as [t] and [d] and de-emphasize the differences between sounds such as [t] and [t^h].

But why would these English stops behave this way? These are not examples of assimilation: the /t/ in *stop* cannot be said to agree with the /s/ that precedes it or the /ɑ/ that follows it any more than it would if it were aspirated. If we view the unaspirated stops as the default (prototypical) allophone, then why would these stops get aspirated when they come at the beginning of stressed syllables? The origin of this process is probably more Hearer-oriented than Speaker-oriented. For the Hearer, what matters is that /t/ sounds different enough from nearby phonemes, namely, /d/ and /θ/. In terms of voice-onset time, [t^h] is further from [d] than [t] is, especially at the beginning of a syllable. The gap between the release and voicing is longer, and there is the (possibly audible) puff of breath. Since aspiration requires more breath, it is easier to achieve in stressed syllables, which are executed with a greater effort. Thus English has settled on a set of realization rules for voiceless stops that maximize the distinctions between voiced and voiceless stops at the beginning of stressed syllables.

Note, however, that in some other contexts, the distinction between voiced and voiceless stops disappears altogether in English. Following an /s/ in the same syllable, only voiceless stops are possible; that is, there are no words like /sdɪl/ or /sbɪn/. And in North American English, the distinction between /t/ and /d/ disappears at the beginning of an unstressed syllable between vowels; both are realized as the voiced tap [ɾ], as we have seen. (Most speakers, however, make a distinction between the *vowels* preceding these consonants, so *writer* and *rider* do not sound quite the same for these speakers.)

As with the examples of assimilation that we discussed, we've seen aspiration of voiceless stops in English makes some sense, at least from the perspective of the Hearer. But it is still a convention of English, something that we should not necessarily expect in other languages. In Spanish, voiceless stops are never aspirated; in Amharic, they are weakly aspirated in all contexts. In Mandarin Chinese, the relevant distinction is between voiceless stops that are always aspirated and voiceless stops that are never aspirated; that is, there are no real voiced stops. And in Hindi, there is a three-way distinction, between stops that are voiced, stops that are voiceless and always unaspirated, and stops that are voiceless and always aspirated. Finally, even when voiceless stops are aspirated in some contexts but not in others in a language, the pattern may be different from in English. In Tzeltal, voiceless stops are aspirated at the *ends* of syllables, but *not* at the beginnings.

This page titled [4.2: Assimilation](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.3: Distribution of Phones

Evidence for Learning Phonology

Exercise 4.3.1

Say you're trying to learn about the phonology of an unfamiliar language, but all you have access to is recordings of a set of sentences. You do not know what the sentences mean, nor do you know where the boundaries between the words are. What could you and could you not learn about the phonology of the language given this information?

To gain a better understand of how phonology works, in this section we'll take the perspective of the learner. What does a learner have to achieve with respect to phonology? The learner has to learn to be both a Speaker and a Hearer, to figure out what the categories (phonemes) are and how they are realized in different contexts. What the learner is presented with is allophones of course, not phonemes directly. There are two kinds of evidence that the learner can use to arrive at what the phonemes are.

First, the learner could just pay attention to what phones tend to occur. As I've said a number of times, there are an infinite number of possible phones and, even within a given language, a very wide, if not infinite, set of possibilities. But for a given language, the phones the learner hears will tend to cluster in particular regions within the space of possibilities. For example, if the language is Spanish, there will be many possible vowel phones, but they will tend to cluster around the vowels [i], [u], [e], [o], and [a]. For example, nothing very close to [æ] or [ey] or to [ω] (an unrounded high back vowel) will occur. There will also be many possible consonant phones, but they will tend to cluster around particular points in the space of possibilities. These will include [ɸ] (a voiced dental stop), [ð], [t̪], and [ɾ], but not [ɹ] or [t̪ʰ] (an aspirated dental stop) or [t̪ʰ] (an ejective dental stop). Of course these tendencies would be specific to Spanish. [æ] *would* occur if the language were English; [t̪] *would* occur if the language were Amharic.

Apparently it's Impossible to Learn What the Phonemes in a Language are Without Paying Attention to Meaning

But listening to what phones occur and what phones do not does not provide any direct information on how they are used contrastively, that is, to distinguish words. For example, a child exposed to English will hear a variety of stops in contexts following vowels. Some of these contrast, for example, [t] and [p], but others do not, for example, [t] and [ɾ]. The only way to know for sure that they contrast is to pay attention to meaning as well as to the patterns of phones that occur. For example, if the Learner can tell that [ræt] and [ræp] mean different things, they will know that [t] and [p] contrast.

So there are these two sorts of evidence that the Learner can use: what sounds tend to occur and what different sequences of phones mean.

What There is to be Learned

But what sort of knowledge of phonology is there to be learned? We've seen that knowledge of the phonology of a language includes the following.

1. Knowledge of what the contrastive categories of the language are, that is, the basic units that are used to make (and distinguish) the words of the language. In spoken languages these include phonemes and supragemental features; in signed languages the contrastive units appear to be syllable types.
2. Knowledge of how the contrastive units are realized as particular forms in different contexts. This knowledge needs to be in two forms, one that enables the Speaker to *pronounce* words and one that enables the Hearer to *recognize* words.
3. Knowledge of how the contrastive units may (and may not) be combined to form words (phonotactics).

In the rest of this section and the next section, we'll be considering how a person would learn the three types of knowledge. To simplify matters, we'll start by looking at a simple imaginary language, the one that a tribe of Lexies has arrived at in an early stage of the evolution of their system of communication. And rather than looking at data on children's production and comprehension, we'll first look at what kinds of information the child might have access to and might be useful in learning about phonology. In fact what I'll be discussing would apply just as well to a linguist who is trying to figure out the phonology of a previously unresearched language.

Of course infants and adult linguists differ in many important ways. For one thing, linguists are conscious of what they are learning about the language; their conclusions will be things they can describe and write down. Infants, on the other hand, are not conscious of any of this and will not even be conscious of the phonological knowledge they have when they grow up. Second, linguists can *elicit* data; that is, they can ask questions to test their hypotheses. Children obviously could not do this even if they were aware of what they were learning. Still, the task of the linguist and the "task" of the infant bear some interesting similarities.

Learning Phonology is like Learning Meaning in Some Ways

[po]	'father'
[pi]	'rock'
[pe]	'sun'
[pobo]	'leg'
[pudu]	'cave'
[pana]	'head'
[mi]	'river, hawk'
[maga]	'rain'
[ti]	'cat'
[tu]	'man'
[tede]	'mosquito'
[tunu]	'milk'
[toso]	'ground'
[su]	'mother'
[zu]	'mother'
[sama]	'eye'
[zama]	'eye'
[nu]	'child'
[nasa]	'tree'
[nini]	'mountain'
[ke]	'woman'
[ku]	'moon'
[kene]	'tiger'
[kudu]	'snake'
[kini]	'path'
[ŋo]	'fire'
[ŋobo]	'hair'
[ŋasa]	'hand'

Figuring Out Phonotactics

Let's begin by thinking about the third kind of knowledge, phonotactics, because it will help us figure out the other two. Examining the general structure of the words, we see that they can consist of one or two syllables and that all of the syllables consist of a consonant followed by a vowel. This means that consonants appear in two different contexts, beginning a word and in the middle of a word following a vowel (and preceding another vowel). It also means that vowels appear in two different contexts, at the end of words and in the middle of words preceding a consonant (and following another consonant). So the next question we might ask is whether there are any constraints on which consonants and vowels can appear in which contexts or on which combinations of consonants or vowels occur in two-syllable words. We see that all of the consonants appear in the word-initial context but that only the following ones appear in the third position in two-syllable words: [b], [m], [d], [n], [s], [g], [ŋ]. We also see that all of the vowels can appear in either of the two vowel positions.

In addition, it's hard not to notice a striking regularity to the vowels: in two-syllable words, the first and second vowels are always the same. While languages are usually not this extreme, they often do have constraints on how neighboring phones must agree on some feature. This is true for clusters of final consonants in English, for example. When English syllables end with more than one stop or fricative, these consonants must agree in voicing; that is, either all (or both) must be voiced or all must be voiceless. For example, /kt/, /sk/, /fθs/, /bd/, and /gz/ are possible, but /kd/, /zk/, and /vðs/ are not.

While this language apparently has no stress, if it did we could also look at stressed and unstressed syllables to see if there are phones that can occur in one and not the other type of syllable. In English, for example, unstressed syllables are more constrained than stressed syllables in terms of what can occur.

So we can summarize what we've learned about the phonotactics of the language as follows:

- Syllables always consist of a consonant followed by a vowel (CV).
- Words consist of one or two syllables. In words consisting of two syllables,
 - the second syllable can only begin with one of the following consonants: [b], [m], [d], [n], [s], [g], [ŋ]
 - the vowels of the two syllables must be the same.

Minimal Pairs and Overlapping Distributions

Exercise 4.3.2

Say a language learner discovers the following forms in the target language:

1. [vam] 'break'; [fam] 'snow'
2. [lo] 'picture'; [lu] 'picture'
3. [kes] 'lip'; [kes] 'radio'

What does the first pair of words tell us about the status of [v] and [f] in the language? What does the second pair of words tell us about the status of [o] and [u] in the language? What does the third pair of words tell us about the language?

Now we need to figure out what the phonemes of the language are and how they're realized. Obviously it's important to know which phones occur (and which do not). As I noted at the beginning of this section, the phones in a language should tend to cluster around particular prototypical places, places that differ from one language to another. The transcriptions of the words above are meant to represent this. So, based on the words heard, the child has the vowels [a], [i], [e], [u], and [o], and the consonants [p], [b], [m], [t], [d], [s], [n], [k], [g], and [ŋ] to deal with. (Note that a lot of the learning process is being left out here; deciding that there are this many phones, no more or less, is no mean feat, and children may in fact not do anything like that early in phonological learning.)

A particular phone P (really a cluster of phones centered on P) is a phoneme in a language only in the sense that it contrasts with the other phonemes in the language, that is, that the difference between P and those other phonemes can make a difference in meaning. This means that we can only establish what the phonemes are by comparing the different phones with one another. But which pairs should we be comparing and what sorts of comparisons should we be making? There's no point in comparing phones that are very different from one another because changing from one of these to another almost certainly changes the meaning. For example, in English we'd never expect the phones [b] and [s] to belong to the same phoneme. Rather what we're interested in are

pairs that are relatively similar. For such pairs it is possible that both phones belong to the same phoneme, or that they belong to different phonemes. "Similar" phones will be phones that differ in only one or two features.

A Minimal Pair is the Clearest Evidence that Two Phones are Separate Phonemes

Let's start with the vowels because there are fewer of them. Of the five vowels, pairs that are somewhat similar include [i,e], [u,o], [a,o], and [e,a] ([a] is a low, central vowel). For each pair, we are interested in whether the difference between the two is enough to make a difference in meaning. The best evidence for this would be two words that differ only in that one has one of the phones, and the other has the other. Such a pair of words is called a **minimal pair**.

We have a minimal pair for [i] and [e], the words [pi] and [pe]. Both of the forms consist of two phones, the first of which is [p]; clearly the only difference is that one has [i], the other [e] in second position. It is important that we not only have two forms that differ in only one way but also that the two forms have different meanings. Otherwise they would not actually be different words. Since [pi] means 'rock' and [pe] means 'sun', and these two meanings are not obviously related to each other, it's clear that [pi] and [pe] are different words. And since the only difference in the forms is the difference between [i] and [e], we can be fairly sure that [i] and [e] are separate phonemes in the language. Let's tentatively call them /i/ and /e/, where the phoneme labels selected are supposed to represent the prototypical allophones. As far as we know so far, these apparent phonemes have only one allophone each, so this is the one we'll select for the phoneme label.

What about [u] and [o], the comparable pair of back vowels of different heights? Looking through the list of words, we find no minimal pairs for [u] and [o]. But this does not necessarily mean that these two phones could not be used contrastively, that is, that they are not separate phonemes. We would have evidence for this if we could show that they are used in the same contexts, that is, that they can appear next to the same phones. If they're used in the same context, then the difference between [o] and [u] can't be due to assimilation or some other process related to context because if this were true, the contexts would have to be different for the two phones. In fact, it would be enough to show that they are both used in one particular context.

The range of contexts that a phone can appear in is called its **distribution**. We already know that all vowels can appear in one-syllable words and as either vowel in two-syllable words, so a vowel is always preceded by a consonant and sometimes followed by a consonant. What we'd like to know is which consonants can come before and after [o] and which can come before and after [u]. Looking at the words with these vowels, we find that, among other consonants, [p] and [t] can come before both vowels and that [b] can come after both vowels. So the indication is that [o] and [u] occur in the same contexts, or at least that their distributions *overlap*. Even though there are no pairs of words distinguished only by the difference between [o] and [u], it appears that there *could* be. For example, based on everything we know about [o] and [u], we could imagine a word pronounced [pu] that would mean something different than the word pronounced [po] 'father'. In other words, it appears that [o] and [u] are separate phonemes. We'll call them /o/ and /u/ tentatively.

We can follow the same procedure for the other vowel pairs. The realization rules for the vowels are simple. Since, as far as we can tell, each vowel phoneme has only one allophone, each vowel is always realized as that allophone.

Complementary Distributions

Establishing the Status of Two Phones Involves Looking at their Phonetic Contexts

Now let's consider the consonants. One possible set of pairs is consonants that differ only by voicing: [p,b], [t,d], [k,g], [s,z]. In many languages such pairs of voiced and voiceless consonants are allophones of the same phonemes. There are no minimal pairs in the list for any of these pairs of phones, so we need to see whether they can appear in the same contexts, as we did for the pair [o,u]. For [p,b], we discover that [p] appears only at the beginning of words, whereas [b] appears only in the middle of two-syllable words, that is, between vowels. In other words, there is no overlap at all in the distributions of [p] and [b]. In this case we say they are in **complementary distribution**; there is no overlap at all in their distributions.

Two similar phones that are in complementary distribution cannot be separate phonemes because we *can't* replace one by another in a form to get a different word. That is, if we're right about the distribution of [p] and [b], we can assume that there could be no form [ba] that would make a minimal pair with the existing form [pa] and no form [popo] that would make a minimal pair with the existing form [pobo].

We can conclude that [p] and [b] belong to the same phoneme. We'll call it /p/, though we have no way at this point of knowing whether [p] or [b] is the prototypical allophone. The realization rules for /p/ are fairly simple. It is pronounced [p] at the beginning and [b] in the middle of words. With [p] as the default allophone, we can see the [b] allophone as resulting from assimilation. In the middle of words, the consonant is surrounded by vowels, that is, voiced sounds, so voicing it (changing it from [p] to [b]) makes it agree with the context on the voicing dimension. For this reason, it makes sense to choose [p] as the default allophone for this phoneme.

Languages tend to be systematic, so we should not be surprised when we find the same sort of distributions for the other stop pairs [t,d] and [k,g]. That is, the voiceless phones in each case appear only at the beginnings of words, while the voiced phones appear only in the middle of two-syllable words. Again we conclude that each pair represents a single phoneme. We'll call these phonemes /t/ and /k/. The realization rules are the same as for /p/, so at the point, we can make a more general realization rule for all three of the stops in the language: *Pronounce the stop voiceless at the beginning of a word, and pronounce it voiced in the middle of a two-syllable word (between vowels).*

Sometimes a Phoneme's Realization Depends on Formality, Rate of Speaking, or Degree of Emphasis on the Word

For [s] and [z] we have what at first glance appears to be a minimal pair, [su] and [zu]. But this is not a minimal pair because the two forms have the same meaning, 'mother'. Apparently these are not different words for 'mother', but alternate ways of saying the same word. It is not clear from the list in what situations the different pronunciations are used. One possibility, similar to what we discovered for the pronunciations of English *at*, is that the pronunciation depends on what precedes the word. Another is that the difference is related to formality, speed, or emphasis. Something like this happens in English with word-final voiceless stops. The /p/ at the end of a word such as *lip* would normally not be aspirated or released. But if the speaker is speaking unusually formally or slowly or with a great deal of emphasis on the word in question, the /p/ can be released and aspirated.

In any case it's clear that for the pair [su] / [zu] the difference between the [s] and the [z] is not contrastive; changing from one to the other makes no change in meaning. We should also notice the same thing going on with the two other forms beginning with [s] and [z]: [sama] and [zama]. Again the difference in the initial consonants makes no difference in meaning.

So far the evidence that we have indicates that [s] and [z] belong to the same phoneme. If we examine the other forms in the list that contain either [s] or [z], we find only three others, all containing [s] at the beginning of the second syllable: [toso], [nasa], and [ɲasa]. There are no words with [z] in this position. So, as far as we can tell from this list of words, there are no positions in words in which [s] and [z] make different words, and we conclude that [s] and [z] are allophones of the same phoneme. Because [s] seems to occur in more contexts than [z], we can consider it to be the prototypical allophone, and we'll refer to the phoneme as /s/. The realization rules for this phoneme only need to specify that it is optionally pronounced as [z] when it begins a word. Note that the rule is different than the ones for the stops, which get voiced when they are in the middle position in words.

There are three other consonants, the three nasals produced in the same three places of articulation as the other consonants, [m], [n], and [ɲ]. There is a minimal pair for [n] and [ɲ], [nasa, ɲasa], indicating that they are separate phonemes. For the other two pairs, [m,n] and [m,ɲ], there are no minimal pairs, but clear evidence of overlapping distributions: [maga], [nasa], [ɲasa]; [sama], [pana]. So we conclude that there are three nasal phonemes: /n/, /m/, /ɲ/.

Based on the evidence in the list of words, we can propose that the Lexie language has five vowel phonemes and seven consonant phonemes. But it's important to note that our baby has heard just this short list of words; there is more to the language. All of the generalizations that we have made about the phonemes and the phonotactics of the language could prove wrong with more evidence. In particular, whenever we concluded that something could *not* occur, the child could later discover that such a thing could occur. For example, we concluded that words consisted of one CV syllable or two CV syllables, but it is always possible that a word not yet heard could have a different structure, such as CVC. Similarly, we concluded that [p] and [b] are not contrastive because [p] always occurs at the beginning of words and [b] always in the middle. But the baby could later encounter (or become aware of) a previously unknown word with a form like [be] or [napa]. In fact, even our minimal pairs are suspect. Let's see how.

We've seen that languages use phonemes to distinguish words from one another. But this is only the usual case. We haven't yet considered another possibility for different words. Notice that in this Lexie language, [mi] can mean 'river' or 'hawk'. It's hard to see how these two meanings are related to one another, so we have to conclude that [mi] is really two different words that happen to have the same pronunciation. Such words are called **homophones**. Many (perhaps all) languages have homophones, probably

mainly the result of historical accidents, changes that happened to bring the forms of two words together. Homophones represent an example of **ambiguity**, a situation in which a form has more than one possible interpretation. We will meet more ambiguity later in the book. Ambiguity presents a potential problem for Hearers because by itself the form cannot be interpreted. Hearers can normally solve this problem by using the context of the ambiguous form, either the other words that it appears with or the situation that it refers to. However, languages normally do not have very many homophones because of the burden this would place on Hearers.

Now consider a minimal pair like [pi] and [pe] again. Because these have different meanings, they are apparently different words, and we used this fact to conclude that [i] and [e] are separate phonemes. But what if [i] and [e] are just allophones of a single phoneme that varies considerably in how it's pronounced, even in a context at the end of words? In that case, [pi] and [pe] would be an example of a pair of homophones, both realizations of a phonemic form that we could write as /pi/. But homophones are rare (for good reason), and it is much more likely that the difference between [i] and [e] is contrastive, as we originally concluded. Still, more evidence would help us decide.

How to Figure Out the Phonology of a Language (Within Limits)

Now let's summarize what we've learned in the form of a set of instructions for discovering the phonology of a language. You can use these as a guide when solving phonology problems concerning real languages. But first, remember this caveat: any sample of words is necessarily incomplete, so learners can't be completely sure of their conclusions. It's better to think of the conclusions as hypotheses. The more data there is, the greater the learner's confidence in the hypotheses.

1. Learn something about the phonotactics of the language. Begin by looking at the pattern of consonants and vowels that make up the words. This should tell you about syllable structure and maybe about constraints on the form of words. It may also be possible to make more generalizations by examining subcategories within consonants and vowels, for example, voiced consonants, and by looking at suprasegmental features such as stress. In particular, stressed and unstressed syllables may have different structures.
2. Within the list of all of the phones you have heard, pick pairs of similar phones. For each, look for evidence that the phones belong to the same phoneme or different phonemes. For each pair, do the following.
 1. Look for minimal pairs. If you find them, conclude that the phones are separate phonemes.
 2. If you don't find minimal pairs, look at the contexts that the phones occur in to see if the contexts overlap. Can they be followed by the same or similar phones? Can they be preceded by the same or similar phones? If there is considerable overlap (you have to use your judgement here), conclude that the phones are separate phonemes.
 3. If the contexts don't overlap, you probably have allophones of the same phoneme. To make sure, and to figure out what the realization rules for the phoneme are, look for the following.
 1. Try to see if the distributions of the two phones are complementary (the usual situation for allophones). If so, the realization rules are straightforward: the phoneme is pronounced one way in one context (or set of contexts), the other way in the other context (or set of contexts).
 2. Look for cases where you have forms with the same meaning differing only with respect to the two phones. If you find them, it may be that the choice of allophone in this context depends on factors such as formality. Alternately, if the phone is at the beginning or end of a word, the choice could depend on what phone comes on the other side of the phone.

But what do we see when we look at actual data from children learning the phonology of their first language and adults learning the phonology of a second language? That's the topic of the next section.

This page titled [4.3: Distribution of Phones](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.4: Learning Phonology

Babies Learning Phonology

Exercise 4.4.1

Since a newborn infant is (equally) capable of learning any human language, what abilities would an infant have to come equipped with to allow it to learn phonology?

The Beginning

As we've seen, languages differ with respect to what is contrastive: all languages treat some distinctions as significant and others as non-significant. In English the difference between [ɪ] and [i] matters — it distinguishes words from one another — but in Spanish it doesn't. In Amharic the difference between [k] and [k'] matters; in English it doesn't. When an infant is born, it is capable of learning any human language. How could it figure out which distinctions are contrastive and which aren't?

Obviously the only information the infant gets to help it is the language that it hears being spoken around it. If it is to figure out which distinctions matter, it has to be able to hear those distinctions. So for example, it has to be able to hear the difference between [i] and [ɪ] or the difference between [k] and [k']. In fact experiments with very young infants indicate that they can perceive all of the distinctions that matter in the world's languages. This is quite impressive since, as we'll see, they tend to *lose* the ability to hear many of these distinctions later on.

Babies Start Learning the Sounds of the Language Around Them Long Before They Understand Any of It

So what happens when an infant starts getting exposed to a particular language? For at least six months after birth, the baby makes the same sorts of sounds regardless of what language this is. (Here's an example of what a baby sounds like at three months.) That is, if we only looked at the baby's *production*, we would not see any signs that learning is taking place. But it is. Within a few months after birth, babies can distinguish the language around them from other languages that they haven't heard. This means that they have already become sensitive to some of the properties that distinguish languages from each other.

What sorts of properties? Experiments with eight-month-olds show that they are capable of learning the frequencies at which different phones occur together. Nine-month-olds can distinguish sequences of sounds that obey the phonotactics of their language from those that don't and can also distinguish sequences that are frequent in their language from sequences that are phonotactically legal but infrequent (for example, /cʌn/ vs. /tʌʃ/ in English). Babies apparently start learning phonotactics at a very early age. As we saw in the last section, knowing the phonotactics of the target language comes in handy in learning the phonemes of the language because it makes clear what the possible contexts are.

Babbling

Starting at around their sixth month, the sounds that babies make start to take on a different character. This stage, which normally lasts about 12 months, is called **babbling**. (Here's an example of what a baby learning English sounds like at nine months.) Babies start producing simple syllables, such as [ba], and these may include a very wide variety of consonants and vowels, including many not found in the target language. Later, several things happen. Babies begin to string the syllables together in sequences, such as [bababa] and [batabatabata]. And the sounds they are producing begin to resemble the target language more and more. What happens with deaf children is somewhat more complicated. They begin this phase like hearing children, producing a range of simple syllables. But the sounds they make never get more language-like and never come to resemble the spoken language around them. Deaf children who are exposed to a sign language from birth also go through a sort of sign babbling phase, however, "babbling" with their hands.

In the Second Half of their First Year, Babies Sound More and More like the Language Around Them

What is going on during babbling? It seems to have three related sorts of functions. First, it should be clear from what we've learned about phonology that producing language is very complex and requires a great deal of coordination. In this sense, babbling may be a form of practice; the baby is figuring out how to use its articulatory apparatus in a fluent manner.

Second, the baby has to learn to tie what it hears to what it says. The auditory and the articulatory properties of linguistic sounds are totally different things, but, as we've seen, a particular phoneme needs to be associated with both. How could the baby learn to make this association? It's possible that during babbling the baby tries out various articulatory positions and movements and then listens to the auditory consequences, associating the behavior with the sounds each time this happens.

Third, the baby has to learn to sound more like the target language. This may work through a mechanism known as reinforcement learning. The baby tries out a particular articulatory pattern, listens to the consequences, and if these sound close to the kinds of sounds it is hearing around it, that articulatory pattern gets reinforced for the baby. The result is that the baby is more likely to produce this pattern later on. If, on the other hand, the sound that is produced sounds very different from the linguistic sounds in its environment, that articulatory pattern fails to get reinforced, or is penalized. In that case, the baby will be less likely to produce the pattern later on.

While this is going on, there are also changes in what the baby can perceive. As it begins to learn the phonemes of the language, it begins to lose the ability to hear distinctions that are not contrastive in that language. One well-known example of this phenomenon is the distinction between [r] and [l], a phonemic distinction in English, but not in Japanese, Lingala, Inuktitut, or many other languages. For many speakers of languages like these, the ability to hear the distinction is lost. We'll see more examples when we consider what happens in second language learning below.

Simplification

Children start producing recognizable words around the beginning of their second year. In the beginning what they produce only very roughly approximates the forms they are hearing. Partly this may come from not having worked out what the phonemes of the language are yet; this may take several years. But it also results from the inability to produce some of the distinctions that the child does hear. Production lags behind perception throughout the learning of phonology.

But the forms the child produces in its second, third, and fourth years do not deviate from the adult forms in random ways. Rather they can be seen as **simplifications** of the adult forms. The kinds of simplification include the following.

- Some phones are inherently more difficult to produce than others, so there tends to be a order in which phonemes are learned. For example, stops and nasals are easier to produce than fricatives. For stops the articulators are brought together completely, whereas for fricatives they must be brought close enough together to yield the characteristic fricative turbulence but not so close as to block the passage of the air completely. The result is that young children may replace fricatives with stops ([mɛti] *messy*) or replace all fricatives with one fricative that they've mastered ([mɛfi] *messy*).
- Syllable structure may be simplified. As we saw when we looked at the phonotactics of different languages, the simplest syllables are those consisting of a vowel preceded by at most one consonant. Thus young children sometimes simplify syllables ending in one or more consonants by dropping them ([kæ] *cat*) or simplify syllables beginning with consonant clusters by dropping all but one of the consonants ([tɒp] *stop*).
- In words of more than one syllable, children may replace one of the phonemes with another one found in the word ([kiki] *kitty*), drop a syllable, or combine two syllables ([bænə] *banana*).

Adults Learning Phonology

Exercise 4.4.2

We learned that the Spanish voiced stop phonemes /b, d, g/ are pronounced as approximants when they follow vowels. Yet English speakers learning Spanish tend to pronounce them as stops in all contexts. Why might this be?

When a person learns a language later than in the first few years of life, their success depends on their age, as well as on a range of other factors, though why age matters and which ages, if any, are crucial remain hotly debated issues. Here we'll just consider what happens in adult language learning. I'll be using the term "second language" to refer to any language other than the learner's native language since in most ways the learning of third and later languages follows the same pattern as the learning of a second language.

The learning of second-language phonology seems to be quite independent from the learning of grammar and vocabulary. Most of us are familiar with people whose grammar and vocabulary are indistinguishable from those of a native speaker but who still have a

noticeable foreign accent. So what I'll be discussing in this section applies only to the learning of pronunciation. We'll look briefly at the learning of second-language grammar later on.

The phonological learning *task* of the adult learner is the same as that of the baby: to figure out what distinctions in the target language are contrastive and to learn how to produce and recognize the different phonemes in different contexts. But what we see in adult phonological learning *looks* quite different from what we see with babies. One clear difference is the amount of variability. Normal children learning their first language end up roughly equivalent in their ability to pronounce and understand the language. Adults, on the other hand, differ dramatically from one another. While the great majority of adults never achieve native-like proficiency in the pronunciation of a second language, no matter how much they are exposed to it, the degree of foreign accent they exhibit varies a lot from one learner to another.

By Comparing the Phonology of the First and Second Languages, We Can Predict the Kinds of Errors that Second-language Learners Will Make

An even more important difference between child and adult language learners stems from the fact that adults already know the phonology of at least one language. This can both help and hinder them in their learning of the new phonological system. In general the influence of a body of knowledge on the learning of new knowledge is called **transfer**.

The crucial issue is the ways in which the first language phonology agrees with that of the second language. When they agree more or less perfectly, we can expect **positive transfer**. That is, knowledge of the first language makes the target language *easier* to learn than it would be for learners with other first languages. For example, as we saw in the section on vowels, Spanish and Japanese vowels are quite similar. In general, it is easier for a Spanish speaker to learn Japanese vowels or for a Japanese speaker to learn Spanish vowels than it is for an English speaker to learn the vowels of either of these languages or for speakers of either of these languages to learn English vowels. Similarly, English and Spanish both have the phoneme /f/, realized in virtually identical ways, whereas Japanese has no such phoneme. In general, then, it is harder for Japanese speakers to learn this aspect of English or Spanish than it is for English or Spanish speakers to learn this aspect of each other's languages.

Much more noticeable in second language phonology are the consequences of differences between the languages. Differences may result in **negative transfer**, that is, interference from the first language to the target language. As you will see in the section on English accents, accents or languages can differ phonologically in several ways. These differences can often predict some features of foreign accent and areas of phonological difficulty for second language learners.

Phonetic Differences

One possible difference is purely phonetic. The first language and the second language both have a similar phoneme P, distinguished from other similar phonemes, but the phoneme differs in the details of how it is produced. Either it is always pronounced differently in the two languages, or it is pronounced differently in some contexts. Learners will tend to pronounce the phoneme as it is pronounced in their first language.

English /r/, Spanish /r/ ([ɾ])

Though English /r/ and Spanish /r/ are pronounced quite differently, they have similar auditory properties, so learners will tend to map them onto each other. English speakers will tend to pronounce Spanish /r/ as [r] (in *marido* and *cortar*, for example), and Spanish speakers will tend to pronounce English /r/ as [ɾ] (in *marry* and *quarter*, for example).

English /k/, Spanish /k/

In both languages, this voiceless stop is distinguished from the corresponding voiced stop, /g/, but English /k/ has an aspirated allophone [k^h] that is never used in Spanish. Learners should have no problems when /k/ does not begin a stressed syllable. But English speakers learning Spanish will tend to aspirate Spanish /k/ when it begins a stressed syllable, as in *que* and *como*, for example, and Spanish speakers learning English will tend not to aspirate English /k/ in any context, including those where it would be aspirated in English, as in *come* and *quick*, for example.

Phonemic Differences

Second-language Learners May Have Difficulty Hearing and Producing a Distinction that Matters in the Second Language

Another possible difference is phonemic. The second language makes a distinction that is not made in the first language. Learners may fail to hear the distinction and will tend to pronounce the two forms in the same way. Because this sort of difference can interfere with communication, it is more serious than problems of the first type.

Japanese /r/, English /l/ and /r/

While English makes a distinction between /r/ and /l/, Japanese has a single phoneme /r/ that is usually pronounced like an alveolar tap ([ɾ]) but sometimes takes the form of an alveolar lateral (similar to English /l/) or a phone that is somewhere in between these two. Which form it takes may be difficult to predict; it depends on the phonetic context, the speaker, and even the situation. Japanese speakers learning English may fail to hear the difference between English /l/ and /r/, so they may not be able to distinguish *right* from *light*. And they will tend to pronounce both /l/ and /r/ as either [ɾ] or [l].

English /k/, Amharic /k/ and /k'/

As we saw in our discussion of ejectives, Amharic and many other languages make a distinction between plain, non-glottalized voiceless stops and ejective voiceless stops (for example, between /k/ and /k'/), a distinction not made in English. English speakers learning Amharic often have difficulty hearing the difference, so they may not be able to distinguish /kɛbɛro/ 'drum' from /k'ɛbɛro/ 'fox'. And they will tend to pronounce /k/ and /k'/ in the same way, like English /k/.

Phonotactic Differences

The languages may also differ in their phonotactics. If the second language has more complicated syllables than the first language, in particular if it allows more complicated clusters of consonants at the beginnings and ends of syllables, it may present special difficulties for the learner. These learners may drop consonants, replace one consonant with another, or add vowels to break up consonant clusters. Another potential problem is a difference in the range of phones that can appear in a particular position, for example, the vowels in unstressed syllables. If the first language is more constrained, the learner may tend to follow those constraints in speaking the second language.

Beginnings of syllables in English and Japanese

As we saw in the section on syllables, English allows a variety of consonant clusters at the beginnings of syllables, while some other languages do not. Except for clusters ending in the semivowel /y/, Japanese allows no more than one consonant at the beginning of a syllable, and Japanese speakers may tend to insert vowels between the consonants at the beginnings of English syllables, for example, [gu'rɛet] for *great*. Japanese, on the other hand, permits more consonant+/y/ clusters than English does. In particular, Japanese allows syllables to begin with /ry/ ([ɾy]), whereas /ry/ is not a possible syllable beginning in English. English speakers learning Japanese tend to insert a vowel between /r/ and /y/. That is, they may pronounce the words /ryuu/'dragon' and /riyuu/ 'reason' the same way, [ri'yuu], and may also be unable to hear the difference between these words.

Frequent [ə] is a Feature of English-accented Japanese or Spanish

Unstressed vowels in English and other languages

English has a strong tendency for unstressed vowels to be pronounced with the vowel /ə/. This is not true for many other languages with stress, for example, Spanish. This means that English speakers will tend to use /ə/ for vowels in some unstressed syllables in Spanish. For example, they may say [ɛnčə'lɑdə] for [ɛnči'laða] *enchilada*. In word-final position, English does commonly have /i/ in unstressed syllables, for example, in *happy* and *hurry*, and English speakers will sometime replace /ɛ/ or /e/ in final unstressed syllables with /i/. We sometimes see this in the pronunciation of words borrowed from Japanese into English. Even though Japanese does not have stress, English speakers impose stress on words borrowed from Japanese, treating some syllables as stressed and others as unstressed. So the word *karaoke* is usually pronounced /kɛri'oki/ in English, whereas the original Japanese word was [kaɾaokɛ].

Suprasegmental Differences

English and Japanese Use Pitch in Very Different Ways, Leading to Difficulties for Second-language Learners

As we saw in the section on syllables, languages can use suprasegmental dimensions, especially pitch, in very different ways. It is used as one component of stress in some languages and as a signal for tone or pitch accent in other languages, and it is the major component of intonation in all languages. English and Spanish speakers have little difficulty learning stress in each others' languages, but they may have great difficulty learning pitch accent in a language like Japanese or tone in a language like Mandarin Chinese or Lingala. The situation becomes more complex because of the way intonation interacts with stress, pitch accent, and tone. Thus English uses a sharply falling intonation pattern on words that are being emphasized in statements, for example, when contradicting the hearer, as in the following conversation.

A: *That cookie looks good.*

B: *It's not a cookie; it's a **candy**.*

The first syllable of *candy* would be pronounced more loudly and with the pitch falling to the next syllable:

caⁿdy

In Japanese, the second part of B's line would be

<i>ame</i>	<i>da</i>	<i>yo</i>
candy	is	assrt
'It's (a) candy.'		

The word *yo* at the end of the sentence makes the sentence more assertive.

The word *ame* means 'candy', but its lexical pitch pattern is one that rises from the first to second syllable: *ame*. In this emphatic context, the word would be pronounced more loudly, but the pitch pattern would not change because it is part of the word, though the pitch rise might be more exaggerated:

me da yo
a

'It's **candy**.'

An English speaker learning Japanese will tend to use the characteristic falling English intonation here to signal the emphasis on *ame*:

a
me da yo

But *ame* with pitch falling from the first to second syllable is a different Japanese word, meaning 'rain', so the English speaker's Japanese would come out as 'it's not a cookie; it's rain'. Here's what the last part of that sentence sounds like in Japanese:

a
me da yo

'It's **rain**.'

Japanese does not have many minimal pairs such as this differing only in their pitch pattern, but pitch errors of English-speaking learners could still be expected to make their language less comprehensible, especially when combined with other errors at the level of individual phones.

This page titled [4.4: Learning Phonology](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.5: English Accents

In this section we'll look at various English accents and how they differ from one another. The point, aside from learning some important aspects of the English language, is to get a deeper understanding of the concepts introduced in this chapter and the last one by comparing several similar phonological systems. We'll see that the ways in which the accents differ correspond to the kinds of knowledge about linguistic form that we've been discussing: the form that particular words take, the distinctions between phonemes, the detailed realization of individual phonemes, the allophones of phonemes that appear in different contexts, and phonotactics.

Accents Revisited

Let's first remind ourselves what an accent is: the set of pronunciation conventions of some speech community. Where we draw the boundaries between accents is pretty arbitrary; if we call General American a single accent, for example, we'll have to deal with the range of variation that exists among speakers within that large community. And any boundaries we draw will be wrong in another sense because the group of people who have one pronunciation convention may not coincide neatly with the group of people who have the other set of conventions that belong to the accent we're considering. For example, the group of speakers who pronounce the words *pin* and *pen* the same includes speakers of Southern US accent but also some speakers of General American, which is in many other ways a very different accent from Southern. The point is that conventions of pronunciation *tend* to cluster together; this is what allows us to talk about "accents" at all.

Another point to keep in mind is that in most countries there is a standard, prestige accent alongside a number of accents associated with particular regions, social classes, or ethnic groups. Each of these non-standard accents can be described in its "broad" form, the form that is most different from the standard in the country where it is spoken, but what many people are speaking much of the time is something *in between* a particular non-standard accent and the relevant standard. In this section I'll be mostly concerned with broad variants of non-standard accents because they illustrate the range of possible differences best.

Ways to Talk about Differences Between Accents

When comparing two dialects or accents, one possibility is to see one of them as deviating from the other. A biased view of non-standard dialects often starts this way: the speakers of these dialects are seen as just making mistakes with the standard when what they say is non-standard. But of course this is not what is actually happening. Speakers of non-standard dialects learned the conventions of these dialects by hearing other speakers speak them, just as the speakers of standard dialects learned the conventions of their dialects. They are no more speaking the standard wrong than the speakers of the standard dialect are speaking their dialect wrong.

But there is one situation in which it does make sense to speak of dialect A as diverging from dialect B. Dialects of a single language always started out as a single dialect at some point in the past, and for a given convention, one of the dialects may have changed while the other preserved the original convention. Some people seem to have the sense that standard dialects are conservative, and that non-standard dialects are more likely to change, that is, to introduce "innovations". Sometimes this does happen. In fact some of the conventions that eventually become standard started out as innovations in non-standard dialects. We can see this process going on in England now as features from London accent are starting to creep into the speech of people in situations where we'd expect the standard accent of England. But it seems just as likely that the old conventions get lost in a standard before a non-standard. In North America, the distinction between /ɔ/ and /ɑ/ is in the process of disappearing in General American and (standard) Canadian English, while this distinction is maintained (conservatively) in all of the major non-standard dialects. Of course if we are not interested in history when comparing two dialects, which is more conservative doesn't really matter, and we can just treat the dialects as different from one another.

Overview of English Accents

Before looking at examples of differences between accents, it might help to have a sense of what the major accents are and where they're spoken. But you can safely skip this subsection if you prefer.

The British Isles

There is no "British" accent. England, Scotland, Ireland, and possibly Wales all have their own unofficial standard accents, and the standards of Scotland and Ireland in particular are as different from that of England as American accents are. The standard, or prestige, accent of England is usually referred to as Received Pronunciation (RP). This is what the royal family, all recent Prime Ministers, and most BBC announcers speak. It is probably what most Americans think of as an "English" accent, though it is spoken as a native accent by no more than about 10% of the English population. It differs most noticeably from General American in the pronunciation of a few vowels and in the way /r/ is treated following vowels. For example, in RP there would be no [r] sounds at all in the phrase *the northern fourth of the park*.

Within England there are many identifiable regional accents, probably more than in the United States in fact. Among these, London accent (sometimes called "Cockney") stands out because it is familiar to many Americans through film and drama characters such as Eliza Dolittle in "Pygmalion/My Fair Lady" and because it has a number of very characteristic features. Many of the vowels in this accent differ considerably from RP (and General American). Other very striking features are the loss of initial /h/ ("e 'as an 'ard 'eart" = "he has a hard heart") and the frequent glottal stops in place of other stops in other accents ("i?ll ta?e a lo? o' time to se?le" = "it'll take a lot of time to settle"). Perhaps the other major accent boundary in England separates the accents of the north from those of the south. Americans may be familiar with northern England English through the speech of the Beatles or the characters in films such "The Full Monty". These accents can be identified fairly easily because they make no distinction between the vowels /ʌ/ and /ʊ/; both are pronounced like /ʊ/.

Scottish and Irish English share one feature with northern England English; the tense vowels /i/, /u/, /e/ and /o/ are not pronounced as diphthongs, as they are in RP (and General American). In addition, these accents are like General American, and unlike most accents of England, in how they treat /r/ after vowels.

The Western Hemisphere

The unofficial standard accent of the United States is usually called General American (GA). This is the accent of much of the Midwest and the West and the most frequent accent for US newscasters, though, interestingly, only five of the last ten US Presidents have spoken it. As the prestige accent, it has been encroaching on some regional accents, for example, in the northeast, but at the same time, changes within GA are creating what amount to new accents. One striking example of this is Northern Cities accent, spoken in cities such as Chicago, Detroit, Cleveland, and Rochester, and distinct from GA in the pronunciation of lax vowels.

Almost everyone is familiar with Southern US accent, spoken by people mainly in the southeastern part of the country. Like London accent, this accent has strikingly different vowels from other English accents. African-American Vernacular English (AAVE) is a dialect associated with an ethnic group rather than a region, though of course you don't have to be African-American to have learned it. The accent associated with this dialect is similar in many ways to Southern US accent, while the grammar has its own characteristic properties.

People from the northeastern US are often easy to identify by their accents; the accent of New York City stands out within this region, again mostly for its vowels. Some other US cities, especially Pittsburgh, are known for particular pronunciation conventions. In Pittsburgh, for example, [a] may be used where GA has /aw/, so *downtown* may be [dan'tan].

Standard Canadian English (except in the province of Newfoundland) is very similar to General American, and it doesn't vary much from place to place. Two features that can help identify Canadians are their pronunciation of /ay/ and /aw/, which we'll learn about later, and a tendency to use rising pitch at the end of some statements as well as questions.

English is the native language of much of the Caribbean, with some features common to the region and others specific to particular islands. Americans may be familiar with Caribbean English through the speech of Jamaican performers of reggae music. As with other accents, there are characteristic vowels in these accents, and in addition, a tendency in the Caribbean, as there is in some US accents, to make no distinction between /t/ and /θ/ and between /d/ and /ð/. Jamaican English in particular also has quite striking intonation patterns.

The Southern Hemisphere

English is the native language of most Australians and New Zealanders and a sizable minority of South Africans. While the standard English accents of these countries tend to approach RP, the broad accents of most English speakers in all three countries

have tense (long) vowels similar to those in London accent. The lax (short) front vowels of Australian and New Zealand English differ from those in other accents. Americans are likely to be familiar with these features from the speech of actors such as the Australian Paul Hogan.

Non-native Accents

English is spoken as a second language by millions of people, especially in regions that were once colonized by Britain in South Asia and Africa. In some of these regions there are particular English pronunciation conventions that derive from the phonology of the local languages. So in the English of South Asians (Indians, Pakistanis, Bangladeshis, Sri Lankans, Nepalese, Bhutanese, and Maldivians), the alveolar consonants /t/, /d/, /n/, and /l/ tend to be replaced by retroflex consonants, an important place of articulation for consonants in the languages of this region. Some of these conventions may be viewed as belonging to a kind of non-native regional or national English standard. These non-native standards are one of the ways in which English is becoming even more of an international language.

Phonetic Differences

Exercise 4.5.1

You learn that the phoneme /e/ is pronounced [yɛ] in Jamaican English. As a speaker of General American, how easy would it be for you to master this aspect of a Jamaican accent?

Probably the most common sort of difference between accents is purely *phonetic*. A phoneme in one accent corresponds perfectly to a phoneme in another accent, so we can consider it to be the *same phoneme*, but it differs in its precise realization, that is, how it is articulated and perceived.

How to Make Your /o/s Sound English or Irish or Scottish


Take the vowel /o/. In GA, this is pronounced as a diphthong beginning as a rounded mid back vowel and ending as a rounded high back vowel (or semivowel): [oʊ]. In RP, on the other hand, this same phoneme has a slightly different realization. It begins as an *unrounded* mid *central* vowel and ends as a rounded high back vowel: [əʊ]. In other accents, such as Irish, Scottish, and northern English, /o/ is not a diphthong at all; it is realized as [o]. But since the set of words in GA with [oʊ] is the same as the set of words in RP with [əʊ] (with perhaps a few exceptions) and the other accents with [o], we can see these as the same phoneme. If you're a speaker of GA, and you want to sound English, one thing you could do would be to simply pronounce all instances of /o/ in your speech as [əʊ], just as a speaker of RP could pronounce all instances of /o/ as /oʊ/ as a part of affecting an American accent.

Another similar example concerns the vowel /ɑ/, as in the words *hot*, *sock*, and *rob*. In RP, this vowel is pronounced in roughly the same position as it is in GA (that is, with the same height and backness), but in RP it is somewhat rounded (leading some Americans to think that the RP vowel in these words is /ɔ/). Sometimes a different symbol is used for the vowel in fact. But this difference between the accents is a bit more complicated than this because, as we'll see below, it applies to only some instances of /ɑ/ in GA.

For English vowels, the pattern of phonetic differences between accents is often more extensive than just the correspondences between individual phonemes. The realization of a number of vowel phonemes in one accent may correspond to different realizations for all of those phonemes in another accent. This may be true for the lax ("short") vowels or the tense ("long") vowels or both.

Let's compare the tense vowels, including the diphthongs, of GA and London accent. These accents have the same set of tense vowel phonemes, which we have been writing using the symbols /i, e, u, o, ay, aw, ɔy/, but each is realized differently in the two accents, in some cases, very differently. The first figure below summarizes what you already know about the tense vowels of GA. Recall that each of these vowels is actually pronounced like a diphthong, though this is not reflected in the symbols used for /i, e, o, u/ and it may be difficult to hear or feel for /i/ and /u/. Each line represents one of the phonemes, and it is labeled with a word containing the phoneme (written in the same color as the line). Circles at one end or the other of an arrow represent rounding, and the arrows next to the words show the direction of the diphthong. The second figure shows the corresponding vowels in London accent. Click on the words to hear my imitation of a Londoner saying them. The colors in the two figures represent phonemes that correspond, though in some cases they differ considerably in their phonetic realization, as you can see.

How "phoned Ray" in London Sounds like "found rye" in Indianapolis



The main point to note here is that there is a clear correspondence between the GA and London vowel phonemes, even though the correspondences might not be reflected in the symbols that we use to represent the phonemes. For example, we have been using /e/ to represent the vowel in *bait*, but [e] is very far from the vowel in this word in London; for London accent, a better symbol would be [aɪ], which of course is the realization of a *completely different phoneme* in GA, the one in the word *bite*. So when we are talking about two phones in different accents, there are two ways we can compare them, phonetically and phonemically/lexically. Phonetically, the vowel [aɪ] in London, as in the word *bait*, is quite similar to the vowel [aɪ] in GA, as in the word *bite*. But phonemically or lexically, the vowel [aɪ] in London functions the same way as the vowel [eɪ], that is, /e/, in GA.

Phonemic Differences

Exercise 4.5.2

You're a young speaker of a Caribbean accent in which there is no /θ/ or /ð/ phoneme (*thing* is /tɪŋ/; *this* is /dɪs/). When you start school, you're expected to learn a prestige accent in which distinctions are made between /t/ and /θ/ and between /d/ and /ð/. In what ways might this be difficult for you?

Another possibility is that two accents may differ in the number of phonemes. That is, a distinction that is made in one accent and used contrastively is not made in the other accent. This means that some words that contrast in one accent may sound the same in the other accent.

I have already mentioned two examples of this phenomenon. Many, perhaps most, speakers of GA and Canadian English do not make a distinction between the phonemes /ɑ/ and /ɔ/; they have a single phoneme instead. The actual phonetic character of the sound varies somewhat; it is more like [ɑ] for Americans but more like [ɔ] for many Canadians. The point is that the speakers do not distinguish words from one another using a distinction between [ɑ] and [ɔ]. Pairs of words like the following are distinct in other English accents, but they sound the same for these speakers.

1. *awed*, *odd*
2. *dawn*, *Don*
3. *cawed*, *cod*
4. *caught*, *cot*

In these accents, there is a distinction between /ɑ/ and /ɔ/ before /r/, for example, in pairs such as *car* and *core*, *part* and *port*, *lard* and *lord*. But in this same context, there is no distinction between /ɔ/ and /o/, so we could consider a word such as *core* to be /kor/ rather than /kɔr/.

Because these pairs of words sound the same in this accent, there is a potential problem for the hearer that does not exist in an accent where the distinction is made. We know that this feature of this accent is relatively new; that is, the earlier distinction made in this and other accents has been lost (and is apparently being lost by more and more speakers). Given the problem that hearers have distinguishing words like the pairs above, how can such a change take place? In fact it turns out that there are very few such pairs. The additional burden on the hearer is apparently small enough that the loss of the distinction is tolerated by speakers and hearers of the accent.

Another example was also mentioned earlier, the lack of a distinction between the phonemes /ʌ/ and /ʊ/ in accents of northern England. For these speakers there is a single phoneme, normally pronounced [ʊ]. So the following words, which sound different in most other accents, are pronounced in the same way by these speakers.

5. *cud*, *could*
6. *buck*, *book*
7. *luck*, *look*
8. *putt*, *put*
9. *stud*, *stood*

In this case, unlike that of /ɑ, ɔ/ in North America, it is the accent that fails to make the distinction that is more conservative; Middle English did *not* make a distinction between /ʌ/ and /ʊ/. In any case, as before, the lack of a distinction does not leave hearers for this accent handicapped because there are not many pairs of words distinguished only by this difference.

As a final example of phonemic differences in English vowels, consider how the vowel /ɑ/ in GA corresponds to vowels in RP. In GA this vowel appears in words where it is spelled "o" — *hot*, *shock*, *stop* — and words where it is spelled "a" — *father*, *part*, *carve*. In RP, on the other hand, these sets of words have different vowels, a short, rounded, low, back vowel in the first set (which I'll write with /ɑ/ even though it differs a little from GA /ɑ/) and a long, unrounded, low, back vowel in the second set (which I'll write /ɑ:/). That is, in RP, the words *father* and *bother* do not rhyme. Most of the words with /ɑ:/ in RP have an /r/ in GA that does not appear in RP, as we'll see below. This means that, even though RP has two vowel phonemes where GA has one, there are few if any words that are distinguished in RP but not in GA. For example, a pair such as *pot* and *part* is distinguished by the consonant (/r/) in GA but by the vowel in RP (GA: /pɑt/, /pɑrt/; RP: /pɑt/, /pɑ:t/).

But RP /ɑ:/ also corresponds to many words that have /æ/ in GA. Some words are pronounced with /æ/ in both accents, for example, *gas*, *bad*, *can*, and *lamp*. Other words with /æ/ in GA are pronounced with /ɑ:/ in RP, for example, *glass*, *rather*, *can't*, and *laugh*. Note that there is in general no way to predict from the context which words will have /æ/ and which will have /ɑ:/ in RP. Thus an American trying to imitate an RP accent will have to remember which words have which phoneme. This is difficult and leads to frequent over-generalization mistakes such as the pronunciation of *gas* as /gɑs/ or *stand* as /stɑnd/ in American attempts at imitating RP.

Differences in the number of English consonant phonemes are not as common, but there are some. In a number of accents, especially in the Caribbean, in London, in AAVE, and in some US cities, the dental fricatives /θ, ð/ do not exist as separate phonemes. Where other accents have these phonemes, these accents have either /t/ and /d/ or /f/ and /v/. So in accents where /θ/ is not distinguished from /t/ and /ð/ is not distinguished from /d/, each of the following pairs of words would sound the same.

10. *tin*, *thin*
11. *tie*, *thigh*
12. *boat*, *both*
13. *true*, *through*
14. *tread*, *thread*
15. *den*, *then*
16. *ride*, *writhe*

In these accents an older distinction has been lost, but as with /ɑ/ and /ɔ/ in North America, the loss apparently does not seriously interfere with communication because there are not too many pairs of words that end up as homophones with the loss of the distinction.

Going From an Accent with Fewer to an Accent with More Distinctions is Difficult

Going back and forth between two accents is more complicated when the number of phonemes differs than when there are only phonetic differences. Say a speaker of GA or Canadian English who does not make the distinction between /ɑ/ and /ɔ/ wants to learn or to imitate the speech of someone from London or New York or Houston, all places where the distinction between these two phonemes is made. The problem is that words with these phones in this person's mental lexicon are all represented in terms of one vowel category, whereas the same words are represented in terms of two different categories in the mental lexicons of speakers of other accents. For each word, say, *caught* or *hawk* or *hot* or *lock*, the speaker will have to figure out which vowel in the other accent is appropriate. But unless the speaker has learned this for each word, this will be impossible. In this situation, speakers often make mistakes, over-extending either one or the other phone. For example, a North American speaker might overuse /ɔ/ in trying to speak with an RP accent, using this vowel for words like *hot* and *lock*. In the same way, a speaker from northern England trying to speak with an RP accent, might over-extend the vowel /ʌ/, using it for words normally containing /ʊ/ such as *sugar* and *cushion*.

Note that a speaker going in the other direction would not have the same problem. A speaker of RP would just have to remember to pronounce both /ɑ/ and /ɔ/ in the same way when imitating GA and to pronounce /ʌ/ and /ʊ/ in the same way when imitating an accent of northern England.

Allophonic Differences

A further possibility is for two accents to differ in the way a phoneme is realized in different phonetic contexts. The allophones of the phoneme may differ, or the contexts in which they apply may differ. Let's consider /t/ again, a phoneme with a wide range of allophones in many English accents. In the context where it is surrounded by vowels and does not begin a stressed syllable, this phoneme is realized as a tap, [ɾ], in GA, Canadian, Australian, and New Zealand accents, for example, in words like *butter*, *settle*, and *city* and phrases like *put it on* and *at a glance*. Speakers of most accents of England never use this allophone, however. Instead, some of them, especially London speakers and others whose speech is under the influence of London accent, use a glottal stop, [ʔ], in this same context. The glottal stop is a possible allophone of /t/ in GA, but only in the context where it follows a vowel and precedes a consonant, for example, in *outright chaos* and *let me go*. In these contexts, even more speakers in England also have glottal stops.

Another place where English accents often differ with respect to their allophones is the pronunciation of /l/. All English accents have an /l/ phoneme, but it may be realized differently. In many accents, including GA and RP, there are two allophones, a "clear" one, for which the tongue body is pushed up and forward, and a "dark" one, for which the tongue body is pulled backwards. Most GA speakers use the clear allophone only when it comes at the beginning of a syllable (possibly at the end of a consonant cluster), for example, in *live*, *play*, and *relate*. (For comparison, here's how these words would sound with very dark /l/s: *live*, *play*, *relate*.) In other contexts, for example, in *full*, *old*, and *silly*, GA speakers use the dark allophone. (For comparison, here's how these words would sound with very clear /l/s: *full*, *old*, *silly*.) (Note that this description of the contexts requires that we treat the syllable boundary in *silly* as occurring between the /l/ and the /i/.) In RP, the clear allophone is generally even clearer than in GA (a phonetic difference), and it is used in word such as *silly* as well, though not in *full* and *old*. That is, in RP, we can say that the clear allophone of /l/ occurs generally before vowels. In other accents, one or the other variant of /l/ may be used always. In Irish English and Caribbean English, the clear /l/ tends to appear in all contexts, while in Scottish English, the dark /l/ tends to appear everywhere. So in these accents we can say that /l/ has only a single allophone.

Long Lax Vowels Give Southern US Accent Much of Its Characteristic Sound

Accents may also differ in their vowel allophones. The Southern US accent has unusually complex lax ("short") vowel phonemes. In fact "short" is not at all appropriate for this accent since these vowels are often longer than the "long" vowels. In particular, each of the front lax vowel phonemes, /ɪ/, /ɛ/, and /æ/, has a wide range of possible realizations, depending on the place of articulation of the following consonant, the backness of the next vowel, and whether the vowel is in a word consisting of one syllable. Each of these vowels has at least one diphthong allophone. Listen to the vowels indicated in bold in the following sentence:

17. *Deb lives in the lab.*

Each is realized as a diphthong. When these same vowel phonemes appear before /k/, however, they have (non-diphthongal) allophones close to the GA vowels.

18. *Mick wrecked the jack.*

Since all of these allophonic differences concern realization and not how words are represented in the lexicon, they are more like the phonetic than the phonemic differences discussed above. If a GA speaker wanted to imitate a London speaker's use of the glottal stop allophone of /t/, they would only have to worry about what context each /t/ occurred in, using a glottal stop whenever the /t/ occurred after a vowel and not at the beginning of a stressed syllable. Words in the GA speaker's lexicon with /t/ would also have /t/ in the lexicon of a Londoner, so there would be no need to remember new properties of individual words.

Phonotactic Differences

Exercise 4.5.3

Some accents (such as AAVE and Caribbean) do not permit consonant clusters such as /st/ and /nd/ at the ends of words, while other accents (such as GA and RP) do. With respect to just this property, would you expect it to be easier for a speaker of AAVE to learn GA or for a speaker of GA to learn AAVE?

The Behavior of /r/ After Vowels is One of the Main Features Distinguishing English Accents from One Another

Accents can also differ from one another in their phonotactics, that is, in the way in which consonants and vowels combine to make syllables. The most noticeable place in English where there is this sort of variation is in the distribution of the phoneme /r/. In most accents, /r/ can occur freely at the beginnings of words, both alone (*rat*) and in clusters (*brat*). In GA, however, there are restrictions on which vowels can occur before /r/ within a word. For many speakers, the vowels /i, e, u, o, ʌ/ cannot occur before /r/. In RP (and many other accents, including AAVE and some southern and northeastern US accents), the restrictions are even more severe: /r/ can only occur before a vowel. That is, in words such as *harm*, *port*, *hurt*, and *weird*, there is no [r] sound at all in these accents. And in words such as *car*, *pour*, *her*, *hair*, *here*, *fire*, and *power*, there is no [r] sound unless the words immediately precede a vowel in the following word. So the following sentence has no [r] sounds in RP, imitated here by me.

19. *Arnold carelessly poured the sour yogurt on the fire.*

In fact the situation in RP is more complicated than this. First, where there used to be /r/ after a vowel, and where there still is today in most accents, RP sometimes has a [ə], resulting in the diphthongs /ɪə/, /ɛə/ and /ʊə/. These diphthongs are considered to be separate vowel phonemes in some descriptions of RP. There are also pairs of homophones in RP that differ in GA and other accents, for example, *farther/father* and *source/sauce*. One further complication is that [r] gets *inserted* in some contexts in RP. But I'll save this for a later section.

Another example of restrictions on the distribution of a phoneme concerns the vowels /ɪ/ and /ɛ/ in Southern US English (and also the English of some neighboring regions), which we've run into before. In these accents, the distinction between /ɪ/ and /ɛ/ is neutralized before [n]; only /ɪ/ occurs in this context. So in these accents, the following pairs of words are pronounced the same.

20. *pin, pen*

21. *tin, ten*

22. *since, sense*

23. *sinned, send*

24. *mint, meant*

As with other cases of the loss of a contrast, there is the possibility of a burden on the Hearer because of the words that are no long distinguished. This loss may explain why speakers of these accents seem to replace *pen* and *pin* with longer expressions such as *fountain pen*, *ballpoint pen*, and *straight pin*. Note that *pen* and *pin* are both nouns referring to physical objects, so there might be some confusion on the part of a hearer for these two words.

Finally, English accents may differ in what sorts of consonant clusters are possible. AAVE and Caribbean accents, for example, have more restrictions on what can occur at the ends of words than GA or RP do. For example, word-final consonant clusters ending in /p/, /t/, /k/, or /d/ in other accents are simpler or different in these accents: *wasp* /wɒs/, *list* /lɪs/, *left* /lɛf/, *act* /æk/, *desk* /dɛs/, *ask* /æks/, *find* /fayn/, *cold* /kol/, *loved* /lʌv/.

Lexical Differences

Exercise 4.5.4

In the United States some words have pronunciations that are stigmatized; some of the pronunciations are associated with particular regional accents. Examples are *get* pronounced /gɪt/ and *once* pronounced /wʌnst/. If you spoke an accent that included these pronunciations, how easy do you think it would be to learn to pronounce such words in the standard (prestigious) way?

A final way in which accents can differ is lexically, that is, in the conventional pronunciation for particular words. For example, a small number of words are pronounced differently by GA and RP speakers (and most other speakers on the two sides of the Atlantic). Examples include *schedule* (GA: /ˈskɛʃəl/, RP: /ˈʃɛʃəl/), *tomato* (GA: /təˈmeto/, RP: /təˈmɑːto/), *laboratory* (GA: /ˈləbrəˌtɔːri/, RP: /ləˈbɔːrətəri/), *lieutenant* (GA: /luˈtɛnənt/, RP: /lɛfˈtɛnənt/), *figure* (GA: /ˈfɪɡyər/, RP: /ˈfɪɡə/, *beta* (GA: /ˈbetə/, RP: /ˈbitə/). Note that none of these differences is related to any more general difference between the accents; that is, it could not be predicted from what we know about the phonetic, phonemic, allophonic, or phonotactic differences between GA and RP.

There are also lexical differences between GA and other US accents, especially broad Southern and AAVE. Some examples are *police* (/pəˈlis/ vs. /ˈpolis/), *wash* (/wɑʃ/ or /wɔʃ/ vs. /wɔtʃ/), *yellow* (/ˈyɛlo/ vs. /ˈyælər/), and *catch* (/kæʃ/ vs. /kɛʃ/).

Other lexical differences may apply to larger sets of words. One way this can happen is with the pronunciation of *parts* of words that recur in many words. The word ending spelled *-ile* occurs in many English words such as *fertile*, *docile*, *mobile*, and *sterile*. This ending is normally pronounced /əl/ in GA but /aɪl/ in RP. A more familiar example is the pronunciation of the word ending spelled *-ing* in words of more than one syllable, that is, words like *playing*, *eating*, and *something*, but not *thing* and *sing*. Some speakers, located in various places in the English-speaking world, always pronounce this ending /ən/. Many other speakers have two pronunciations for the ending, one reserved for more formal situations, the other for more informal situations. For RP speakers and many GA speakers (including me), the formal pronunciation is /ɪŋ/ and the informal pronunciation /ən/. For other GA speakers (a group that appears to be growing), the more formal pronunciation is /ɪn/.

Even though these differences in pronunciations of word parts such as *-ile* and *-ing* apply to large numbers of words, they still need to be seen as lexical differences since there is nothing in what we know about the phonetic, phonemic, allophonic, or phonotactic differences between the accents that would allow us to predict the different pronunciations from the contexts of the phones.

Learning to change your pronunciation of particular words, for example, if you want to make your pronunciation more standard or want to imitate a different accent, is not too difficult, as long as there are no differences of other types (phonetic, phonemic, allophonic, phonotactic) to worry about. On a word-by-word basis, you just have to remember the new pronunciation in terms of the phonemes that are part of your native accent.

Suprasegmental Differences

English Accents Sound More Similar When They're Sung Because Pitch Differences are Lost

In the section on syllables, we saw that languages vary in terms of how they use the dimensions of pitch, loudness, and duration. One very noticeable difference between English accents is in the details of how these dimensions interact with the structure and the meanings of sentences, that is, in their intonation. Because of the inherent difficulty of describing intonation, however, these accent differences are apparently not as well studied as differences at the level of consonants and vowels.

One difference between GA and RP is in the typical pitch pattern used for *yes/no* questions, that is, questions that can be answered with *yes* or *no*, rather than with phrases like *Felix* or *on Tuesday*. In GA, the usual pattern for these questions involves a pitch rise on the stressed syllable of the last stressed word in the sentence followed by a continuing high pitch on succeeding syllables. In RP, the syllables leading up to the stressed syllable of the last stressed word in the sentence are relatively high, and the pitch falls on that stressed syllable and then rises to a high pitch again, remaining high for the rest of the sentence as in GA. Consider the following question in the two accents.

1. Has she writ^{ten} to you ?
2. Has she writ^{ten} to you ?

This difference in intonation is similar to phonetic, rather than phonemic, differences at the level of consonants and vowels because it does not involve more distinctions made in one or the other accent. It is just the *realization* of the *yes/no* pattern that differs for the two accents.

We've seen that accents within a language can differ in all of the characteristic ways that languages differ from one another. What makes accent differences special is the fact that the different dialects are related to one another; they ultimately derive from the same dialect in the past. This means that either phonemes in one accent correspond directly to phonemes in the other accent, or, if phonemes have split or merged in one of the accents, one phoneme in one accent corresponds to multiple phonemes in the other. These correspondences are easy to observe because of the correspondences between words in the different accents. The word *cat* is [k^hæt], [k^hat], [k^hɛət], and [k^hɛt] in different accents (GA/RP, Scottish/Northern English, Northern Cities, Australian/New Zealand), and from the lexical correspondences we can discover the vowel correspondences. One interesting kind of information that we can infer from these sorts of correspondences is the history of the dialects, how the parent dialect turned into the different daughter dialects. Of course the same is true for related languages. That is, we could learn something about the history of Spanish phonology by looking at correspondences between words in modern Spanish and related languages such as French and Portuguese. We'll see more about how this works in the section on phonological change.

This page titled [4.5: English Accents](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.6: Phonological Change

Phonetic Change

Exercise 4.6.1

What do you think would be the consequences if, for some reason, the phoneme /i/ in English started to become more and more similar to the phoneme /ɪ/? Think in terms of what this would mean for the Hearer. How might the English phonological system deal with such a change?


One way the pronunciation of a language can change over time involves changes in how particular phonemes are pronounced but not in the number of phonemes. As with other language change, it is usually not clear how the change begins, but the prototypical phone for some phoneme starts to move. In the simplest case, this is all that happens. For example, about 150 years ago the vowel /o/ in some dialects of English, already a diphthong, began to shift so that its beginning was more central and less rounded, resulting in the characteristic [əʊ] of today's Received Pronunciation.

More often a change in one phoneme affects others. This is because the change may either make that phoneme more similar to another or open up a region in the phonetic space where there is no phone. In the former case, the changing phoneme may "push" another phoneme away as it comes close to it. In the latter case, the changing phoneme may "pull" another phoneme into the region where it used to be. Both kinds of changes favor the Hearer because they keep the phonemes as far apart as possible. These processes are best known from the history of vowel systems. The vowels of English have undergone several such changes and in some English dialects are undergoing them now.

Sometimes a Whole Set of Vowels Will Shift in the History of a Language

Probably the most famous example of such a set of changes is the Great Vowel Shift of Middle English. I won't go into it in detail, but what happened was that the pronunciation of all of the tense (long) vowels of English changed, in some cases quite dramatically. For example, the vowel in words such as *fine* had been pronounced [i]; eventually it became the [ay] of Modern English.

Instead we'll look in detail at some changes going on in the vowels of one English dialect today. A quite striking set of changes is happening in some cities of the Northeast and the Midwest in the US (for example, Rochester, Cleveland, Detroit, Chicago, and Milwaukee). This is called the Northern Cities Vowel Shift; in its most extreme cases, it applies to the set of all of the English lax (short) vowels except for /ʊ/. You can read more about the Northern Cities Vowel Shift in this paper. The figure below diagrams the changes.



As a Result of the Northern Cities Vowel Shift, Some Words in Chicago Sound Like Different Words in Indianapolis

The diagram shows what is happening in vowel space (front vowels on the left). Each arrow indicates the direction of change for one phoneme. The phoneme label and example word appear in the position where the vowel started, that is, its position in General American. The end of each arrow shows where the vowel ends up in the cases where the shift has progressed the furthest. For example, the word *cot* following the shift is pronounced something like the word *cat* in General American. The order of the changes is indicated by the numbers. Different speakers, and to some extent different cities, can be seen as being currently at different points within the set of changes. For examples, for some speakers, only changes 1 and 2 might have taken place, whereas for others all of

the changes might have taken place. There is also considerable variation, so we should not expect everybody in the Northern Cities to follow exactly this pattern.

The first change to happen was a movement of the /æ/ vowel higher. (Incidentally a similar change has happened in other accents of the US Northeast, but it is normally confined to only some contexts, for example, in *glad* but not *back*.) As the vowel moved higher, it also tended to become a front-to-central diphthong. The diagram shows the most extreme change; more moderate changes occurred within the speech of many speakers. Note that once /æ/ has shifted like this, it is the same phoneme as it was in the sense that it is still distinguished from all of the other vowel phonemes in this English dialect and is still used for the same set of words as before (*back*, *glass*, *fancy*, etc.). But phonetically it is no longer [æ], so we could choose to represent the phoneme with a different symbol. Keeping the symbol the same, however, reminds us how this phoneme corresponds to the /æ/ of other dialects.

Apparently the next change to take place was the movement of /ɑ/ forward. This is an example of a "pulling" change; the movement of /æ/ left a gap in the vowel space that /ɑ/ moved to fill so that the vowels remained roughly equally spaced. Again for most speakers the change was not as dramatic as shown in the figure; for many the vowel is closer to [a] (a low central vowel) than to [æ].

The next change to take place seems to have been the movement of /ɔ/ downward. As you'll see in the section on English accents, this change has happened for many North American speakers, but in other accents, the /ɔ/ and /ɑ/ vowels are no longer distinguished. In the Northern Cities, the two vowels remain distinct, and the movement of /ɔ/ can be seen as another example of a "pulling" change since the movement of /ɑ/ opened up a place for /ɔ/.

Next the vowel /ε/ became more central. Again this is an example of a change that seems to be occurring more generally in North America, though apparently only in some contexts, for example, in the word *level*. In the Northern Cities accent, it was probably a response to the rising /æ/. As this vowel became higher, it came to resemble /ε/, and there was the potential for confusion since many words in English are distinguished by the distinction between /æ/ and /ε/ (for example, *bat* and *bet*). As a result, /ε/ shifted so that it would be more distinct. This is an example of a "pushing" change; the /æ/ pushed the /ε/ into another region of the vowel space.

Next /ʌ/ became more back. This can be seen as both a pulling and a pushing change, pushing because /ε/ became more confusable with /ʌ/ as it moved back, and pulling because /ɔ/ opened up a gap in the vowel space when it moved down.

Finally /ɪ/ also moved back. This is an example of a pushing relationship. As /æ/ rose, for some speakers it seems to have reached the point where it became potentially confusable with /ɪ/, and /ɪ/ moved back to make room for it. This is also a change that seems to be happening more generally in North America, though again apparently only in some contexts, for example, in the word *liver*.

Change in Some Contexts

Allophones Often Emerge in the History of a Language

Another possibility is that a phoneme will come to be pronounced differently in some contexts but not others. In other words, the realization rules for that phoneme change. Often the changes are examples of assimilation. Here are some examples from the history of English.

Old English /k/ before /i/

In Early Old English /k/ came to be **palatalized** when it occurred before /i/; that is, the point of articulation moved forward from the velar to the palatal region. This is an example of anticipatory assimilation; the /k/ changes to be more like the /i/, for which the high tongue position is near the palatal place of articulation. Eventually /k/ in this context became similar to [č], that is, an alveopalatal affricate.

Final unstressed vowels

At several times in the history of English, final unstressed vowels have been dropped. In Old English, which had no /ə/ phoneme, there were many words such as /sʊnʊ/ 'son' with final unstressed vowels that are quite unlike unstressed vowels in Modern English. At some point, unstressed vowels in words such as these became reduced to /ə/, and later this vowel when it was final ceased to be pronounced altogether. This is the reason English has so many "silent e's"; the orthography has been conservative and fails to represent all of the phonological changes.

Phoneme Loss

A further possibility is that two phonemes will merge as a result of change in one or the other or both. Obviously this can only happen when the difference between the two phonemes is not so significant in the language, that is, when the phonemes do not distinguish many words. In the section on English accents, we'll see several examples of this. In many North American dialects, for example, the vowels /ɔ/ and /ɑ/ have merged in recent years. This does not create a serious problem for the Hearer because there are not many minimal pairs such as *awed/odd* and *caught/cot* that are distinguished only by these phonemes.

Phoneme Creation

Exercise 4.6.2

We saw in the last section how phonemes could be lost. Given what you know about allophones, how might the opposite process take place? That is, how might allophones of the same phoneme (for example, [t] and [tʰ] in English) turn into separate phonemes?

Phonemes are Both Lost and Created, Apparently with Roughly the Same Frequency

If phonemes can be lost, it stands to reason that they can also be created. Otherwise languages would tend to have fewer and fewer phonemes, making them more and more difficult for the Hearer. There are at least two ways that new phonemes can emerge in the history of a language. In both of the ways we'll look at, the phoneme starts as the allophone of an existing phoneme.

One way in which an allophone can turn into a phoneme results from the borrowing of words from another language in which that phone is a separate phoneme already. We have seen that [v] was an allophone of the phoneme /f/ in Old English, not a separate phoneme. But following the Norman conquest of England in the 11th century, English borrowed many (Norman) French words. Some of these words contained [v] (a separate phoneme in French), and some of these were in positions where the [v] allophone of /f/ did not occur, for example, at the beginning of words (*very*). Once [v] was appearing in positions where [f] could appear, that is, once the distributions of [v] and [f] overlapped, it was a separate phoneme in English. After this the distinction between /f/ and /v/ could be used to distinguish words from each other, for example, *fine* and *vine*.

Phonemes may also emerge out of allophones when other changes combine to make the contexts for different allophones overlap. This is what happened in Old English in the case of /č/, originally an allophone of /k/ before /i/, as we saw above. When other changes caused [k] also to appear sometimes before [i], the contexts for [k] and [č] overlapped, and they were now separate phonemes, distinguishing some words from one another. These changes are summarized in the table below, which also illustrates the emergence of another phoneme in Old English, /ü/, a high front rounded vowel. (You are familiar with this phone if you happen to know French, German, Dutch, Swedish, Hungarian, Finnish, or Mandarin Chinese.) The symbol [-] represents vowel endings that are irrelevant for the discussion, and "→" represents a sound change. The table shows what happened over a period of several hundred years. The change is illustrated with two examples, the Old English words for 'kin' and 'chin'.

	'kin'	'chin'	Change
1	[kunni]	[kinn-]	
2	[kunni]	[činn-]	Palatalization: [k] → [č] before [i]
3	[künni]	[činn-]	Vowel fronting: [u] → [ü] when the next vowel was [i]
4	[künn]	[činn]	Deletion: some final vowels dropped
5	[kinn]	[činn]	Derounding: [ü] → [i]

Originally the words for 'kin' and 'chin' began with the same consonant phoneme, realized as [k] in both words. Then, in period 2, in a change already discussed above, /k/ in the context of a following /i/ came to be realized as [č]. At this point [č] existed in the language, but only as an allophone of the phoneme /k/. That is, the allophones [č] and [k] were still in complementary distribution. In period 3, the vowel /u/ came to be fronted in the environment of an /i/ later in the word. This is an example of anticipatory assimilation because the /u/ takes on the frontness value of the following /i/. At this point [ü] was still an allophone of the phoneme /u/, however, since it occurred only in the context of an /i/ in the next syllable. Next, in period 4, some final vowels in the language were dropped. This leaves the [ü] without the context that originally motivated it. In other words, [u] and [ü] now occur in overlapping contexts, and because the distinction between them matters for the meaning of the word, [ü] has become a phoneme in the language. Finally, in period 5, as Old English was changing to Middle English, the phoneme /ü/ was lost, merging with /i/. This made it possible for [k] to occur before /i/, as it once had in the language. But the original change that caused [k] to become [č] in this context hundreds of years before no longer applied. Thus at this point [k] and [č] occurred in overlapping contexts; namely, both could occur before /i/. Since the distinction between [k] and [č] also mattered for the meaning of the word, the two phones had become separate phonemes in the language. In fact, the words for 'kin' and 'chin' already constituted a minimal pair for these two phonemes.

Inferring Phonological Change

Exercise 4.6.3

Say there are two related languages A and B. In A there is a contrast between /t/ and /θ/; in B there isn't. There are two possible histories that could have resulted in this situation, starting from the ancestor language of A and B. What are they?

How can linguists figure out what changes have occurred in the history of a language? Recent changes are not a problem; there may still be older speakers whose speech is a reflection of the period before a change took place. For earlier changes, we sometimes rely on written records, though this presents several problems. First, as we have seen, orthography never does a very good job of representing phonology, especially allophonic differences, and it tends to lag behind, representing earlier pronunciation rather than current pronunciation. Second, some writing systems, such as Chinese characters, which are also used in Japanese, do not represent phonology in any direct way at all. Third, most languages are not written, and languages that are written today were not always written.

Inferring the Pronunciation of an Extinct Language May Rely Mainly on What We Know About its Modern Descendants

For these reasons, written records can never be adequate for a full picture of phonological change. Linguists have developed another technique for inferring the past. If a change has taken place in a particular dialect or language, there are likely to be other related dialects or languages where the change has *not* taken place. So by examining a set of related dialects or languages, it is sometimes possible to infer how some of them have changed and what the dialect or language that is the ancestor of the whole set was like. Consider the following example from English; the table compares the forms of several words in General American and Received Pronunciation with those in the English of Northern England and Ireland.

	GA, RP	E. of Northern England, Irish E.
<i>put</i>	/pʊt/	/pʊt/
<i>look</i>	/lʊk/	/lʊk/
<i>but</i>	/bʌt/	/bʊt/
<i>luck</i>	/lʌk/	/lʊk/

We see here (and more examples would make it even more obvious) that the English of Northern England and of Ireland makes no distinction between /ʊ/ and /ʌ/; there is instead a single phoneme pronounced [ʊ]. We know that these dialects (and many others) share a common ancestor dialect with General American and Received Pronunciation. The problem is that from the data here alone, there are two possibilities for the history of these phonemes. Either the ancestor English dialect made a distinction between /ʊ/ and

/ʌ/, and this distinction was *lost* in the English of Northern Ireland and Ireland, or the ancestor dialect did not make the distinction and the distinction *emerged* in General American and Received Pronunciation. If we look at other dialects, we discover that the distinction is made almost everywhere except in Northern England and Ireland. This means that, if the second alternative is right, the distinction would have to have emerged a long time ago, when all of those other dialects still shared a common ancestor. If the first alternative is right, we would expect the dialects of Northern England and Ireland to be closely related to one another, that is, to constitute a subgroup within English dialects where the distinction was lost.

To figure out which alternative is right, we can look for several other sources of information. We could try to determine from linguistic or other evidence whether the dialects of Northern England and Ireland are closely related. In fact we would discover that they are not very similar, no more similar to each other than either is to RP. Or we could look for evidence from another dialect that we know diverged from the common ancestor of all of these dialects even earlier. Unfortunately this is not very helpful in this case since the evidence is somewhat mixed. Finally we could try to come up with an explanation for how the distinction could emerge, similar to what happened in the case of /ɛ/. The story is too complicated to go into here, but it is possible to see the the split between /ʊ/ and /ʌ/ as beginning with allophonic variation and ending with the present phonemic distinction. This in fact is apparently what happened. That is, it is the dialects of Northern England and Ireland that are more like the ancestor dialect with respect to these vowels and it is the ancestor of the other dialects that changed.

Here is another example. The modern dialects of Japanese, spoken on the main Japanese islands, and Ryukyuan, spoken in the Ryukyu islands of southern Japan, are the descendants of a single language, spoken perhaps 1000 years ago. We can infer what changes have taken place in the different dialects and what the ancestor language looked like by comparing the modern dialects. In the table below are some example words in two of the dialects. We assume that an analysis of the modern dialects has already determined what the phonemes in these dialects are, so phonemic notation is used. The symbol /ϕ/ represents a voiceless bilabial fricative, and long vowels are doubled. Tone patterns are ignored (both dialects use pitch contrastively).

	Standard Japanese	Central Okinawan
'body'	/karada/	/duu/
'bone'	/hone/	/ɸuni/
'branch'	/eda/	/ida/
'breath'	/iki/	/ʔiiči/
'dance'	/odori/	/udui/
'daytime'	/hiru/	/ɸiru/
'fog'	/kiri/	/čiri/
'grass'	/kusa/	/kusa/
'hair'	/ke/	/kii/
'man'	/otoko/	/wikiga/
'open'	/akete/	/ʔakiti/
'organ'	/kimo/	/čimu/
'sash'	/obi/	/ʔubi/
'shoulder'	/kata/	/kata/
'sleeve'	/sode/	/sudi/
'south'	/minami/	/nisi/
'stomach'	/hara/	/wata/
'sun'	/hi/	/tiida/
'rain'	/ame/	/ʔami/
'rock'	/isi/	/ʔisi/
'west'	/nisi/	/ʔiri/
'where'	/doko/	/maa/
'wine'	/sake/	/saki/

For any pair of related dialects or languages, some words with the same meaning will have arisen from the same form, and others won't. Even for very closely related dialects such as General American and Received Pronunciation, we will find pairs such as /'ɛləˌvetər/ and /lɪft/ (*elevator*, *lift*), forms with the same meaning but different origins. When we are interested in phonological change, we should only take into account forms that are obviously related. For the above example, words that we should ignore include the words for 'body', 'man', 'south', 'stomach', 'sun', 'west', and 'where' because the words with these meanings in the two dialects clearly have different origins. On closer examination, we see that the Okinawan word for 'south' is identical to the Japanese word for 'west'. Since these meanings are related, it appears that the origin of the forms is the same, so we can also use this pair for comparison.

The Key to Figuring Out the Phonological History of a Group of Modern Languages is Establishing Correspondences Between the Phonemes in the Languages

Once we have found comparable pairs of words, we need to set up correspondences between pairs of phonemes or combinations of phonemes. When we find differences, we will look more closely to see what changes might have occurred in one or the other

dialect. We will focus here only on the phonemes that differ in the two languages and ignore some details such as vowel length in Okinawan. The table below summarizes these. "V" represents any vowel, and "#" represents the beginning of a word.

Standard Japanese	Central Okinawan
/i, e/	/i/
/o, u/	/u/
/k/	/k, č/
/#V/	/#V, #?V/

In each of these cases, one dialect has two forms where the other has one. As in the example of English /ʌ/ and /ʊ/, the change is either a merging of two phones into one or a splitting of one phone into two. Let's consider the vowels first. Note first that there are two similar patterns: a high vowel in Okinawan (/u/, /i/) corresponds to that same vowel and a lower vowel in Japanese (/u, o/, /i, e/). Because languages tend to be systematic, we would expect whatever holds for one of these to be true for the other as well.

Let's first consider the possibility that the parent language had only three vowels, /a, i, u/, like Central Okinawan, and a change took place in the ancestor of Standard Japanese, resulting in five vowels, /a, i, e, u, o/. As before, the explanation for the emergence of new phonemes is more complicated than the explanation for the merging of phonemes. As we've seen earlier in this section, new phonemes normally begin with an allophone occurring in some contexts but not others. So by this story, [e] would have first appeared as an allophone of /i/ in some contexts. But what contexts? As you know by now, allophonic variation often involves assimilation, where a phoneme agrees with features of preceding or following phonemes. The relevant context for a vowel could include the preceding or following consonant, or perhaps the vowel in the preceding or following syllable. But given the phonemes of these dialects, it is hard to see how any contexts would have led /i/ to be realized as [e].

What about the alternative, that the parent language had both /i/ and /e/, as in modern Standard Japanese, and that the difference disappeared in the ancestor of Okinawan? Here the story is simpler. For some reason, /e/ in this dialect began to rise, and apparently because there were not so many minimal pairs distinguished only by the difference between /i/ and /e/, /e/ merged with /i/, becoming a single phoneme. The same thing would have happened for /o/ and /u/ in this dialect.

Of course this is still a hypothesis. We could strengthen it with data from another dialect that we know diverged from the parent language of the other two relatively early. Unfortunately I'm unaware of any such data.

The Order in Which a Set of Phonological Changes Takes Place May Have Important Consequences for a Language

Now let's consider the case of /k/ and /č/. There are three examples of /č/ in the Okinawan data. In all cases /č/ appears before /i/. This leads us naturally to the hypothesis that /k/ came to be realized as [č] in Okinawan when it appeared before /i/, a process of palatalization (exactly the sort of change we saw for Old English above). But there is a problem with this hypothesis: in the Okinawan word /ʔakiti/, we have /k/ (that is, [k]) before /i/. Notice, however, that this /i/ in Okinawan corresponds to /e/ in Japanese. That is, in all cases where we have /ki/ in Japanese, we have /či/ in Okinawan. The solution is to propose that the two changes in Okinawan happened in a particular order. First, /k/ came to be realized as [č] when it appeared before /i/. At this point [č] might have been just an allophone of the phoneme /k/, and 'open' would still have been /ʔakete/ (or /akete/) in Okinawan. Then *later*, /e/ moved and merged with /i/. If the first change had stopped taking place at this time, we would now have /ʔakiti/ in Okinawan. At this point, since both [č] and [k] could appear before /i/ (and affect the meaning of the word), [č] had become a separate phoneme in the language, that is, /č/.

Note how this proposed sequence of events is similar to that postulated for the history of English. In both cases [č] first emerged as an allophone of /k/ in the context of a following /i/ (palatalization). Later another vowel — [ū] in Old English, [e] in Okinawan — came to be pronounced [i], but the old palatalization rule was no longer in effect so words in which /k/ had preceded this other vowel now had [ki]. Since both [či] and [ki] could now occur, and in different words, /č/ and /k/ had become separate phonemes.

This page titled [4.6: Phonological Change](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.7: Phonology in the Wild

If you've listened carefully to your own pronunciation of English words since you started learning about phonology in this book, you may have noticed that the pronunciation given doesn't correspond to the way you sometimes say the words, that your pronunciation varies with the situation.

In Chapter 1, we already saw that the conventions characterizing a particular dialect can change depending on the context the language is used in. What is appropriate in one context may not be in another. This applies to pronunciation, as well as to vocabulary and grammar.

The dimension we will be concerned with here is sometimes referred to in terms of how "careful" the speech is. The "care" referred to is care on the part of the Speaker. To what extent does the Speaker make an effort to accurately produce each of the phonemes and suprasegmental features of the words? To make sense of this idea, we will have to assume that each word in a dialect has a "careful" or "canonical" pronunciation, that is, how the word would sound (or look in a sign language) if produced in isolation or with some emphasis within a sentence and in a relatively formal setting. In general, as the word gets less emphasis and the setting gets more casual, we find a tendency for Speakers to deviate from the careful pronunciation. These deviations are Speaker-oriented; that is, they can all be seen as making the pronunciation easier in one way or another; they are **simplifications**. Simplification is possible because in the casual situations where it is most common, the Hearer knows the Speaker well and is better able to predict what the Speaker is saying than a stranger would be. In this section we will look at some examples of the simplifications that occur in casual English. We will see that they can often be described in terms of the same sorts of processes we have seen elsewhere in this chapter. Because simplification sometimes results in phones that do not clearly belong to one or another English phoneme, I'll be using the "[]" notation for pronunciations.

Preliminaries

Before we look at the simplifications that happen in English as speech becomes more casual, we need to look at some basic features of English phonology.

English is a Language in Which Syllable Stress Matters a Lot

First, in English, as in many languages in which stress plays a major role, there are significant differences between stressed and unstressed syllables. Stressed syllables permit all of the possible vowel phonemes, whereas unstressed syllables (in my accent) are mainly limited to /ə/, /ɪ/, and /i/, with /ə/ by far the most common. We can see these differences most clearly when we look at how the pronunciation of a syllable changes when it becomes stressed or unstressed. Consider the second syllables in the following related pairs of words.

1. *melody*, *melodic*
2. *repeat*, *repetition*

In *melodic* the second syllable is stressed, and the vowel is pronounced /ɑ/. In *melody* it is unstressed, and the vowel is pronounced /ə/. In the second pair, the second vowel is /i/ when it's stressed and /ə/ when it's unstressed.

Like all languages, English has phonotactic constraints on what sequences of vowels and consonants can occur. As in many languages, English phonotactics interacts with morphology, a topic that we will begin looking at in Chapter 7. For now, it is enough to know that English has words like *dog* and *believe* that consist of just one "morpheme". That is, you can't break them down into smaller meaningful units. And English has words like *dogs* and *believed* consisting of more than one morpheme: *dog* + -s and *believe* + -ed. The reason this matters for phonotactics is that the constraints are different for words consisting of one and words consisting of more than one morpheme. Words with just one morpheme normally do not contain sequences of more than one vowel. But words ending in the morpheme -ing may have such sequences.

3. *playing*, *trying*, *showing*

Similarly the limits on possible sequences of consonants within one-morpheme words are relaxed in words ending in morphemes such as -s and -ed

4. *bumped*, *asked*, *thanked*, *acts*, *fifths*, *sixths*

By Adding Suffixes to English Words, We Can Come Up with Longer Sequences of Consonants, as Many as Four in *Sixths*

These words end in sequences of three or four consonants that do not occur in words with one morpheme: /mpt/, /skt/, /ŋkt/, /kts/, /fθs/, /ksθs/.

General simplifications

As we have just seen, phonotactics is to some extent a matter of degree. There are constraints that apply to English words of one morpheme, but these are relaxed in words of more than one morpheme, which may result in sequences of vowels or consonants that would otherwise seem odd to English speakers. Consider first cases where the result is a sequence of more than one vowel, as in the words in 3 above. While these may be pronounced as such in careful speech, in casual speech, they may become a single vowel. This is especially likely when the *-ing* suffix would have the pronunciation [ən] as opposed to [ɪŋ]. So these three words could be pronounced [plen], [trayn], and [ʃon] in casual speech. In each case the /ə/ has been deleted to simplify the pronunciation.

In Casual Speech, English Speakers May Simplify Sequences of Vowels or Consonants by Dropping One of Them

The same thing can happen with sequences of consonants that are "odd" for English, as in 4 above. In some accents, these words are never pronounced with the sequence of three or four consonants that is implied by the spelling. In others, such as my accent, they are pronounced this way in careful speech but lose one of the consonants in casual speech, especially when the next word also begins with a consonant. So *acts* becomes [æks], *thanked* becomes [θæŋt], and *fifths* becomes [fɪθs].

Other possible simplifications may occur across the boundaries between words. Consider what happens when an alveolar consonant ends up before a /y/, as in two places in the following sentence.

5. *Write your name on this yellow sheet.*

Speaking carefully, most people would pronounce the two parts shown in bold as [ty] and [sy]. But when we speed up and allow ourselves to simplify, these may become [č] and [š]. This is an example of assimilation. The alveolar and palatal consonants combine to yield single consonants that are at the postalveolar place of articulation, which is in between the original places. Note, however, that this process must somehow be constrained. For example, in my accent the [ty] of *that yellow sheet* would never become [č], no matter how casually I'm speaking (instead the [t] would become a glottal stop, [ʔ]). The constraints appear to be quite complicated, and I won't have more to say about them, except that they are clearly related to the stress of the words on either side of the boundary.

A third tendency is for the appearance of the same consonant twice with an unstressed vowel in between to be simplified through the deletion of the vowel and the reduction to a single consonant. Here are some examples, with the careful pronunciation first, then the simplified pronunciation.

6. *probably*: ['prɒbəbli, 'prɒbli]

7. *necessary*: ['nɛsə, sɛri, 'nɛ, sɛri]

8. *terrorism*: ['tɛrə, rɪzəm, 'tɛ, rɪzəm]

Simplifications Specific to Particular Words

We have seen that there are simplifications that can happen in English speech to particular sequences of phonemes, independent of what word they occur in. Most of the familiar simplifications that characterize English casual speech, however, apply to particular words and not to others.

One possibility is a word that for whatever reason has a sequence of phonemes within it that is unusual for English. A good example is the word *sandwich*. In some accents, such as mine, this has the careful pronunciation ['sændwɪč]. But the sequence [ndw] in the middle of this word is very rare in English, and for many people the word has the alternate pronunciation ['sæmwɪč] in casual speech. (Of course for many other people this may be the only pronunciation of this word.) Note, however, that this tendency seems to be confined to this word. At least in my accent, it is not possible to simplify the sequence [ndw] similarly in a word such as *bandwidth*.

Why *Probably* Turns Into ['prɒli]

The frequency of a word has a clear effect on its tendency to be simplified. In 6 above we saw a simplified pronunciation for the word *probably*, but because this word is common, this may also represent the careful pronunciation for some speakers. In fact the word can undergo further simplifications: to ['prɒli] and even [pray].

Words that are not very informative also have a strong tendency to be simplified, probably because they are more predictable from the linguistic context than words that convey more information. Since many such words in English are also very common, they appear more often in their simplified form than in their canonical, careful pronunciation, which is normally only appropriate when they are stressed. Here is a partial list of these words, showing for each at least two pronunciations, ranging from the canonical, stressed, careful pronunciation to the most unstressed and casual.

- 9. *you*: [yu, yə]
- 10. *he*: [hi, i]
- 11. *him*: [hɪm, ɪm]
- 12. *them*: [ðɛm, əm]
- 13. *our*: [aʊr, ɔr]
- 14. *of*: [ʌv, əv, ə]
- 15. *to*: [tu, tə]
- 16. *for*: [fɔr, fər]
- 17. *out*: [awt, at]
- 18. *am*: [æm, əm, m]
- 19. *is*: [ɪz, əz, z, s]
- 20. *are*: [ɔr, ər, ə]
- 21. *have*: [hæv, əv, v]
- 22. *did*: [dɪd, dəd, d]
- 23. *will*: [wɪl, wəl, əl, l]
- 24. *can*: [kæn, kən]
- 25. *would*: [wʊd, əd, d]
- 26. *while*: [wayl, wal]
- 27. *because*: [bɪ'kɔz, bɪ'kʌz, kʌz, kəz]
- 28. *how*: [haw, ha]

Notice that many of these variants are predictable if we recall that unstressed vowels tend to be pronounced as [ə]. But not all vowels get reduced to [ə]. In some words the diphthongs [aw] and (less often) [ay] can lose their second element, becoming simply [a] (or [ɔ] before [r]). The pronunciation of *did* as [d] also agrees with the general tendency for a repeated consonant to merge into a single consonant. However, in other words, an initial or final consonant is dropped in a way that would not be predictable from the general tendencies discussed above. Thus it appears that some of these simplifications are conventional.

Finally, frequently occurring sequences of words are often simplified, especially when they consist of words from the group just illustrated. Some of these simplifications are so frequent that they represent the careful pronunciation, and they have even entered English orthography as contractions such as *it's*, *they're*, *doesn't*, and *I'd*. Most of these pronunciations are predictable from the simplified forms we have already encountered, but a few, such as *don't* (from [du nɔt]) and *won't* (from [wɪl nɔt]) are not. For some there are multiple pronunciations, varying in how casual and unstressed they sound. Here are a few of these.

- 29. *I'm*: [aym, am, əm]
- 30. *they're*: [ðɛr, ðər]
- 31. *we're*: [wɪr, wər]
- 32. *wouldn't*: ['wʊdn̩t, 'wʊdn̩]

Also very common are these combinations with *have* and *to* as the second word.

- 33. *would (should, could) have*: ['wʊrə]
- 34. *going to*: ['gʌnə]
- 35. *got to*: [gɔrə]

- 36. *want to*: [wʌnə]
- 37. *have to*: [hæftə]
- 38. *has to*: [hæstə]
- 39. *supposed to*: ['spostə]
- 40. *trying to*: ['traynə, 'tranə]

Notice how in 35 the two [t]s merge into one, pronounced as the tap [ɾ], and how they disappear altogether in 36.

Sequences of Words May Have Their Own Conventional Casual Pronunciation: *I'm Going to* [amənə]

Finally, let's see how some frequent sequences of three or more words are pronounced in very casual speech (in my accent).

- 41. *You wouldn't have thought so*. [yə,wʊdŋə 'θɔt ,so]
- 42. *What did you think?* [ˌwʌʃə 'θɪŋk]
- 43. *How has he been?* [ˌhæzi 'biːn]
- 44. *I am going to look*. [ˌamənə 'lʊk]
- 45. *I don't know*. [ˌaɪə 'no]

Let's summarize what we found for simplified speech in English. First, how likely a word or sequence of words is to be simplified depends on at least on these factors.

- How frequent the form is
- How little information the form carries
- How casual the situation is

Second, the simplifications that occur involve assimilation; the reduction of vowels, often to [ə]; the merging of sequences of the same consonant; the deletion of [ə] and some initial or final consonants. Many of these processes are general processes in the language. In some cases, however, the simplifications are conventions associated with particular words and must be learned separately.

This page titled [4.7: Phonology in the Wild](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.8: Problems

Phonetic Contexts and Assimilation

1. English vowels in most dialects have various possible degrees of length. In the following words, relatively long variants of the vowels are indicated with a following [:]. Based on these examples, say what the phonetic context is for the long allophone of English vowels.
 - a. *hat* [hæt]
 - b. *had* [hæ:d]
 - c. *gas* [gæs]
 - d. *jazz* [jæ:z]
 - e. *mate* [met]
 - f. *made* [me:d]
 - g. *roast* [rost]
 - h. *rose* [ro:z]
 - i. *hoop* [hup]
 - j. *tube* [tu:b]
 - k. *buck* [bʌk]
 - l. *bug* [bʌ:g]
2. The Japanese phoneme /s/ has two allophones: [s] and [ʃ]. Based on the following words, say what the phonetic contexts for the two allophones are.
 - a. [saya] 'pod'
 - b. [kasa] 'umbrella'
 - c. [senkyo] 'election'
 - d. [mise] 'store'
 - e. [sono] 'that'
 - f. [heso] 'navel'
 - g. [suʃi] 'sushi'
 - h. [hanasu] 'speak'
 - i. [ʃiku] 'spread'
 - j. [ʃima] 'island'
 - k. [ʃite] 'doing'
 - l. [kuʃi] 'skewer'
 - m. [saʃimi] 'sashimi'
 - n. [meʃi] 'rice'
 - o. [sasemaʃita] 'caused'
3. In Modern English, as you know, the fricatives [f, v, θ, ð, s, z] are all separate phonemes. But in Old English, although all of these phones occurred, they made up only three phonemes, each with a voiceless and a voiced allophone: [f, v], [s, z], [θ, ð]. The voiceless allophones are the more general (default) forms. Given the following words, (i) say what the phonetic context for the voiced allophones is, and (ii) say how the change from voiceless to voiced in this context is an example of assimilation. Hint: the context includes both what precedes and what follows the consonants. ([:] indicates vowel length, and phonetic details of vowels are not indicated because they are irrelevant.)
 - a. [fæst] 'firm'
 - b. [full] 'very'
 - c. [æfter] 'after'
 - d. [klif] 'cliff'
 - e. [heθvon] 'sky'
 - f. [seva] 'mind'
 - g. [hævde] 'had'

- h. [hweθrvan] 'return'
 - i. [æ:vre] 'always'
 - j. [sunu] 'son'
 - k. [la:st] 'track'
 - l. [hu:s] 'house'
 - m. [hors] 'horse'
 - n. [ræ:zan] 'to attack'
 - o. [i:zern] 'iron'
 - p. [ræ:zde] 'attacked'
 - q. [bizgu] 'occupation'
 - r. [θæ:əw] 'custom'
 - s. [wraθ] 'angry'
 - t. [so:θ] 'true'
 - u. [θiyeθ] 'receives'
 - v. [kweðan] 'to say'
 - w. [swi:ðre] 'right hand'
 - x. [wraðu] 'support'
 - y. [furðor] 'further'
4. The following words are from one dialect of Tzeltal. ['b] is a voiced glottalized bilabial stop, roughly a [b] accompanied by a glottal stop. Recall that [t'], [p'], and [k'] represent voiceless ejective (glottalized) stops.
- 1. [bi] 'what'
 - 2. [bu t'il] 'as'
 - 3. [hba] 'myself'
 - 4. [šbaht^h] 'goes'
 - 5. [sba] 'him/herself'
 - 6. [t'uhbil] 'beautiful'
 - 7. [ilbil] 'seen'
 - 8. [tahb] 'twenty'
 - 9. [ti'bal] 'meat'
 - 10. [ho'bel] 'San Cristóbal' (a city)
 - 11. [ma'ba] 'not'
 - 12. [ča'b] 'honey'
 - 13. [hayeb] 'how many'
 - 14. [tut^h] 'small'
 - 15. [tulel] 'to harvest'
 - 16. [htat^h] 'my father'
 - 17. [čitam] 'pig'
 - 18. [pat^h] 'back'
 - 19. [naht^h] 'tall'
 - 20. [sit^h] 'fruit'
 - 21. [t'ut'] 'greedy'
 - 22. [t'anal] 'heaped'
 - 23. [t'ulel] 'to pour'
 - 24. [yut'il] 'inside'
 - 25. [naht'] 'long'
 - 26. [pat^h] 'back'
 - 27. [pohp^h] 'mat'
 - 28. [spuy] 'his snail'
 - 29. [hpik^h] '8000'

30. [k'opoh] 'spoke'
31. [sp'uy] 'squashed'
32. [p'ihp'inel] 'to spread'
33. [snop'] 'seized'
34. [lap'ap'] 'sticky'
35. [hp'itp'on] 'throbbing'
 - a. The Tzeltal phoneme /b/ has two allophones, [b] and [b̥]. Using the words above, which are representative of the contexts in which the allophones occur, say what the complementary distribution of the allophones is.
 - b. The Tzeltal phoneme /t/ has two allophones, [t] and [t̥]. Using the words above, which are representative of the contexts in which the allophones occur, say what the complementary distribution of the allophones is.
5. Recall that the Spanish phonemes /b/ and /g/ each have two allophones, stops ([b] and [g]) and approximants ([β] and [ɣ]) and that the stops are used when the consonant begins a word after a pause. But they are also used when these consonants follow a nasal consonant. In fact if we look at a lot of Spanish words in context, we see that the only nasal consonant that occurs before /b/ is [m], and the only nasal consonant that occurs before /g/ is [ŋ]. In other words these sequences are possible: [mb], [ŋg], and these are not: [mβ], [nb], [ng], [nɣ], [ŋɣ]. Some words ending in nasals even change their pronunciation to maintain these constraints. So consider the words *un* and *con*, which normally end in a dental nasal [ɲ]. When they are followed by /b/ or /g/, however, they take the form of [m] or [ŋ], for example, *un vaso* [um'baso] 'a glass', *con gusto* [kon'gusto] 'with pleasure'. Explain what is going on, that is, why the nasal consonant changes and why /b/ and /g/ are not realized as approximants after nasals. Your explanation should be in terms of assimilation.

Distribution of Phones

1. In addition to the voiced stops /b/, /d/, and /g/, the voiced fricative /z/, and the nasals /m/, /n/, and /ŋ/, Lingala has a set of **prenasalized** voiced stops and fricatives, that is, stops that begin with nasalization. I'll write them with a superscript nasal consonant symbol preceding the stop symbol, for example, /^mb/ for the voiced bilabial stop beginning with nasalization and bilabial closure ([m]). The prenasalized stops that occur in all dialects of Lingala are /^mb/, /ⁿd/, /ⁿz/, and /^ŋg/.
 - a. Say why you think these are the only prenasalized stops and fricatives that occur and not, for example, /^mz/ or /ⁿd/ or /ⁿb/.
 - b. Given the following words, say how you know that (i) [b] and [^mb] belong to separate phonemes, (ii) [m] and [^mb] belong to separate phonemes, (iii) [n] and [ⁿd] belong to separate phonemes, (iv) [g] and [^ŋg] belong to separate phonemes. Recall that the best way to establish that two phones belong to separate phonemes is to find a minimal pair for them. ([ˈ] over a vowel marks a high tone; low tone is unmarked.)
 - I. [kozi^mba] 'to trick'
 - II. [^mbá^ŋgu] 'fast'
 - III. [bé^ŋga] 'call!'
 - IV. [gú^mbá] 'fold!'
 - V. [kozima] 'to be extinguished'
 - VI. [núka] 'gather'
 - VII. [^mbé^ŋga] 'vocation'
 - VIII. [ⁿdúka] 'dam'
 - IX. [ⁿdáko] 'house'
 - X. [ko^ŋgala] 'to be wild'
 - XI. [^ŋgú^mbá] 'spleen'
2. Old English had both short and long low front vowels, [æ] and [æ:]. From the following examples, say how you can know that these two phones belong to different phonemes.
 - a. [sæ:] 'sea'
 - b. [θæ:r] 'there'
 - c. [æ:t] 'food'
 - d. [mæ:st] 'most'
 - e. [græ:y] 'gray'
 - f. [dræ:van] 'to drive'

- g. [klæ:ne] 'clean'
- h. [bæk] 'back'
- i. [æt] 'at'
- j. [wæter] 'water'
- k. [fæstan] 'fasten'
- l. [mæst] 'mast'
- m. [næyl] 'nail'
- n. [hwæt] 'what'

3. In some accents of southeastern England, the tense (long) vowels include /ʊu/, /ɜ:/ (a long rounded, high, central vowel), and /ʌu/. From the following examples, say how you know these are separate phonemes in this accent. Note: they do *not* correspond directly to phonemes in other accents such as General American.

- a. *moan* [mʊun]
- b. *tow* [tʊu]
- c. *nose* [nʊuz]
- d. *sole* [sʊul]
- e. *mown* [mʌun]
- f. *soul* [sʌul]
- g. *knows* [nʌuz]
- h. *toe* [tʌu]
- i. *news* [nɜ:z]
- j. *moon* [mɜ:n]
- k. *two* [tɜ:]
- l. *soon* [sɜ:n]

4. In sign languages the main dimensions along which syllables differ are handshape, location (the place on the body or in space where the sign is made), movement (the motion of the articulators in space), and orientation (the direction that the palm "points"). Given the following ASL signs, say how you know that location and movement are **contrastive** dimensions in ASL.

'airplane'
'dry'
'fly (airplane)'
'hamburger'
'massage'
'piano'
'summer'
'worthless'

5. Below are some words from Argentine Spanish containing the phones [s] and [h]. (Each has a different meaning, but the meanings are left off because they might make it easy for students who know Spanish.) Syllable boundaries are marked with a space, and [x] represents a voiceless velar fricative. Based on these words only, are these two sounds in complementary or overlapping distribution in the language? If the distributions are complementary, say what the contexts for each phone are. If the distributions are overlapping, say how you know. (Hint: [h] is the more restricted phone.)

English examples:

- [p^h] (aspirated, voiceless bilabial stop), [p] — complementary, [p^h]: at the beginning of a stressed syllable, [p] elsewhere
- [s], [ʃ] — overlapping, *sip/ship* is a minimal pair
 - a. [sa 'lar]
 - b. ['ka sah]

- c. [soy]
- d. ['se so]
- e. ['syen]
- f. ['swer te]
- g. ['swa reh]
- h. [gon 'sa leh]
- i. [dye si 'sye te]
- j. [rre 'swel to]
- k. ['xwi syo]
- l. [xweh]
- m. ['ka xah]
- n. ['fyeh tah]
- o. ['xeh to]
- p. [pah]
- q. [suh]

6. In Amharic, consonants can be simple, for example, [t] and [m], or long, for example, [tt] and [mm]. Given the following words, say whether consonant length is a contrastive dimension in Amharic, and explain how you know.

- a. [yɪmɛtal] 'he hits'
- b. [mɛtta] 'he hit'
- c. [mɛla] 'scheme'
- d. [tɛdɛrrɛgɛ] 'it was done'
- e. [bɛrr] 'door'
- f. [nɛgɛ] 'tomorrow'
- g. [mɛlla] 'it got full'
- h. [nɛgga] 'it dawned'
- i. [yɪmmɛttal] 'he got hit'
- j. [zɛr] 'seed'
- k. [k'ɛllɛlɛ] 'it got easy'

7. Amharic /b/ has two allophones, [b] and [β] (a voiced bilabial fricative). Given the following words, which are representative of words containing this phoneme, say what the complementary distribution of these allophones is. (Remember from the last problem that consonant length is contrastive in Amharic.)

- a. [laβ] 'sweat'
- b. [rɛhaβ] 'hunger'
- c. [nɛβɪr] 'leopard'
- d. [nɛbbɛrɛ] 'was'
- e. [bɪrd] 'cold'
- f. [bal] 'husband'
- g. [bɛlla] 'he (it) ate'
- h. [tɛβɛlla] 'he (it) was eaten'
- i. [kɪβɪr] 'honor'
- j. [tɛβabbɛru] 'they were united'
- k. [aβro] 'together'
- l. [gɛβs] 'barley'
- m. [wɪβɛt] 'beauty'
- n. [tɛsɛbbɛrɛbbɛt] 'it was broken by it'
- o. [kɛβt] 'livestock'
- p. [t'ɪβɛβ] 'wisdom'
- q. [t'ɛbbaβ] 'narrow'
- r. [ambɛssa] 'lion'

s. [arba] 'forty'

t. [albɛlam] 'I don't eat'

8. In Tzeltal /t/ (with the context-sensitive allophones [t] and [t^h]) and /t'/ ([t']) are separate phonemes. Say how you can know this from the words in problem 4 in 4.8.1 above.
9. In Tzeltal /p/ (with the context-sensitive allophones [p] and [p^h]) and /p'/ ([p']) are separate phonemes. Say how you can know this from the words in problem 4 in 4.8.1 above.

This page titled [4.8: Problems](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

5: Composition - Combining Words

[5.0: Introduction](#)

[5.1: Attributes and Attribution](#)

[5.2: Modification](#)

[5.3: Compositionality and Idiomaticity](#)

[5.4: Problems](#)

This page titled [5: Composition - Combining Words](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.0: Introduction

As we've seen, words allow people to refer to a potentially very large number of things in the world, either directly with proper nouns, or indirectly through semantic categories of utterance roles (such as hearer) with common nouns and personal pronouns. But even the flexibility that comes with common nouns is limited. No matter how many categories a speech community labels with nouns, there will always be multiple members of particular categories that need to be distinguished from each other, for example, one apple from another. Coming up with new categories and new labels for them is a slow process, certainly not fast enough to cope with the minute-to-minute demands of communication, where the difference between one apple and another may matter a lot. This chapter is about the way people cope with this need by making use of one of the most fundamental properties of human language, compositionality, the power to combine words into phrases whose meanings are combinations of the meanings of the words. The focus is on noun modification, the use of words in combination with nouns to restrict the meanings of the nouns. In later chapters, we will see how compositionality allows us to refer to events and states as well as things.

This page titled [5.0: Introduction](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.1: Attributes and Attribution

Attributes of Things

Exercise 5.1.1

Apples come in various colors, sizes, and flavors. If we want to distinguish two apples from one another, we need to be able to talk about these dimensions. One possibility is more nouns, one for each subcategory of apple. What is the advantage of doing this instead of using the noun we already have, *apple*, along with words like *red*, *big*, and *sour*?

Let's review what our Lexies have so far. They have a set of nouns that allow them to refer to important individuals and to members of categories of things in their world. And each of these nouns is represented in their mental lexicons as an association between a form and a meaning. Specifically, within the class of common nouns, the most important type of word we've seen so far, each word form is associated with, or **designates**, a category of things. Note that we will need to distinguish the notion of designation from the notion of reference. A Speaker produces an expression like *the apple* or *that apple* to refer to a particular apple in the world or in the imagination of the Speaker or Hearer. But within this expression the noun *apple* itself cannot be said to refer to a particular apple; rather it designates a whole category, the category apple, which includes many possible individual apples. It is only in combination with a word like *the* or *that* that it is used for reference.

Let's place some new communicative demands on the Lexies and see what they come up with. The common nouns that they have are very powerful because they allow them to refer to a potentially infinite set of things in the world, that is, all of the members of the categories that are the meanings of the nouns. But what if they need to distinguish one member of a category from another, one apple from another, one tree from another, one tiger from another? This kind of need comes up in two kinds of situations. If you want to announce to the rest of your tribe that you've found a tree, it may be important for them to know that it's tall (perhaps because tall trees are a useful source of wood). That is, in one kind of referring situation, the Speaker *introduces* a new referent to the Hearer. But say the others already know about a group of trees and you want them to focus on the one that stands out as tallest. That is, in another kind of referring situation, the Speaker calls the Hearer's attention to a referent that they are already aware of.

Nouns by Themselves are Limited in What Things They Allow Us to Distinguish

The problem is that the nouns you have don't allow you to distinguish the members within each of the categories. One way to deal with this problem would be more semantic categories and more nouns. So the Lexies could have different nouns for apples of different colors or flavors, for trees of different heights or widths, for tigers of different sizes or ferocity. Let's look again at the categories that nouns designate and see if this makes sense.

Recall that each of the things in the world can be thought of as having a value on each of a set of dimensions such as size, color, shape, taste, and consistency. Each of the categories that we have common nouns for represents a strong tendency for particular sets of values on different dimensions to co-occur. An apple isn't just an object of a particular shape; it has a characteristic range of sizes, tastes, consistencies, and locations. In other words, the category apple is a whole cluster of co-occurring features, and it useful to know that something belongs to the category because you can predict a lot about it. Even if you don't know what it tastes like or what it looks like inside the peel, you can predict these things.

But even though apples tend to have particular sizes, they do vary in size. What can you say given that something is a relatively small apple? If being small doesn't imply having a particular taste or consistency or size, then the category *small_apple* wouldn't be useful in the same way the category apple is. If being unusually ferocious for a tiger doesn't tell us much of anything about its size, color, or location that we don't already know, then having a category *ferocious_tiger* won't be worth much.

It's not that the size of an apple and the ferocity of a tiger don't matter for the Lexies; they very likely do. A big apple means more to eat. A ferocious tiger means something to stay away from. It's that size for apples may work roughly the same way for raspberries and rocks and that ferocity for tigers may work roughly the same way for bears and crocodiles.

The point is that people are not as likely to form coherent categories for things like small apples, ferocious tigers, and tall trees as they are for things like apples, tigers, and trees. And because they don't have coherent categories like *small_apple*, they are not likely to have nouns for such things. In any case, having a different noun for all the distinguishable combinations of size, color, sourness, and texture in apples would greatly increase the number of nouns that the Lexies would have to learn and remember.

A Major Advance: Referring with More than One Word

You probably can tell where we're headed. There is an alternative to having separate nouns for large and small apples, for large and small raspberries, and for large and small rocks. We can divide the work of referring to something like a small apple into two parts, one part dedicated to the category, apple, the other part dedicated to the dimension along which members of the category vary, size in this case. That is, we can refer to an apple with two words: *small apple*.

So now we are seeing each of the members of a category as having particular properties that I'll call **attributes**. Each attribute is really a value on some dimension such as size, color, or consistency. But seeing things this way is not just a convenient way of talking about language. The fact that people can refer to things with expressions like *small, sour apple* means that people have the capacity, when they see or think about an apple, to separate out its attributes from its category. A small, sour apple isn't just an undifferentiated object for us. It's apple, it's small, and it's sour. Being able to keep these aspects of an separate, at least for the purposes of talking about the object, is an impressive ability, one that is so intimately tied up with the language that it's hard to say which came (or comes in each developing child) first.

Attribution

Exercise 5.1.2

If you hear somebody say, *please hand me the bleg apple*, what can you infer about the function of the unfamiliar word *bleg* in the sentence? What is it that tells you this?

An expression like *red apple* consisting of one or more words that make up a unit in the sense that they designate a unified concept is called a **phrase**. I'll have a lot more to say about phrases and how you can know whether a group of words is a phrase in the next chapter. In this chapter we'll only be concerned with phrases consisting of two words, one of them a common noun such as *apple* that designates a category of things. But what about the other word in a phrase such as *red apple*?

Form and Meaning (Again): Syntax and Semantics

With such a phrase we can speak of two relations, the relation between the words in the phrase and the relation between the concepts that the words designate. The first kind of relation is **syntactic**; it is a relation between linguistic forms. The second kind of relation is **semantic**; it is a relation between concepts designated by linguistic forms. Let's start with the semantic relation. One of these is a category of objects, apple. The other is an attribute that characterizes some members of this category, red, one possible value (more precisely, range of values) on a conceptual dimension, color. The syntactic relation between the words involves at least two components. First, it has an order: the word designating the attribute comes before the word designating the conceptual category. Second, the words belong to their own categories. The word designating the conceptual category, *apple* in the example, is one of the set of common nouns. The word designating the attribute, *red* in the example, is one of a set of words called **adjectives**.

I will refer to both the semantic and syntactic relation in this case as **attribution**, a subtype of the more general relation called **modification**. In other words, we can say that the adjective has an attributive function in the phrase and that the adjective modifies the noun. And we can say that the phrase attributes redness to the apple that is being referred to. Recall again that the main theme of this book is the intimate association between form on the one hand and meaning and function on the other. In fact, most of the rest of the book will be about the association between syntactic relations or patterns and semantic relations or patterns. Attribution is our first example. The close association between the syntactic attribution pattern and the semantic attribution pattern allows a Hearer to infer something about the meaning of a novel word, for example, *bleg*, in the *bleg apple*. *Bleg* appears to be an adjective, designating some sort of apple attribute.

Back to our Lexies. Because they realize that coming up with common nouns for all of the things that need to be distinguished is unreasonable, they decide instead to come up with a set of adjectives to designate attributes of things. These include words similar to the English words *red, small, shallow, wet, dark, smooth, sour, young, pregnant, angry, kind, and dead*. Notice how greatly this extends the capacity of their language. With 50 nouns and 10 adjectives, they can produce 500 different combinations of adjective and noun. Of course some of these are not likely to be produced or to make much sense if they are produced: *shallow apple, sour cat, pregnant river*. But many of the adjectives will apply to large subsets of nouns: *red apple, red raspberry, red stone, red skin; wet apple, wet stone, wet tree, wet skin*.

But the increased power goes beyond just the number of combinations that are possible. Speakers and Hearers who know the meanings of the adjectives and the nouns can produce and understand combinations of the words that they've never heard before, in fact, that have never been uttered by anyone before. This is one of the senses in which language is **productive**; that is, it permits new expressions to be produced and understood. The possibility of combining words into phrases was a revolutionary leap for our ancestors, probably the single most important feature of human language. I will have more to say about productivity later in this chapter and later in the book, especially in Chapter 8.

Before we go on, we should keep in mind the difference between describing of a property of language and explaining how people actually make use of this property. While everyone agrees that people can produce and understand (in some sense of "understand") combinations of words they haven't heard before, there is a lot of disagreement about how this works in the mind (and the brain). In particular, there is the question of how explicit the knowledge is that permits this ability. Do English speakers have an explicit rule that tells them how combinations of adjectives and nouns are to be interpreted? Or do they remember many examples of these combinations, along with their contexts, and then somehow combine all of the relevant examples when they're faced with producing or understanding a new combination? For the moment we'll be focusing mainly on the perspective of linguistics, that is, describing what's going on from the perspective of the language rather than the language user. But we'll need to return to Speakers, Hearers, and Learners later on.

Scalar Adjectives

Exercise 5.1.3

Compare the meaning of *big* in *big elephant* and *big mosquito*. What do these examples tell you about the meaning of adjectives such as *big*, *soft*, and *sour*?

Let's consider in more detail how adjectives designate attributes. The simpler cases are the ones where the interpretation of the adjective seems to be almost completely independent of the noun that it modifies. *Pregnant* is pregnant, whether we are talking about a pregnant woman or a pregnant tiger. But there are not many such adjectives, either in English or in other languages. More often the interpretation of the adjective depends in one way or another on the context: on the modified noun, on other words that occur before or after the phrase, or on the situation in which the phrase is uttered.

Most Adjectives Have Relative, Rather than Absolute, Meanings

Consider a continuous dimension, such as size, darkness, age, or crispness. There are many, potentially an infinite number of, possible values on such dimensions. To designate an attribute for a thing on dimensions like these, we could have words for particular ranges of values, for example, 10-20 cubic meters on the size dimension and 2-3 years on the age dimension. That is, attribute words could have an **absolute** interpretation. But this isn't the way adjectives for these dimensions usually work. In English, the crispness dimension has adjectives for the two poles of the dimension, *mushy* (or *soft*) and *crisp*. But neither of these adjectives has an absolute meaning; the precise meanings of *mushy* and *crisp* depend on how they are used.

Compare the meanings of *a mushy apple* and *a mushy orange*. In both cases, *mushy* means something like 'closer to the mushy end of the crispness dimension than some **standard** of comparison'. But the standard differs. In the first case, the standard is either the crispness of a prototypical apple or some set of apples familiar to the Hearer or present in the utterance context. In the second case, the standard is the crispness of a prototypical orange or some familiar set of oranges. *Mushy* in the second example probably describes a lesser degree of crispness than in the first case.

Adjectives such as *big*, *little*, *crisp*, *mushy*, *dark*, and *light* that designate values on continuous dimensions are called **scalar adjectives**. Scalar adjectives do not normally designate absolute values or ranges of values. Rather their meanings are **relative** to a standard provided by the context. The Hearer is faced with the problem of determining what this standard is. The relative nature of scalar adjectives allows us to use the same adjectives for things with all sorts of values on the relevant dimensions. More importantly, scalar adjectives are probably relative because it is the relative value of things that matters. Everyone who knows elephants knows that they are large creatures. To distinguish one elephant from another, however, we need a word that designates not absolute size but elephant size. *Big* and *little* do this for us.

Consider how adjectives are like and unlike nouns. Just as a noun like *raspberry* designates a category of things, an adjective like *pregnant* designates a category of things (all of the animals that are currently pregnant). But most adjectives designate strange sorts

of categories. While knowing that something is a raspberry tells us many things about it, knowing that something is red or crisp tells us very little; we know about only one of its properties. At first glance, it may seem that adjectives are simpler than nouns because we have only one feature to worry about. But the fact that each adjective lumps together so many things that are so different from one another makes them more complex than nouns in another way. It may be for this reason that adjectives seem generally harder for children to learn than nouns (this is true in English in any case; in many languages it's hard to compare nouns and adjectives because there are so few adjectives). It may just be hard for children to focus in on the single dimension that matters for a word like *thin*, *red*, or *smooth*. For more on this possibility, see this paper that I wrote with Linda B. Smith.

This page titled [5.1: Attributes and Attribution](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.2: Modification

Noun + Noun Phrases

Exercise 5.2.1

Consider the following phrases, each consisting of two nouns. How does the meaning of the phrase depend on the relation between the semantic categories represented by the nouns, and how does this relation vary for the different phrases?


- *apple tree*
- *apple seed*
- *apple pie*
- *apple bowl*
- *apple sweater*
- *apple company*
- *apple festival*

In this section, we'll look at another kind of modification in English, noun + noun phrases. This should give you a better idea of how modification works. It also provides a good example of how Hearers often have to rely on the context of a phrase to interpret it.

Some languages, including English and Japanese from our list of languages, allow phrases consisting of two nouns together, phrases such as *apple pie*, *desk chair*, and *axe handle*. Like simple nouns, these phrases designate categories of things. In each case, the category is a subcategory of the category that is designated by the second noun. That is, an apple pie is a kind of pie, a desk chair is a kind of chair, and an axe handle is a kind of handle. But how is an apple pie different from a pie and a desk chair different from a chair?

Let's start our consideration of noun + noun phrases with a phrase which may be unfamiliar to you but which you can probably interpret anyway, *pie apple*. Most likely you read the phrase to mean a kind of apple that is used in making pies. How does this meaning derive from the two words? Since there are two nouns in the expression, there are two categories involved in the meaning of the phrase, the categories of things designated by the two nouns. The category designated by the second noun, apple, is restricted by the category of the first noun, pie. Specifically, the phrase designates a subcategory of apple whose members are used in making some members of the category pie.

We can diagram the meaning of the phrase as in the figure below. If we think of a category as the set of all members of the category, then we can represent that set by a circle, with each point in the circle corresponding to a possible member of the category. (People probably don't represent categories as sets in their minds, but this does provide a convenient way of visualizing categories and the relations between them.) Because there are two nouns in *pie apple*, there are two large circles in the figure. The smaller circle within the pie circle represents the category of pies containing apples, that is, the category designated by the phrase *apple pie*. The smaller circle within the apple circle represents the category that the phrase *pie apple* designates, the category of apples that are used in pies. The arrow represents the conceptual relation between the apples in this subcategory and the pies that they are used in; the label on this arrow, use, is meant to show what kind of conceptual relation this is.




This sort of characterization also works for phrases such as *pancake flour*, *watch screw*, and *tire rubber*. In each case, the phrase designates a subcategory of the category designated by the second noun, and that subcategory consists of all of those members of the larger category that are used for at least some members of the category designated by the first noun. For example, *tire rubber* designates a subcategory of the category rubber whose members are used in making tires (members of the category tire).

The Meaning of a Noun + Noun Phrase Involves Two Different Categories of Things and a Conceptual Relation

But what about a phrase such as *mountain goat*? We can't think of goats as being used for mountains. Rather the relevant conceptual relation seems to be one of location. This relation also applies for phrases such as *corner table*, *wall clock*, and *river stone*. In each of these cases, the phrase designates a subcategory of the category designated by the second noun, and that subcategory consists of all members of the larger category that are located at or near some member of the category designated by the first noun.

But there are still other possible conceptual relations for noun + noun phrases. In *crow feather*, the members of the designated subcategory are *part* of members of the first noun category. In *tire company*, the members of the subcategory *deal* in members of the first noun category. In *apple pie*, the relevant relation is one of *content*. It may be helpful to compare the meaning of *apple pie* with the meaning of *pie apple*, diagrammed in the figure above. For *apple pie*, it is a subcategory of pie that is designated, and that subcategory consists of all members of the larger category that contain members of the category designated by the first noun, apple. This is illustrated in the figure below.

The Same Two Nouns in Different Orders Can Mean Very Different Things



In fact it appears that there are few limits on what sort of conceptual relation can be behind the meaning of a noun + noun phrase in English, as long as the relation is seen by the Speaker as somehow accessible to a Hearer. If this is so, then how do we describe, in as general terms as possible, the meaning of noun + noun phrases, the kind of knowledge that would be needed by English Speakers and Hearers to produce and understand these phrases? Here is an attempt.

An English phrase consisting of a noun A designating a thing category C_A followed by a noun B designating a thing category C_B designates a subcategory of C_B whose members are related in a particular way to at least some members of C_A .

Attribution and Modification

How can we characterize attribution (as in *mushy apple*) in similar terms? Here is one way to state it.

An English attributive phrase consisting of an adjective Adj designating an attribute Att followed by a noun N designating a thing category C designates the subcategory of C whose members have attribute Att.

If both attribution and the noun + noun pattern are examples of noun modification, we can now state generally what modification is. But first let's make sure we know what *kind* of thing modification is. It has both a syntactic component — the relationship between the words — and a semantic component — the relationship between the concepts and how they combine to give the meaning of the whole phrase. I will refer to such a dual relation, including both a form component (the syntactic part) and a meaning component (the semantic part), as a **grammatical relation**. Most of the rest of this book is about one grammatical relation or another.

Back to our definition of modification. Syntactically, noun modification is a relation between a modifier and a **head** noun. In English the modifier precedes the head if it is an adjective or another noun. Semantically, noun modification involves the category of things designated by the head noun and whatever concept is designated by the modifier. The two modifying concepts we have seen are attributes (designated by adjectives) and further categories of things (designated by nouns). Specifically the modifying concept functions to narrow down the category designated by the head noun. The whole phrase designates this narrowed category.

Speakers Can Apparently Create New Categories as They are Speaking, and Hearers Apparently Know How to Apply the Categories

What sort of ability is implied by the fact that Speakers and Hearers can produce and understand expressions in which a noun is modified? Each such expression represents a category of things that is narrower than the category designated by the head noun. When the expression is a familiar one, such as *green apple*, *happy face*, *apple pie*, or *table lamp*, this narrower category will also be a familiar one. But when the expression is new, when the Speaker is producing it for the first time or the Hearer is hearing it for the

first time, as might be the case for *fuzzy apple* or *artichoke pie*, a new category is being created on the fly. That is, people apparently have the ability to readily combine concepts to form new categories and to understand how a new category can apply to particular individual things in the world. For example, given an instance of a pie made out of artichokes, a Speaker could create a new category for such things and refer to it with the phrase *artichoke pie*, and a Hearer, when presented with the phrase, could create the category in their memory, even without any previous experience with such pies.

Interpretation and Ambiguity

Exercise 5.2.2

What different interpretations might the phrase *apple girl* have? What might allow a hearer to choose from among the different interpretations?

But how do Hearers interpret the kinds of noun + noun phrases we've been looking at? If many different conceptual relations can associate the two semantic categories with one another, then how do Hearers know which one is intended? For familiar combinations, such as *apple pie*, *dirt bike*, and *flower pot*, this is not a problem. In these cases, the Hearer has probably learned through experience with the phrases which interpretation is most likely.

Let's take a potentially confusing example, *llama blanket*. This could be a blanket that you put on the back of a llama, a blanket that your llama likes to sleep on or with, a blanket made from the wool of llamas, a blanket decorated with images of llamas, even a blanket in the shape of the outline of a llama. Such a phrase is **ambiguous**; it has more than one possible interpretation. Ambiguity is rampant in human language. In fact it is probably best to see the words and sentences that people produce as representing only fragments of what Speakers are thinking and what they would like Hearers to be thinking. In a sense every utterance is ambiguous. A phrase like *llama blanket* is just more ambiguous than most.

Hearers Have to Rely on More than Just their Knowledge of a Language to Understand Sentences in the Language

So generally the problem that the Hearer faces is coming up with a single interpretation out of the various interpretations that are consistent with what is said. Clearly the Hearer has to rely on information outside of what the words and the grammar of a sentence alone convey. There are several sources for this information. One is the meanings of other phrases and sentences in the context of the ambiguous one. For example, say *llama blanket* occurred in this sentence: *No, no, I meant the llama blanket; the one you gave me has pictures of geese on it*. In this case the Hearer would have reason to believe that the phrase referred to a blanket with pictures of llamas on it.

Another source of information is the non-linguistic context of the words, what is present in the environment, who the Speaker is, and what the Speaker might want the Hearer to know. For example, if there are several blankets in the Hearer's view, each in the shape of a different animal, the Hearer might expect *llama blanket* to mean 'blanket in the shape of a llama'.

A final source of information is the Hearer's knowledge of the world, in this case, what the Hearer knows about llamas and about blankets. The fact that wool can be made from the hair of llamas means that the interpretation of the phrase as 'blanket made from llama wool' is possible, though of course only if the Hearer knows this fact about llamas.

In sum, then, Speakers use nouns like *apple* and *blanket* when they want to refer to members of categories like apple and blanket. But sometimes these categories are too broad for the distinctions that they want to make. In this case they narrow down the categories, in some cases to categories they previously had not considered, using modifiers such as adjectives or other nouns. An adjective narrows down the head noun category by restricting it to members with a particular attribute (green, rough, etc.). A noun modifier narrows down the head noun category by restricting it to members with some relation to the members of another category, the one that is designated by the noun modifier. A variety of relations is possible, however, and the Hearer attempts to figure out which one is intended using possible experience with the particular noun + noun combination, the linguistic context, the non-linguistic context, and knowledge of the world.

Modification, however, is just one of the ways in which meaningful elements are combined to form larger meaningful expressions. In the next section, we'll extend the discussion to this more general property of human language.

This page titled [5.2: Modification](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.3: Compositionality and Idiomaticity

Compositionality

Exercise 5.3.1

What kinds of knowledge would be required for the interpretation of a phrase the hearer has never heard before such as *purple apple* or *apple soup*?

In this section we'll look more generally at what is involved in producing and understanding meaningful combinations of words. People's ability to do this is based on what is probably the most important property of human language, **compositionality**, the property that the meaning of a phrase is derived from the meanings of the words in the phrase and the grammatical relation that joins them. The details of this property and how it works in the production and understanding of language are somewhat complicated though; in fact, there's quite a lot of controversy about these details. Most of the complexity revolves around what conventions Speakers and Hearers need to store in long-term memory to make things work and the degree to which language actually is compositional.

Let's begin by again considering the interpretation of an unfamiliar phrase. Say a Hearer hears or reads the phrase *bitter grape* for the first time, and say there is no help from the context about what this means. The Hearer is familiar, however, with the words *bitter* and *grape*. To come up with an interpretation for the new phrase, the Hearer would first need to retrieve the meanings for *bitter* and for *grape* from long-term memory, specifically from the Hearer's lexicon. Recall that the lexicon is a sort of list of conventions for words, how they are pronounced (and written if the language has a writing system and the person is literate) and what they mean. The meaning of *grape* would be some representation for the category of things grape, and the meaning of *bitter* would be some representation for the attribute bitter that things may have when they are tasted.

The Interpretation of a Phrase Involves Both Lexical and Grammatical Conventions

But how are the meanings of the two words to be combined? The knowledge of how to do this comes from another convention, the grammatical convention that spells out how attribution works in English. Specifically, as we discovered in the last section, this convention specifies that the meaning of a phrase consisting of an adjective modifying a noun is a subcategory of the noun's meaning whose members all have the attribute that is the adjective's meaning. We can think of this as a rule to be followed in producing or understanding a phrase or as a pattern that describes the meaning of the phrase in terms of its parts. As with nearly every other aspect of language, there is a lot of disagreement about whether Speakers and Hearers actually have such an explicit rule in their minds (or brains). For now, let's just assume that we're trying to come up with a way of describing their behavior.

So the Hearer puts all of this together by inserting the meanings of the individual words into the grammatical pattern: the meaning of the phrase *bitter grape* is the subcategory of grape whose members all have the attribute bitter. One use of the phrase *bitter grape*, the only one we've considered so far, would be to refer to a particular grape, as in the slightly longer phrase *the bitter grape*. The Hearer would understand this expression to refer to some member of the category designated by *bitter grape*.

To summarize what we've seen so far, the meaning of a phrase is the **composition** of the meaning of its parts, that is, the meanings of the parts substituted into a general rule for how the meanings of the parts of a phrase of that type are combined. This involves two kinds of conventions, lexical conventions, that is, knowledge about the meanings of words (*grape*, *bitter*), and grammatical conventions, that is, knowledge about how meanings are combined for particular grammatical relations (such as *attribution*).

But the example we considered was relatively simple because each of the words has a meaning that is more or less independent of its context. How we interpret the noun *grape* and the adjective *bitter* does not seem to depend much on what other words are present in the utterance. Such meanings are said to be **context free**; they don't rely on their linguistic context, that is, the other words in their environment, for their interpretation.

Many cases are not this simple. We have already seen one type of example, involving adjectives whose meanings seem to change with the noun they modify. *Bitter* is a scalar adjective, but one whose meaning seems to vary little with the context. A scalar adjective such as *mushy*, on the other hand, apparently requires a standard for its interpretation. As we learned in the section on attribution and adjectives, the standard could be the default value on that dimension for the head noun category, or it could come from the context of the utterance containing the adjective. So for a phrase like *mushy grape*, the hearer would have to combine the

category grape with an adjective meaning like *mushy relative to some standard* and would have to figure out what the relevant standard is.

But even in this last case, the basic information required to figure out the meaning is available in the lexicon (the meanings of *mushy* and *grape*) and the grammar (the attributive rule for combining the meaning of adjective and head noun). This knowledge would only need to be supplemented with the additional ability to find the appropriate standard (either the default for the head noun category or the average over some set of members of that category in the context). So we are still talking about pure compositionality, the possibility of coming up with the meaning of the phrase on the basis of the lexical knowledge of the words in the phrase and general grammatical knowledge.

But there are further cases where it seems that additional linguistic knowledge is behind the interpretation of a phrase. Consider how a hearer would handle the phrase *grape pie*. Remember that in such noun + noun phrases, there are a number of possible relations between the semantic categories that the two nouns designate, including location, use, part-whole, and contents. Most people would probably interpret *grape pie* to mean a pie containing (or made from) grapes. The question is how they arrive at this interpretation. One possibility would be to figure out the meaning on the basis of what seems most likely. It is very unlikely, for example, that a pie is meant to be seen as a *part* of a grape (as in *rabbit foot*), that a pie is being *used* for a grape (as in *bread flour*), or that a pie is located *at or near* a grape (as in *lake fish*).

As discussed in the last chapter in the context of ambiguity, Hearer almost certainly use knowledge of this sort about the world in interpreting noun + noun phrases. But there is a further kind of knowledge that they could use. Consider other *familiar* noun + noun phrases that overlap with *grape pie*. There are some familiar noun + noun phrases with *grape* as the first word and a word designating a food or drink as the second: *grape soda*, *grape jello*, *grape jelly*. There are also familiar noun + noun phrases with *pie* as the second word and a word designating a fruit as the first: *apple pie*, *cherry pie*, *blueberry pie*. In all of these cases, the relevant conceptual relation is basically the same: a member of the category designated by the second noun *contains* or *is made from* members of the category designated by the first noun.

An Unfamiliar Phrase May be Understood by Analogy with Overlapping Familiar Phrases

If a Hearer had at least some of these more common phrases in long-term memory, they could be used to help interpret the new phrase, *grape pie*. That is, in attempting to match the new phrase against the other familiar phrases in memory, the Hearer would discover a number of partial matches, and they would lend weight to the interpretation in which the grapes are part of the contents of the pie. The new phrase is understood by **analogy** with familiar phrases. Note that this is no longer a case of pure compositionality because the phrase is not being understood purely in terms of the context-free meanings of the individual words. Additional knowledge about what other similar phrases mean is being used.

This possible use of analogy in interpreting phrases — "possible" because the idea that people actually do it this way is still somewhat controversial — has implications for what is in the lexicon. If a Hearer is to use more familiar phrases in interpreting new phrases, then those more familiar phrases have to be stored somewhere. So the assumption behind this idea is that the lexicon consists not only of words (that is, their forms and meanings), but also of common phrases (that is, their forms and meanings). It turns out that this isn't a big leap, since, as we'll see below, the lexicon will *have* to contain some phrases anyway. And storing frequent phrases in the lexicon can speed up comprehension. That is, rather than going to the trouble of combining the meanings of the component words, a Hearer could access the meaning directly from the combination. So, presented with a phrase like *student visa* for the first time, you could easily figure out what it means. But it would be even easier if you could just look up the meaning directly from the words. Something like this may be possible after you hear the phrase a few times.

So let's summarize the possible sources of information that a Hearer could use in interpreting a phrase. In the ideal, purely compositional case, the Hearer can get by with the meanings of the individual words in the phrase, that is, lexical conventions, and the grammatical rule appropriate for the particular grammatical relation about how to combine the meanings, that is, a grammatical convention. But the Hearer may also supplement pure compositionality with knowledge of the context and what is possible or likely in the world. Finally, the Hearer may rely on further lexical conventions, knowledge of other similar phrases and what they mean, by a process of analogy.

Idiomatcity

Exercise 5.3.2

What knowledge would a Hearer need to have to interpret the phrase *couch potato*, as in the sentence, *he never gets anything done; he's nothing but a couch potato*?

Expressions Differ in Terms of How Much of their Meaning Can be Derived from the Meanings of their Parts

Now consider an example that deviates even more from pure compositionality, *blackboard*. Note that even though this is written as one word, it clearly consists of the meaningful parts *black* and *board*. The question is: how much of a relationship is there between the meaning of *blackboard* and the meaning of its component parts? Certainly a blackboard is a sort of board, and it is normally black, though we might refer to a green board that you write on with chalk as a "blackboard" as well. But a blackboard is more than just a board that's usually black; it's a board that's usually black but also has the special function of providing a surface where writing or drawing in chalk can be done. The point is that even though the word *blackboard* consists of meaningful parts *black* and *board*, you wouldn't be able to figure out the full meaning of the word from those parts. This word goes beyond pure compositionality and must be stored in the lexicon so that Speakers will be able to produce it in appropriate situations and Hearers will be able to understand it when they hear or read it.

Contrast this with *black board*, written as two words. This is a relatively infrequent combination, one you may never have heard, in fact, so it is not likely to be stored in the mental lexicon. To understand *black board*, a Hearer could apply strict compositionality, looking up the meanings of *black* and *board* and combining them according to the rule for attribution.

In the case of *blackboard*, there are aspects of the meaning that cannot be predicted from the parts and the grammatical combination rule. Such a word is at least partly **idiomatic**. Idiomatcity is the tendency of phrases to take on meanings that go beyond the meanings of their parts. That is, idiomatcity is in opposition to compositionality. As we'll see below when we consider more extreme examples, idiomatcity is a matter of degree. But in all cases, the aspects of meaning that are not derivable from the parts of the phrase and that Speakers and Hearers are expected to know must be stored in the lexicon. That is, they are linguistic conventions. The extent to which languages are idiomatic, that is, the extent to which they deviate from pure compositionality, is a matter of considerable controversy. I'll focus on clearly idiomatic examples, but some researchers believe that extra meaning is inherent in a great number of frequent phrases and that the lexicon stores far more information about meaning than it would if language were more purely compositional.

Note that English actually distinguishes the compositional phrase *black board* from the idiomatic word *blackboard*. Not only is the idiom written as one word (a relatively trivial writing convention); it is stressed on the first element rather than the second, as would be normal for non-idiomatic adjective + noun phrases. This is also true for other similar examples: *greenhouse* vs. *green house*, *bluebird* vs. *blue bird*, *softball* vs. *soft ball*. But English is not completely consistent with this pattern. *Slow motion* has an idiomatic component; it refers not just to motion that is slow but to normal motion which is displayed in a slowed-down fashion. However, the stress is on *motion* rather than *slow*, and we write the expression as two words.

Are expressions like *blackboard*, *softball*, and *slow motion* words? We could choose to decide whether an English expression is a word on the basis of whether it is *written* as one word, that is, without spaces within it. But the writing conventions are not very consistent, and in any case they are a *reflection* of whatever it is that makes something a word rather than the basis for calling it a word. Besides, using the conventions regarding placement of spaces would not help at all with languages such as Japanese that have writing systems that do not use spaces to separate words or with the many languages that are not written at all.

One Way to Create New Words in English is By Combining Existing Words into Compounds

Clearly expressions such as *blackboard* and *slow motion* are like words such as *board* and *motion* in that there is something arbitrary about their form-meaning relationship that must be spelled out in the lexicon. Those that are stressed on the first element are also wordlike in another sense; English has a strong tendency to stress the first syllables of words, especially nouns. But these expressions differ from words like *board* and *motion* in that they have subparts which have *something* to do with their meanings. That is, they are considerably less arbitrary than words like *board* and *motion*. So are they words? The best answer seems to be "sort of". That is, like many other linguistic concepts, wordhood (whether or not something is a word) is a matter of degree. Just as there are "good" and not so good members of the category apple (a Red Delicious vs. a crabapple) and good and not so good

consonants ([t] vs. [w]), there are good and not so good words (*motion* vs. *slow motion*). Wordlike expressions such as *blackboard* that are made up of words are called **compounds**. Compound nouns in English include not only expressions with adjective and noun components but also expressions that started out as noun + noun phrases, for example, *coffee pot*, *tape recorder*, *headline*, and *surfboard*.

Sour Grapes, Cold Turkey, and Couch Potatoes May Not be Food at All

Now let's consider expressions whose meanings deviate even more from what we would predict from their parts. The phrase *live wire* can be used to refer to a wire that is connected to an electric power source; in this sense the phrase is close to purely compositional. But it can also be used to refer to a very alert or active person. Of course there is no way this second meaning could be predicted from the meanings of the words *live* and *wire*. In fact in a sense the phrase used with this meaning breaks the attribution rule, the rule that the phrase designates a subcategory of the meaning of the head noun, *wire* in this case. The phrase doesn't designate a subcategory of *wire* at all; it designates a subcategory of *person*.

So *live wire* is even more idiomatic than *blackboard* or *greenhouse*. Other examples of adjective + noun combinations of this sort are *sweet tooth*, *sour grapes*, and *cold turkey*. Hopefully you recognize the idiomatic uses of these expressions as examples of the sorts of semantic extension that you learned about in Chapter 2. In each case there is an original meaning derived directly from the meanings of the individual words, and this provides the basis for an extended meaning. For *live wire*, the extension is metaphoric since it is based on the similarity between a kind of wire and a person. For *sweet tooth*, the extension is metonymic since it is based on the part-whole relationship between a certain kind of tooth and a person with a craving for sweets.

English also has a number of highly idiomatic noun + noun phrases. A *couch potato* is not a potato at all but a person who spends a lot of time on the couch. A *bookworm* is not a worm but a person who loves books. Both of these are examples of metaphoric extensions. Some insults are noun + noun phrases involving both metaphor and metonymy. A *chowder head*, for example, is a person whose head seems as if it is full of soup. This is metaphoric in the sense that the person's brains are similar to chowder and metonymic in the sense that the person is referred to through a body part.

Though expressions such as *live wire* and *couch potato*, like *blackboard*, are wordlike in the sense that their meanings have to be stored in the lexicon, many of them seem less like words than compounds such as *blackboard* do. For the adjective + noun expressions in this category, English does not usually put the stress on the adjective as it does for compounds such as *blackboard*. Perhaps having these idiomatic phrases sound like ordinary compositional phrases rather than like words makes them more colorful; hearers are reminded of their original meanings as compositional phrases. In any case, we have been thinking of the lexicon as a place to store *words*. Now it is clear that some whole phrases need to be there as well. We need a new, more general term to refer to units stored in the lexicon; the usual term for this is **lexical items**. Lexical items include words such as *board* and *apple*, compounds such as *blackboard* and *softball*, and idiomatic phrases such as *live wire* and *cold turkey*. As noted above, they may also include phrases that are compositional but appear frequently, phrases such as *apple pie* and *green apple*.

Let's sum up what we've learned about how the meaning of a phrase depends on the meanings of its parts. For at least some phrases, language appears to be completely compositional: a Hearer can arrive at a meaning for the phrase based on stored meanings for the parts and a grammatical combination rule, supplemented, if necessary, by knowledge of the world and of the context where the phrase occurs. But this purely compositional picture breaks down in three sorts of situations. First, even if the meaning of a phrase *could* be derived from its parts, it may be more efficient to store it in the lexicon as a unit if it is frequent (e.g., *pretty face*, *rough road*, *apple pie*, *course grade*). Second, the Hearer may interpret the expression in part by analogy with other phrases that are stored in the lexicon. For example, *lip ring* might be understood by analogy with *earring* and *nose ring*. Finally, phrases may have a more or less idiomatic meaning, an aspect of their meaning which is not predictable from the parts of the phrase. Because there is no way to know these aspects of meaning, idiomatic expressions must be stored in the lexicon. They include words or phrases with a relatively limited degree of idiomaticity such as *bluebird* and phrases that are highly idiomatic such as *cold turkey*. For this last group, many people may not even be aware of the connection between the original meaning and the idiomatic meaning.

This page titled [5.3: Compositionality and Idiomaticity](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.4: Problems

Compositionality and Idiomaticity

Note: if you're not a native speaker of English, you might want to collaborate with a native speaker for this assignment since it requires understanding idiomatic expressions.

The purpose of this assignment is to show that you understand compositionality and idiomaticity, that is, that you can distinguish what aspects of the meaning of a phrase can be derived from the meanings of the words and the grammatical relation and what aspects of the meaning come from somewhere else.

For each of the following noun + noun phrases (or words), say first whether it is familiar to you. If it is familiar, say what the conceptual relation between the two categories is and what, if anything, is idiomatic about the meaning. If it is unfamiliar, say what you can know about the meaning from composition alone, say what conceptual relation you think holds between the noun meanings, and say how you guessed the relation: using world knowledge or by analogy with other noun + noun phrases.

For the conceptual relations, choose from the following list of possible relations (A and B represent members of the categories designated by the first and second noun respectively). If you are not sure of the interpretation, pick the relation that seems the most plausible.

1. B is a part of A.
2. B contains or is made from A.
3. B is used for A.
4. B is located at/near A.
5. B is a location or an occasion for A.
6. B resembles A.
7. B deals with A.
8. Some other relation (please specify what).

Examples:

- *tooth decay*

This is a familiar phrase for me.

The conceptual relation is 4 or 1.

There is nothing idiomatic about the phrase, but from world knowledge, I know a lot about how this kind of decay differs from, say, decaying fruit or meat.

- *potato bath*

This is an unfamiliar phrase for me.

From composition alone, I know only that this designates a subcategory of bath that has some conceptual relation to potatoes. The relation could be 2; that is, it could designate a sort of bath that potatoes are added to (maybe because of what they do for the skin).

I guessed what the phrase means using my world knowledge: how people sometimes add plants (rose petals, citrus fruits) or other food items (egg) to bath water or to shampoo.

- *seahorse*

This is a familiar word for me.

The relation is 4, but this phrase is idiomatic because it doesn't designate a subcategory of horse at all, but rather a fish that looks a little like horse.

1. *pancake*
2. *carrot cake*
3. *turnip cake*
4. *turnip beetle*
5. *turnip scientist*
6. *turnip face*

7. *tennis ball*
8. *handball*
9. *water polo*
10. *web site*
11. *webhead*
12. *web party*
13. *road hog*
14. *hall hog*

This page titled [5.4: Problems](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

6: Sentences

- [6.0: Introduction](#)
- [6.1: States and Events](#)
- [6.2: Situation Schemas and Semantic Roles](#)
- [6.3: Constituency and Noun Phrases](#)
- [6.4: Subjects](#)
- [6.5: Direct Objects](#)
- [6.6: Adjuncts](#)
- [6.7: Sentence Functions](#)
- [6.8: Problems](#)

This page titled [6: Sentences](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.0: Introduction

In the last chapter, we saw how words are combined to yield phrases with new meanings. But so far we have only looked at words or word combinations that refer to things. People also talk about situations, about the states that things are in (*I feel nauseous, this book is a gem*) and about events involving things (*I tripped, I read the book you lent me*). To do this, they put words together in phrases such as *the book you lent me* and combine these with other phrases to make sentences. In doing so, they relate the parts, or roles, of the situation they are describing with the roles in the sentence used to describe it. As with the other aspects of language we have studied, each language has conventions for how these two kinds of roles are associated with one another and for how the sentence roles themselves are marked. Many of these conventions are tied to verbs, which organize the structure and the meaning of sentences. In this chapter, we'll see how different subtypes of verbs behave and how the properties of verbs and sentences vary across languages. We'll also see how sentences not only describe actual situations in the world but also allow speakers to get information about these situations (*did you read that book I lent you?*) or to cause them to happen (*please return that book I lent you*). When you're done with this chapter, you should have a better idea what sentences are, how they work, and how they allow us to talk about what's going on in the world around us.

This page titled [6.0: Introduction](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.1: States and Events

Now that our Lexies have the capacity to put words together in novel combinations, they have the full range of the basic features of language. Since we'll be focusing on grammar in the rest of the book, let's change their name to Grammies at this point.

Although these Grammies can put words together in novel combinations to refer to things that combine conceptual features of different pre-existing concepts, they are still limited in one important way. They can only *refer* to things, that is, point to things with language. They still can say nothing *about* things. For example, they refer to a particular rock, using a phrase like English *the big rock*, but they can't say where it is. They can refer to a particular person, using a phrase like English *the young man* or *Clyde* or *you*, but they can't say where the person is or what they are doing. They can use a noun such as *rock* or *person* to designate the category of a thing or an adjective such as *big* or *moldy* to designate an attribute of a thing but they can't make a claim about a thing belonging to a category or having a property or ask whether a thing belongs to a particular category or has a particular property.

To allow speakers to do these things, modern languages have sentences, which are the topic of this chapter. In this first section of the chapter, we'll look at the kinds of situations in the world, and in the minds of Speakers and Hearers, that sentences designate.

Though I'll be using some English sentences to illustrate these kinds of situations, it's important to make the distinction between the language we use and what that language is about, just as it was with nouns and the categories of things they designate. In this section, our only concern will be the concepts that sentences are about, not the sentences themselves, so this section is not really about English or any other particular language.

States

Exercise 6.1.1

An object such as an apple or a tiger has a set of relatively stable properties (or attributes) and a set of relative temporary properties. What might these two kinds of properties be for apples and tigers?

People May Notice and Talk About Relatively Temporary or Unpredictable Properties of Things

Consider the photograph below and what an observer of the scene might see in it. First, there are several things in the scene, the boy, the soccer ball, the book, the grass. Remember from the section on reference in the Meaning chapter that what makes an object such as a soccer ball an object is that it has boundaries around it and a stable set of properties, its shape, its size, its weight, its composition, the pattern of colors on its surface, the way it feels when you press on it. But an object such as a ball can also have some relatively temporary properties, and these are the sorts of things that we might notice about it. For most objects, location is such a property. In the scene in the picture, we might notice that the ball is on the grass, that it is next to the book, that it is to the boy's right. Similarly, we could notice the location of the book or the boy.



For people there are many relatively temporary properties of interest in addition to location. The orientation of their body and the parts of their body is one; we might notice not only that the boy in the picture is on the grass but that he is lying, rather than standing or sitting, there. And we might notice that he is resting his chin in his right hand. Facial expression is another temporary property of people that we pay attention to; we might notice that the boy is smiling or that he has a dreamy look on his face.

Even though these aspects of the scene are not permanent properties of the objects in the scene, they are still relatively stable; that is, they continue for some length of time without changing significantly. We call such situations **states**. The fact that the book is next to the ball, the fact that the boy's left hand is resting on the book, the fact that the boy is lying on the grass are all states. Since the more stable properties of objects are also lasting in this same sense, we can also consider them to be states, though we don't

notice or call attention to them as often. So the fact that the soccer ball is black and white might not attract our attention because most soccer balls are. But we might notice permanent properties of the boy that are not properties of all boys, for example, the shape of his nose, his mouth, and his head or the color of his skin and his eyes.

For now the important point to note about states, in addition to their relative stability, is that they are always states *of* something. That is, when we observe the world, we see the things in the world as being *in* particular states. Some states are states of just one thing; the fact that the boy is lying down is a property only of the boy. But states may also relate more than one thing; the fact that the ball and the book are next to each other is a fact about two different things. The things that a state is concerned with are called the **participants** of the state. So the boy is the only participant in the lying state, and the book and the ball are both participants in the `next_to` state that relates them. States that have more than one participant are called **relations**.

Events

Of course not all aspects of the world are as stable as the orientation of the boy's body and the relative position of the ball and the book in the photograph. More often what we notice in the world around us is what does *not* stay the same. For example, look at the photograph below. If we were watching the scene shown in the photograph, we would probably notice the movement of the boy's arm and the ball leaving the boy's hand and traveling toward us. Each of these involves a *change* of state: the orientation of the boy's arm changes; the location of the ball relative to the boy and relative to us changes. We refer to such situations as **events**. Events differ from states in involving change, in being unstable rather than stable.



Like states, events have participants, and like states, they may relate more than one participant. So in the picture above, the boy and the ball are participants in the throwing event. Since a throwing event has two participants in it, the person doing the throwing and the thing that is thrown, throwing is an example of a relation. Since states and events are similar in many ways, it will be convenient to have a term that includes both of them; I will call them **situations**.

Construal

Exercise 6.1.2

Each of the following English expressions is used to designate a state. What's the corresponding expression that designates the change of state leading to that state?

- *wear*
- *be standing*
- *be asleep*
- *be awake*

But the line between states and events is not completely clear-cut. In particular, what counts as a state and what counts as an event depends on the observer. Consider a car traveling down a highway. If we are focusing on the movement of the car relative to the ground beneath it or relative to the places alongside the highway, we conceive of what is going on as an event, the event of the car's moving along the highway or passing buildings and signs. But if we are focusing instead on the fact that the car is staying between the middle and the edge of the road, we conceive of what is going on as a state, the state of the car's being on the road or within a particular lane.

There is another way in which the same situation in the world may be viewed as either a state or an event. Given an object in a particular state, we can either focus on that relatively lasting state or on the change of state (that is, event) that resulted in the state. For example, for a door that is open, we can focus on the state of the door being open or on the opening of the door that led to that

state. Or for an animal that is dead, we can focus on the state of its being dead or on the dying that led to that state. (For the expressions in the box above, the corresponding change-of-state expressions are *put on*, *stand up*, *fall asleep* or *go to sleep*, and *wake up*.)

We Can Speak of States and Events Either as Situations Existing in the World or as Ways We Have of Construing What Exists in the World

The important point here is that a given scene may be conceived of, or **construed**, differently. This is important because, as we will see, languages offer Speakers different options for representing the different construals they have of situations.

Different degrees of attention to different things in a scene can also lead to different construals. In the scene in the picture above, we could focus exclusively on the ball, ignoring the boy. This would be reflected in a sentence such as *the ball is moving fast* or *the ball has just been thrown*. On the other hand, we could focus exclusively on the boy, ignoring the ball. This would be reflected in a sentence such as *the boy has just swung his arm* or *the boy is throwing*. Again we'll see that languages have different options for how particular participants in an event can be highlighted or downplayed, depending on the Speaker's construal of the situation.

So far the examples have assumed that a person thinks or talks about a situation during that situation, but we can also think or talk about a situation at other times relative to it. We can talk about a state while it is true (*the ball is next to the book*), after it was true (*the ball was next to the book*), or before we assume it will be true (*the ball will be next to the book*). We can talk about an event while it is going on (*the boy is throwing a baseball*), after it has happened (*the boy threw a baseball*), or before we assume it will happen (*the boy is going to throw a baseball*). I will not have much to say about these time-related differences and how they are reflected in language, but most languages apparently have ways to make these distinctions.

Like us, our Grammies notice states and events in the world. This is because states and events matter to them. The angry expression on the face of another person tells them it's best to avoid that person for awhile. The fact that clouds are forming in the sky leads them to expect rain. The roaring sound they hear tells them there is a tiger nearby. But not all members of the tribe may have noticed a particular significant state or event, and in this case there is reason to be able to communicate it. Just as we saw that it was impractical to have a separate word for every combination of attribute and thing category, it is impractical to have a separate word for every possible state or event. Rather, as with nouns, our Grammies come up with the idea of a separate word for each category of state or event, words such as *sit* and *throw*. But each of these state or event categories can have different types of participants. This means that it will often be necessary to combine the state/event words with words that refer to the participants. That is, it will be necessary to have *sentences*.

This page titled [6.1: States and Events](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.2: Situation Schemas and Semantic Roles

In this section, we'll look at categories of states and events. We'll see that these categories are more complicated than categories of things because they are in a sense *about* things. Without things, there can be no notion of state or event. Part of characterizing a particular state or event category is characterizing what sorts of things it can be about, how many things there are (two or three in the case of a relation), and what role the things play in the whole situation.

As with categories of things, most languages have a category of words that designate categories of events, and, less often, categories of states. These words are called **verbs**. Examples are *kiss*, *write*, *like*, *smell*, and *go*. In this section we'll be guided in our consideration of the different categories of states and events by some of the English verbs used to designate them, but our focus will still be on the semantics, that is, on the kinds of things that all languages need to be able to designate. In the succeeding sections, we'll be looking at how events and states are actually designated in sentences in English and other languages.

One word of warning. While there is agreement on the basic ideas in this section, there is no general agreement on the details. The particular categories described here are those that I happen to think do the best job of characterizing the states and events that languages designate.

The Do_to Schema

Exercise 6.2.1

What do the meanings of the following verbs have in common: *kiss*, *push*, *punish*, *carry*, *bathe*? (Think about the participants in the events and how they relate to the action and to each other.)

Let's begin by considering the event that we might designate with the following sentence.


1. *Lois kissed Clark.*

This event has two participants associated with it, Clark and Lois, and they are associated with each other in particular way through a relation we can call kiss. This particular kissing event shares a lot with other instances of kissing. We can abstract over all of these instances of kissing by thinking of a kissing **schema**, a kind of template for kissing events.

Each Kind of Event or State Carries with It a Set of Expected Participants Who Play Particular Roles in It


We have already seen an example of a schema in the context of the discussion of utterances. Remember that each utterance has a number of *roles*, form, Speaker, Hearer, location, and time. In the kissing schema, there are two main roles, one for each participant in the event. We can call these roles the *kisser* and the *kissee* of the event. I will refer to roles within the schemas for states and events as **semantic roles**. Semantic roles are defined by their meaning, not by their linguistic form. Later in this chapter, we will see that there is a set of sentence roles, including subject, for example, that are defined by their form. It is important to keep the two kinds of roles distinct.

As with utterances (which after all are a kind of event), I will diagram situation schemas using boxes to show the roles. The figure below shows the kiss schema.



In the schema for the KISS relation, the two roles are not filled by any particular people; we just consider them in terms of the parts they play in the relation. KISS is an asymmetric relation; that is, the parts played by the people filling the two roles are quite different. The KISSER is the participant who consciously initiates the action, who makes contact with the KISSEE. The KISSEE is the destination of the action. While the KISSEE is affected by the action, this participant does not play a conscious part in the event.

In a particular instance of kissing, a particular kissing event, the two roles are filled by particular people. So sentence 1 above describes the event that we can diagram as follows.




Now consider another event, the one described by the following sentence.

2. *Lois pushed Clark.*

Again there is a relation, push, with two roles, which we can call pusher and pushee. And again the relation is asymmetric: the pusher initiates the action and makes contact with the pushee, and the pushee is affected by the action but plays no conscious part in it. The details of how the two push roles fit into their relation differ from how the two kiss roles fit into their relation, of course. A pusher exerts more force than a kisser, and a pushee normally moves more as a result of the action than a kissee does.

But clearly kiss and push share some properties. Both have two core roles, including one that consciously instigates the action and another that is passively affected by the action. So we can be even more abstract, thinking of a schema for the larger category of actions of this type. Let's call this schema *do_to*, and let's call the two roles in *do_to* the agent and the patient. The agent is the conscious initiator of the action, and the patient is the participant that is passively affected by the action. We can diagram the *do_to* schema as follows.



Other examples of relations that belong to the *do_to* schema include those that we would express in English using the verbs *hit*, *touch*, *injure*, *cure*, *carry*, and *bathe*, as well as more abstract relations expressed by verbs such as *appoint*, *insult*, and *abandon*. Note that the agent must be animate (that is, human or animal), or at least construed as animate; otherwise, it can't consciously initiate the action. But the patient can be either animate or inanimate, though some relations, such as *appoint* require it to be animate or even human.

To make it clearer what *do_to* involves, let's consider some event types that are *not* instances of this general category. Events designated by verbs such as *love*, *see*, *know*, and *understand* do not belong to the *do_to* schema both because there is no conscious initiator of the process and because there is not necessarily a participant that is affected by the process. That is, they have no agent and no patient. For example, we normally see or know something without a conscious effort on our part, and the thing that we see or know is not affected by our seeing or knowing. The event described in the sentence *Clark became a reporter* is not an instance of *do_to* because there is only one participant, and the state described in the sentence *Clark resembles Superman* is not an instance of *do_to* because there is no event to be consciously initiated by a participant in the first place.

It is important to note at this point, however, that a category such as *do_to* is similar to the other kinds of categories we've seen in the book. That is, there will be prototypical instances of the category (for example, the event designated by the sentence *Lois hit Clark*), and other cases where it is not so clear which category an event belongs to. We will also see cases where an event can be described as a blend of more than one category.

The Do Schema

Consider the events designated by the following sentences.

3. *Clark worked all day.*
4. *Lois played all day.*
5. *Jimmy just eats.*
6. *Perry was singing.*

In each of these cases, there is a conscious initiator of some process, that is, what we could call an agent, but no second participant that is affected by the process. We obviously cannot call these events instances of *do_to*, but they resemble *do_to* because of the presence of an agent. I will call the schema for such events, as well as for *do_to* events, the *do* schema. This more general schema

is characterized by the presence of an agent, that is, something animate that consciously initiates some process. When there is also a patient, the event belongs to *do_to*, a subcategory of *do*.

The Experience Schema

Exercise 6.2.2

Each of the following pairs of sentences describes the same situation. How does the construal differ for the two sentences?

- *Lois is looking at the bird. Lois sees the bird.*
- *Lois is smelling the rose. Lois smells the rose.*
- *Lois is listening to a rustling in the backyard. Lois hears a rustling in the backyard.*

Sentences that Look Similar in English May Describe Quite Different Kinds of Events or States

An event or state may involve two participants, as in the *do_to* schema, but with a very different relation between the two than for *do_to*. Consider the situations designated by the following sentences.

7. *Clark saw Lois and Jimmy.*
8. *Clark heard something strange.*
9. *Clark likes Lois and Jimmy.*
10. *Clark knows Perry.*
11. *Clark understood the word.*

In each of these cases, there is an animate participant, Clark, who experiences something that he doesn't directly control and therefore doesn't initiate, so he is not an agent. I'll call this participant the experiencer. The second participant can be almost anything, and, unlike the patient in the *DO_TO* schema, it is normally not affected by the experience, though it does in some sense lead to the experience. I'll call this participant the theme. We'll see this term later. In general, it is used to refer to a role that is central to the meaning of the state or event but does not participate directly in it. Situations like these belong to what I'll call the experience schema.

One point to note about instances of experience is that, though they are probably best viewed as states, they also have properties of events. To the extent that we see these experiences as unchanging situations, we are viewing them as states. To the extent that we see them as mental processes that take place or we focus on the beginning of the process, we are viewing them as events.

Note how the meanings of sentences involving sensory experience such as 7 and 8 above differ from similar sentences in which the experiencer plays a more active role.

12. *Clark looked at Lois and Jimmy.*
13. *Clark listened to the music.*

In these cases, Clark is consciously attending, either visually or auditorally, to something. That is, Clark is the agent of the process. However, in these cases the other participant is not affected by the action of the agent, so we cannot call it a patient. Rather the second participant can be seen as a kind of target for the agent. I will again call it the theme to emphasize its similarity with the theme of the experience schema. In both cases, this participant is providing the input to an experience. In a sense, then, sentences 12 and 13 designate events that are a blend of the experience and *do* schemas.

The Happen and Move Schemas

Exercise 6.2.3

Each of the following pairs of sentences could describe the same situation. How does the construal differ for the two sentences?

- *The water boiled. Clark boiled the water.*
- *The ball rolled. Clark rolled the ball.*
- *The shirt got clean. Clark washed the shirt.*

Just as we can have events with an agent but no patient, we can have events with a patient but no agent. Consider the events designated by these sentences.

14. *The lake froze.*
15. *My watch stopped.*
16. *This shirt is fading.*

In these examples, there is a single participant that undergoes a kind of **change of state**. Because this participant is directly affected by the process, it is a patient. I will refer to this as the happen schema; something happens to the patient. Note that we can see *do_to* as a subcategory of happen as well as a subcategory of do because it has a patient as well as an agent role.

When the patient of a happen event is animate, this schema resembles the experience schema, except that there is no second participant.

17. *The cricket died.*
18. *I got sick.*

An important subcategory of happen involves a change in the location of the patient. I'll call this the move schema. It is illustrated by the following sentences.

19. *The ball rolled off the ledge.*
20. *A rock fell in my lap.*
21. *The tractor crossed the field.*

In addition to the patient, the move schema has three other roles that specify the movement. When something moves, it has a place where the movement begins, the source; a place where the movement ends, the goal; and a way of getting from the one to the other, the path. For a given construal of an instance of move, any of these may be seen as worthy of attention. Sentence 19 above mentions the source, sentence 20 mentions the goal, and sentence 21 mentions the path.

When the patient of a move instance is animate, it often has control over the movement and the event becomes a subcategory of the do schema rather than the happen schema. This possibility is illustrated in the next two sentences.

22. *Clark jumped off the ledge.*
23. *Lois crossed the field.*

It is Possible for the Same Participant to Fill Two Different Roles at the Same Time

In such cases the patient is also an agent; for example, in the event described by sentence 22, Clark is both affected by the process and the initiator of the process. But there may also be an agent that is separate from the patient; in such cases we have instances of *do_to* as well as move. This possibility is illustrated in the next three sentences.

24. *Clark put the book on the desk.*
25. *Clark took the book from the table.*
26. *Clark pushed Simon off the ledge.*

The Transfer and Information_transfer Schemas

Another common schema combines move and *do_to*, but in this case there is a second animate participant, the recipient, which functions as the goal for the patient. The source is also animate. I'll call this the transfer schema. It is illustrated in the next four sentences.

27. *Lois gave Clark a CD.*
28. *Lois gave a CD to Clark.*
29. *Lois sold Clark her computer.*
30. *Clark borrowed a computer from Lois.*

A transfer is actually quite complicated. Before the event there is a state in which the source of the transfer is in control of the patient. For example, for sentence 27 Lois owns the CD before she gives it to Clark. After the event who controls the patient has changed; it is now the recipient of the transfer. For example, for sentence 27 Clark owns the CD after Lois gives it to him. Either

the source or the recipient can be the agent of the event. In sentences 27-29, the agent is the source; in sentence 30, the agent is the recipient.

Another type of transfer is also possible, one involving the "moving" of information from the agent to the recipient. In *information_transfer*, there is really no patient, however, because the information itself is not affected by the transfer. I'll refer to the information as the theme. And the agent doesn't really lose the information in the transfer, as is the case with the transfer of physical objects. Here are some examples of *information_transfer*.

31. *Lois told Clark her password.*

32. *Lois informed Clark of the meeting.*

33. *Lois asked Clark the time.*

Sentence 33 is a little complicated. Here the information that is transferred, the theme, that is, is not the time itself but rather Lois's desire to know the time.

The Be Schema

Exercise 6.2.4

How is the noun *doctor* used differently in the following two sentences? That is, how is the Hearer expected to use the information provided by the word (the category of people, doctor)?

- *A doctor came in the room.*
- *Carla is a doctor.*

Finally, we have what is in a sense the simplest schema, the one for prototypical states. I will call this the *be* schema and will refer to the participant that is in the state as the theme. Instead of verbs for each particular type of state, the different states are defined mainly by the nouns and adjectives we already discussed in Chapters 2 and 5. Here are some examples.

34. *Lois is sick.*

35. *Lois is a Hoosier.*

36. *This is a keyboard.*

37. *The keyboard seems sticky.*

38. *The keyboard smells funny.*

Nouns are Not Only Used to Refer to Things. They May also Designate a State that a Given Thing is In

Though we are not focusing in this section on the forms that languages use for the different schemas, it is important to see how the nouns and adjectives at the ends of these sentences are used differently than they were in the examples we saw in Chapters 2 and 5. Previously we only considered nouns and adjectives as a means of referring, of pointing to particular individuals in the world or the mind of the Hearer. But *a Hoosier* in sentence 35 is not meant to point to particular instance of the category hoosier. Rather it functions to call the hearer's attention to a state that Lois is in, the state of belonging to the category hoosier.

Similarly, compare *sick* in sentence 34 with *sick* in the sentence *the sick child is improving*. In this second sentence the adjective has an *attributive* function; it narrows down the range of possibilities for who is being referred to so the Hearer can figure out what child is meant. But in sentence 34 *sick* is used to call the attention of the Hearer to the state that Lois is in, the state of being sick. *Sick* in sentence 34 and *Hoosier* in sentence 35 are said to perform a **predicative** function rather than an attributive or referential function.

There are two very important subcategories of the *be* schema that are relations. One of these involves possession or control of one participant over another. I'll call this the *have* schema. Here are some examples.

39. *Lois has a digital camera.*

40. *Clark has a lot of cash.*

I'll refer to the two roles in *have* states as the possessor and the theme.

Another important subcategory of *be* is also relational; the state concerns the spatial or temporal relation between the two things. We can call this schema *be_at*. Here are some examples.

41. *Lois is behind Clark.*
42. *Lois is next to the car.*
43. *Bloomington is south of Indianapolis.*
44. *The concert is on Monday.*
45. *The concert is after the class.*

There are two participants in the states designated by each of these sentences. It is normally possible to reverse the way the relations are expressed. So corresponding to sentences 41 and 43, for example, we have the following possibilities.

46. *Clark is in front of Lois.*
47. *Indianapolis is north of Bloomington.*

Such reversals do not change the relation between the two participants at all. Hence there is no good reason for distinguishing the two participants, and I'll refer to both as themes. Of course this does not mean that sentence 41 is completely synonymous with sentence 46. The difference has to do with the perspective that the Speaker is taking on the state, which participant is being singled out to say something about. This sort of factor won't concern us in this chapter though.

Note that when the relation is a temporal one, as in sentences 44 and 45, the two "things" that are related are actually events, in 44, the concert and a particular Monday. But by referring to the events with nouns, the Speaker of the sentence is treating them as abstract things that are "located" at different points in time.

Other Semantic Roles

Exercise 6.2.5

How does the semantic role of Clark differ in the following three sentences?

- *Lois sent Clark a message.*
- *Lois did Clark a favor.*
- *Lois did Clark an injustice.*

Participants Can Vary in Terms of How Central They are to the Particular Kind of Event or State

We can identify a number of other semantic roles that can occur together with the basic schemas described above. The roles we have seen so far are those that are central to each of the schemas; I will refer to the things playing these roles as the **core participants** in the states or events. But there is also the possibility of some relatively **peripheral participants**, those that are not central to the state or event that is being thought or talked about. For example, all states or events can have allocation or time. Here are some examples that make these roles explicit.

48. *Lois kissed Clark behind the barn.*
49. *Lois kissed Clark before the class.*
50. *Lois is sick around Jimmy.*
51. *Lois was sick on Tuesday.*

In instances of the do schema, especially the do_to schema, there is the possibility of a participant that is *used* by the agent to achieve the process. This role is called the instrument. Here are some examples illustrating instruments.

52. *Lois broke the window with a stick.*
53. *A stick broke the window.*
54. *Lois used outdated software to write her report.*

The sticks referred to in sentences 52 and 53 play the same role in the breaking events; both are instruments. The main difference is that the agent is not mentioned in 53. But we can assume that even in this case, there was an agent behind the action; that is, the stick didn't break the window of its own accord. In sentence 54, the software is the instrument for the writing of the report.

In instances of do, and again especially do_to, another possible role is an animate participant that *benefits* from the action, a beneficiary. Here are some examples in which Clark is the beneficiary.

52. *Lois fixed the window for Clark.*

53. *Lois fixed Clark a great meal.*

54. *Lois sent Perry a letter for Clark.*

Note that in sentence 56 there is both a recipient, Perry, and a beneficiary, Clark.

An event may also *harm*, or otherwise adversely affect, someone other than the patient. I'll call this participant the sufferer (there is no generally agreed on name). Here are some examples, but, as we'll see later, English does not offer as direct a way as languages such as Japanese and Amharic to refer to the sufferer.

55. *Lois lost her checkbook.*

56. *Lois and Clark went and eloped on me.*

In sentence 55, the checkbook is the patient, and the loss of the checkbook has an adverse effect on Lois. In sentence 56, Lois and Clark (together) are the agent, and their action has an adverse effect on the speaker.

A final possibility is a semantic role that is not really peripheral but that results from viewing an event as having a more direct and a less direct cause. The agent is the participant mostly directly responsible for the process, but there may be another participant that causes the agent to carry out the process. I'll call this the causer. Here are some examples.

57. *Lois got Clark to write the report.*

58. *Lois had the report written.*

Notice that in sentence 58, no agent is mentioned — we do not know who actually wrote the report — but we know that somebody other than Lois did it, so Lois must be seen as the causer rather than the agent.

Summary of Semantic Roles

Let's summarize the different semantic roles we have met in this section. Remember that for each role there are clear-cut, prototypical examples and other, less clear cases where a participant may share properties with multiple roles.

agent

The core participant that **directly causes the event**. Prototypically human, or at least animate, in which case there is conscious intention behind the act, and the question *what did agent do?* is appropriate.

It is easily confused with an experiencer, which is also animate. The difference is that an experiencer isn't consciously behind the experience; it just happens to the experiencer.

patient

The core participant that **is affected by the event**. That is, the patient changes in some way. A patient can be animate or inanimate.

For "going" events, the patient can also be the agent (the agent moves itself). Participants in addition to the patient may also be affected by the event; these include recipient, beneficiary, and sufferer.

experiencer

The core participant that **experiences the event** or state without having direct control over it. Since experience requires a nervous system, an experiencer must be animate.

Since the experiencer is affected by the state or event, it can be seen as a subcategory of patient. Some kinds of events, such as "looking" and "listening", combine experience with conscious action; I am treating the participants that cause these events as agents, but they overlap with experiencers.

possessor

The core participant that **possesses or controls a second participant** in a state. The other participant is the theme. The possessor is prototypically animate.

theme

A core participant that **neither causes the state or event nor is directly affected by it**. The notion of theme is a weakness in this account of semantic roles because it is defined almost completely in negative terms; it is really a place to put core

participants that don't belong to some other role. Spatial relation and temporal relation states have two themes.

causer

A core participant that **causes an agent to cause the event**, though the agent may not be explicitly referred to. A causer is normally human.

source, goal, path

Three roles that **define the spatial aspects of a move, transfer, or information_transfer event**. In a transfer or information_transferevent, the source and goal are the animate participants in the event; the goal is the recipient.

recipient

The animate goal of a transfer or information_transfer event. In a "taking", the recipient is also the agent.

instrument

The participant that **is used by the agent to achieve the event**.

Where there is a separate causer and agent, the agent is a sort of instrument for the causer, but the agent can also have its own instrument: *Lois had Clark take off the tire with the wrench*.

beneficiary, sufferer

Animate, relatively peripheral participants that **are affected either beneficially or adversely by the event**.

These need to be distinguished from patients, which are core participants but can also be animate and affected either beneficially or adversely, for example, in *Carla killed Frank*.

location, time

Inanimate, peripheral participants that **situate the state or event in space or time**.

The set of schemas and semantic roles that I have described is not just an arbitrary way of dividing up the states and events that people think and talk about. The distinctions between the different schemas and roles are reflected in the forms that different languages use to designate different kinds of events and states. This is the topic of the rest of this chapter.

This page titled [6.2: Situation Schemas and Semantic Roles](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.3: Constituency and Noun Phrases

Constituents

We have seen that the situations in the world that people tend to notice and talk about are construed as states or events involving a small number of participants. We can expect the sentences that designate states or events to have a word that specifies the kind of state or event, usually a verb, and a phrase for each of the participants that the speaker chooses to refer to. Each of these makes up a part, **or constituent**, of the whole sentence. Before we can look at how different languages put sentence constituents together to designate states and events, you will need to know more about what makes up a constituent. In this section we'll be focusing on the constituents that refer to the participants in the state or event. And for these constituents we'll be focusing on their internal structure, what words can make them up.

Let's return to where we left our Grammies. Once they have the insight that it is to their advantage to be able to talk about events and states, the next step might be to come up with words for the different categories of states and events, verbs, that is. In the simplest case, then, we can imagine describing events and states with verbs alone, for example, *recovered* for a recovery event or *gave* for a giving event. In some modern languages, such as English, the grammatical conventions don't permit sentences like this consisting only of a verb, but in languages such as Japanese, they are possible. So the following are grammatical Japanese sentences.

1. *naotta* 'recovered'
2. *ageta* 'gave'

The first sentence means that somebody recovered; the second means that somebody gave something to somebody else.

But even in Japanese, these one-word sentences can only go so far. Unless it is clear from the context, the hearer may have trouble figuring out who the patient (the person recovering) is for sentence 1 and who the agent and the recipient and what the patient is for sentence 2. Such sentences are massively ambiguous and only interpretable when the context makes it very clear who the participants in the events are. Something else is needed.

Noun Phrases

Exercise 6.3.1

Look at the function of the word *Clark* and the phrase *the guy who is out to save the world* in these sentences.

1. *Clark got sick.*
2. *The guy who is out to save the world got sick.*

Even though these expressions look quite different, how are they similar in their function?

Compositionality Applies to Sentences as Well as to Modifier + Noun Combinations

Faced with this sort of ambiguity, the Grammies hit upon the same idea that they had used earlier on to make it easier to refer to things, the principle of compositionality. Remember from the last chapter how this works. Words combine into phrases, and the meaning of the whole phrase depends on the meanings of the individual words. Here's a simple English example.

3. *Fred recovered.*

We know from Chapter 2 that a proper noun like *Fred* can refer directly to an individual, making this the easiest way to make the Hearer aware of who is being talked about. By compositionality, we get the meaning of sentence 3 by figuring out the conventional meanings of each of the words, *Fred* and *recovered*, which are available in our lexicon, and then applying the appropriate grammatical rule to combine the meanings. We'll talk about the grammatical rule in the next section. For now, the important point is that the meaning of sentence 3 can be derived from the meanings of *Fred* and *recovered* and what we know as speakers of English about the grammar of English sentences.

But we also know from Chapter 2 that it is impractical to have a name for every thing we might want to refer to, and in any case, no speaker can know all of the names there are. So if the speaker of sentence 1 hadn't known Fred's name, the sentence would have come out differently, perhaps like this one.

4. *The teacher recovered.*

Even though the phrase *the teacher* in sentence 4 and the word *Fred* in sentence 3 consist of completely different words, they are doing the same work in the two sentences, referring to a particular individual. But there are many other ways the Speaker could refer to Fred, depending on the utterance context, what the Speaker knows about Fred, and what the Speaker believes the Hearer knows about Fred. Here are a few possibilities.

5. *He recovered.*

6. *That teacher recovered.*

7. *My teacher recovered.*

8. *That tall teacher recovered.*

9. *The teacher behind me recovered.*

10. *The teacher who got fired recovered.*

Expressions Referring to Things Have Their Own Internal Structure

In each sentence the part before *recovered* refers to the person who recovers (the patient). These phrases do not all mean the same thing, but in one sense they all *do* the same thing. And with respect to just form and not meaning, we can see that they all share the property of being possible phrases before a verb like *recovered* in a sentence. All such phrases are called **noun phrases** (abbreviated "NP"). An NP need not actually contain a noun; for example, the NP in sentence 5, *he*, does not. But the prototypical NP does contain a noun. This noun is called the **head** of the NP; this is a term we already saw in the last chapter.

English Noun Phrases

Exercise 6.3.2

How are the words in boldface in the following phrases similar to one another?

1. **a** friend
2. **a** new friend
3. **the** friend
4. **the** old friend
5. **that** friend
6. **that** particular friend
7. **your** friend
8. **your** crazy friend

What else can an English NP contain? One possibility is an adjective, for example, in *the tall teacher*. In English this always comes before the noun. In fact an NP can have more than one adjective: *the tall sick teacher*. We can think of an NP as having a set of positions or slots in which words can appear. There is a position for the head noun, and a position for one or more optional adjective modifiers. We can denote these possibilities as follows, where adjective and noun are abbreviated as "Adj" and "N".

Adj* N

The "*" after "Adj" means that some number of adjectives, including zero, are possible.

But note that *tall teacher* is not really a grammatical NP; it requires a word such as *the* or *a* before the adjective. If we looked at a lot of English noun phrases, we'd discover that these words always precede the adjectives in a noun phrase if there are any adjectives. We'd also discover that there are a number of other words that can appear in this same position, that is, at the beginning of an NP and before any adjectives. These words include *that*, *this*, *some*, *my*, and *your*. We'd also discover that English doesn't permit more than one of these words in a given NP. We can't say things like *the my boss* (though such combinations are possible in some other languages). This position in the NP is called the **determiner** position, and we can also use this word to refer to the words that can fill that position. So now we can denote the structure of NPs that have a common noun such as *teacher* as follows, using "Det" to mean determiner.

Det Adj* N

Semantic Roles and Syntactic Roles

In plain English, this expression says that an NP can consist of a determiner followed by some number of adjective modifiers (including none) followed by a head noun.

We are by no means finished characterizing the structure of English NPs, but it is useful to summarize what we have so far. English NPs have three positions in them that can be filled by words. Just as we can speak of the roles in an event or state, we can speak of the roles in a phrase. But since with a phrase we are concerned with linguistic form rather than meaning, these are **syntactic roles**, rather than semantic roles. As with semantic roles, we can distinguish the role from the category of things that fills the role. For example, the source and the goal in an instance of move are both filled by things belonging to the category place. For phrases, we can talk about either the categories of the words that make up the phrase or the roles that the words play in the phrase. For NPs, the roles we have seen so far are the following, where "PreMod" means "pre-modifier", a modifier that comes before the head. Note that for the determiner, we use the same word for the role and the category.

Det PreMod* Head

If we look at more NPs, we will see that this structure doesn't fit them. In sentence 9 above, the NP takes the form *the teacher behind me*. Here the phrase *behind me* functions as a modifier within the NP. Just as an adjective can modify a noun by narrowing down the category of things to those with a particular attribute, a phrase such as *behind me* can narrow down the category of things to those members of the category that are in a particular location.

In sentence 10 there is another example of a modifier following the head noun. In this sentence the NP is *the teacher who got fired*. Within the NP the phrase *who got fired* behaves like a modifier; it narrows down the category of things designated to members who got fired (that is, who are the patients of a firing event).

The NPs in 9 and 10, then, can both be denoted as follows, where "PostMod" means post-modifier, a modifier that follows the head noun.

Det Head PostMod


How Syntax and Semantics are Different and the Same

There are two important points to note about the syntactic roles in NPs. First, unlike for semantic roles, there is a sequential order to these roles. The expression above does not only mean that an NP can have a determiner, a head noun, and a post-modifier; it means that they appear in that order. Second, the roles can be filled either by single words — the determiner in *the teacher who got fired* — or by whole phrases — the post-modifier in *the teacher who got fired*. That is, the constituents of phrases can themselves be words or phrases. In fact, an NP can have another NP inside it; for example, the NP *the teacher in front of me* contains the NP *me*. Since an NP can contain an NP, we should not be surprised to see NPs that contain NPs that contain NPs, for example, *the teacher next to the door that leads to the exit*. In fact there is no obvious limit to this process.

We can also make more complex NPs in English by adding additional post-modifiers, as in *the teacher that got fired that I told you about*. It is also possible to combine pre-modifiers and post-modifiers in the same NP, as in *the tall teacher that I told you about*. If we want to capture all of these possibilities in one schema, we can write the following.

Det PreMod* Head PostMod*

This says that an NP can consist of a determiner followed by some number of pre-modifiers (including none) followed by the head noun followed by some number of post-modifiers (including none). We can also diagram this as we did with schemas for events and states. The dotted lines indicate constituents that can be absent, and the arrows indicate the order of the constituents.



Let's summarize what we've learned about English NPs. They serve as constituents in sentences, where (in one of their functions) they refer to the participants of the state or event that the sentence designates. That is, they have a particular function to perform within sentences. They can consist of as few words as one, and there is no clear limit on their maximum length. They can be described as consisting of constituents, parts that are either words or phrases in their own right and that fill particular syntactic roles in the structure of the NPs. For compositionality to work, there should be a grammatical rule for each of these roles that specifies

how the meaning of the whole phrase depends on that role. In the last chapter, we saw informally what some of these rules would look like.

Note how NP form and meaning resemble each other. On the meaning end, we have the thing referred to with a number of properties, all of them localized in a single bounded region of space. On the form end, we have a group of words occurring together within a sentence, that is, not separated by words that belong to other constituents. This is a weak example of iconicity because it involves an aspect of form (coherence) that applies to meaning as well. The tendency for the words that make up a constituent to occur together is quite strong among the languages of the world, though there are some exceptions where other principles can override this tendency. The tendency for words that behave as units on the form end to behave as units on the meaning end of language is what I will call **constituency**.

Before we leave English NPs until the next section, it is worth saying that we are far from characterizing adequately the structure of English NPs. The schema above implies that all NPs with a noun head have a determiner, but there are perfectly grammatical English NPs such as *fresh cheese* and *rocks* that do not. It is also possible for a pre-modifier to be a phrase itself, as in *the extremely tall, slightly bald teacher*. Finally, the determiner can also be a phrase, as in *my sister's teacher*. But all of this goes beyond the scope of this book.

Lingala Noun Phrases

All languages seem to resemble each other in having phrases for referring to the participants of events. Most languages have a category of common nouns that can play the role of head in these phrases; that is, these languages have NPs like English. Many, but not all, languages also have a category of adjectives that can modify nouns. And probably all languages have at least some determiners, though many languages have no words corresponding to the English words *the* and *a*.

Let's look at NPs in Lingala as an example of a language that differs considerably from English. Here are the Lingala sentences that correspond to two of the English sentences above.


11	<i>Fred</i>		<i>abiki</i>	
	Fred		he/she:recovered	
	'Fred recovered.'			

12	<i>molakisi</i>	<i>molai</i>	<i>yango</i>	<i>abiki</i>
	teacher	tall	that	he/she:recovered
	'That tall teacher recovered.'			

Languages Tend to Have the Same Sorts of Constituents But to Differ in Their Conventional Order Within Phrases

From sentence 11, we see that Lingala is like English in allowing proper nouns to refer to things. From sentence 12, we see that Lingala is like English in allowing NPs with adjectives and determiners. It turns out that Lingala has very few adjectives, using alternative expressions in many places where English uses adjectives, but this won't concern us here. It also turns out that Lingala has no determiners corresponding to English *the* or *a*, but, as in the example, it does have words corresponding to English determiners like *this* and *that*.

Note how the order of the constituents in Lingala differs from English, however. We would have to look at a lot of NPs to make sure, of course, but the order of the constituents in 12 holds throughout the language. It also turns out that, unlike in English, there are never modifiers before the head noun. So the structure of Lingala NPs, at least those with common nouns as heads, appears to be as follows. The determiner is dashed to indicate that this constituent is often absent in NPs.



So one important way in which languages can differ is in the conventional order of the constituents. Each language will tend to have a preferred order, but the order may vary from language to language.

Armed with some idea of what NPs look like, we are ready to return to events and states. In the next section, we'll be concerned with how verbs combine with NPs to designate events and states.

This page titled [6.3: Constituency and Noun Phrases](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.4: Subjects

Let's return to our Grammmies and their attempts to talk to each other about events and states. We saw that verbs go part of the way. They allow them to distinguish different kinds of states and events from one another, to say that an event belonging to a particular category such as eat has taken place or that a state belonging to a particular category such as be_hot is true. But, as we've also seen, verbs by themselves can't make it clear who the participants in the event or state are. For this the Grammmies need to be able to produce and understand sentences consisting of verbs and NPs.

Given a sentence, what should Hearers be able to figure out about it?

1. They should be able to tell what category of event or state is being described; the verb (and in some cases a predicate noun or adjective) should make this clear.
2. They should be able to tell who the relevant participants in the event or state are, at least those that the Speaker chooses to refer to. This should be derivable from the different NPs in the sentence, though I haven't said much about how they are interpreted by Hearers.
3. They should be able to tell what the semantic roles of those participants are in the event or state.

It is this third property of sentences that we'll focus on in this section and the next two sections. We'll look at some examples of how English solves this problem and how other languages may differ from English. We'll return to this issue again in Chapter 8.

Talking About Be States

Exercise 6.4.1

What does the word *is* seem to be doing in the following sentences? What does it mean?

- *This is a digital camera.*
- *He is disgusted.*
- *Holly is my sister.*

In this section we'll look only at events or states that have only one core participant, that is, happen events and the simplest be states. In English the most common pattern for be states with only one participant uses the verb *be* (which takes the different forms *am*, *is*, *are*, *was*, *were*, etc.) followed by a predicate NP or adjective. A few other verbs such as *seem*, *taste*, *smell*, *look*, *sound*, and *feel* are also possible. These verbs can be followed by adjectives or by *like* followed by an NP. Here are some examples.

1. *Clark is mad.*
2. *The woman is a teacher.*
3. *Your brother looks terrible.*
4. *That stuff tastes like soap.*

Note how *be* (*is*, *are*, *am*, *was*, etc.) functions. It has very little content; it just marks this as a be state, leaving the nature of the state to be specified by the adjective or NP that follows. (It also has the possibility of marking the *time* of the state; this is the difference between *is* and *was*.) Spanish, Hindi, and Amharic are like English in this regard; each has a verb that functions like English *be* (in fact Spanish has two such verbs). Mandarin Chinese also has such a verb, but it is used only in sentences with predicate nouns like 2, not in sentences like 3 that have predicate adjectives in English.

American Sign Language, Tzeltal, and Inuktitut have nothing like English *be* in either kind of sentence. Here is the Tzeltal sentence corresponding to sentence 2.

5	<i>jp'ijubteswanej</i>	<i>te</i>	<i>antz</i>
	teacher	the	woman
	'The woman is a teacher.'		

Notice that in the English sentences the single participant, the theme, is referred to by an NP that appears right before the verb: *Clark*, *the woman*, *your brother*, *that stuff*. This position, or syntactic role, called the **subject**, is basic to English sentences. (To

help keep syntactic roles distinct from semantic roles, I'll write them in lowercase.) The subject of an English sentence can usually be identified as the NP that directly precedes the verb. For most personal pronouns, the subject in English also has a special form that it does not take elsewhere in the sentence: *I* (rather than *me*), *he* (rather than *him*), *she* (rather than *her*), *we* (rather than *us*), *they* (rather than *them*). So we say *I am alive*, not *me am alive* (though in some dialects, *her and me are alive* is grammatical).

In summary, English subjects are distinguished from other NPs in two ways:

- by their position in the sentence
- for personal pronouns only, by their form

Other languages also have a syntactic role that we can call the subject. As in English, it is the role that refers to the themes of be sentences like those above. But as the most basic syntactic role in the sentence, the subject is also associated with one of the core participants in the other semantic schemas that were discussed in the section on semantic roles. As we'll see later, languages tend to agree with one another in what semantic role the subject refers to for these other kinds of sentences, but the agreement will not be at all perfect.

Languages Differ in Terms of How Hearers Identify the Subjects of Sentences

We've seen that for English the subject is marked by its position within the sentence and (for a small number of cases) by its form. These two possibilities apply to other languages as well. In Spanish, as in English, the subject tends to appear before the verb. But in Spanish this is just a strong tendency; the subject can also appear after the verb, a position it can never be in in English. In Tzeltal, as you can see from sentence 5, the subject normally appears *after* the "verb" (there is really no verb in the sentence but the predicate noun *jp'ijubteswanej* 'teacher' acts something like a verb). As in Spanish, though, this is just a tendency; the word order in Tzeltal is relatively free.

Spanish, like English, has a set of special forms for some of the personal pronouns when they are used as subjects; for example, *yo* 'I' is used only as a subject. Note that we can now add to our discussion of the dimensions along which personal pronouns vary that we began in the section on meaning differences between languages. Remember that we isolated a small number of dimensions, including person, number, and gender, that distinguish personal pronouns in many languages. Just as *I* and *you* are distinguished by person, and *he* and *she* are distinguished by gender, we need a further dimension to characterize the difference between *I* and *me*. This dimension is called **case**. The case of an NP, including a personal pronoun (which is just a very simple kind of NP), specifies its syntactic role in the sentence.

English Personal Pronouns Have Two Different Cases, Nominative and Objective

The case for the subject in English and most (but not all) other languages is called the **nominative case** (abbreviated **nom**). In English, nominative case is marked explicitly only for personal pronouns. *I*, *he*, and *we* are nominative personal pronouns; *me*, *him*, and *us* are not. The pronouns that are *not* normally used as subjects, *me*, *him*, *her*, *us*, and *them*, are called **objective**.

Remember from our discussion of the dimensions that distinguish personal pronouns that languages differ in terms of which dimensions are relevant for a given set of words. Just as gender is relevant for English, Spanish, Japanese and Amharic pronouns but not relevant for Mandarin Chinese, Hindi, Lingala, Tzeltal, Inuktitut, or American Sign Language pronouns, case is relevant for English, Hindi, Spanish, and Inuktitut pronouns but not for Lingala, American Sign Language, or Japanese pronouns. Lingala, ASL, and Japanese use the same forms for personal pronoun subjects and non-subjects.

Talking About Happen Events

Exercise 6.4.2

What does the verb *get* seem to mean in the following sentences?

- *Jill got sick.*
- *My shoes got wet.*
- *This lecture is getting boring.*

Like the be states discussed above, happen events have a single core participant, the patient. Most languages use the subject for this role too.

Many types of happen events have particular English verbs associated with them, for example, *fall*, *boil*, and *die*. Here are some examples.

6. *The milk boiled.*

7. *My goldfish died.*

Note that there a clear relationship between some be states and some happen events. For each type of state, there is an event that *results* in that state. There are two ways this relationship is reflected in English sentences. One possibility is for a particular verb to be used with the happen event and an adjective that is related to that verb to be used in the be state. Here are some examples of this possibility. For each, the sentence describing the event appears first, then the sentence describing the state.

8. *Rex died. Rex is dead.*

9. *The water froze. The water is frozen.*

10. *Lois woke up. Lois is awake.*

States and Changes of State May Not Be Clearly Distinguished in Some Languages

In some languages, such as Amharic and Japanese, the relationship between the state and the event is made even more explicit. The event and the state use the same verb; in fact the two may be indistinguishable from one another. Here are two Amharic examples; only the context can make it clear whether the change of state or the resulting state is intended.

11	<i>Rex</i>	<i>motwal</i>
	Rex	has:died
	'Rex has died' or 'Rex is dead'	
12	<i>Lois</i>	<i>nek'twal</i>
	Lois	has:woken:up
	'Lois has woken up' or 'Lois is awake'	

Another possibility in English (and many other languages have a similar possibility) is for an adjective or NP to be used in both the state and event sentence, with the verb *be* for the state and the verb *become* or *get* for the event. Here are two examples, again with the event first.

13. *Clark got sick. Clark is sick.*

14. *Jimmy became a journalist. Jimmy is a journalist.*

Notice how the relationship between these two sentence patterns in each case simplifies matters for the language learner. Instead of learning a separate word for each type of state and a change into that type of state, the learner is only responsible for one word (an adjective such as *sick* or a noun such as *journalist*). In order to produce both kinds of sentences, the learner only has to learn the general rules that relate the two kinds of sentence patterns. Here is an informal way to state those rules.

be adjective (state) ↔ *get* adjective (change of state)

be a noun (state) ↔ *become* a noun (change of state)

This is yet another example of productivity, the property of language that permits speakers to make new expressions using a general rule. In this case, the rule is based on a generalization that is made in English and many other languages: for every type of state, there is a corresponding change of state that results in that state.

In this section we've seen some of the ways that languages describe events or states with only one core participant. The focus has been on the syntactic role used for that participant, the subject. There will be lots more to learn about subjects, but that will have to wait until we've looked at how other participants are handled within sentences. That's the topic of the next two sections.

This page titled [6.4: Subjects](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.5: Direct Objects

Composition and Sentences

Exercise 6.5.1

Why would it be easier for a young child unfamiliar with the details of English grammar to understand a sentence like *the monkey ate the orange* than a sentence like *the whale ate the shark*?

Now what happens when the Grammmies need to refer to events or states with two core participants? Let's start with instances of the *do_to* schema because the semantic roles are very clear-cut and because children seem to learn about these sentences relatively early.

Say an instance of eating has taken place, the agent (the eater) is Clyde, and the patient (the eaten thing) is a mango. Now one Grammie wants to tell another about what happened. They have a verb for eating — let's call it *eat* — and nouns for the two participants — let's say *Clyde* and *mango*. The speaker could just string the words together: *eat!, mango!, Clyde!*. Assuming these three words are in the hearer's lexicon, the hearer could figure out that an instance of eating is being described and that the participants in that eating are Clyde and a mango.

But the two participants in an eating play very different roles. One is an agent, the other a patient. Understanding what the speaker is saying involves also figuring out "who did what to whom", that is, which participant fills which semantic role. In our eating example, this would not be difficult. The hearer could use knowledge of what eating is to *infer* that Clyde is the agent (eater) and that the mango is the patient(eaten thing). Eaters must be animate; eaten things are often inanimate and are almost always members of the general category of food, which includes the more specific category mango.

Solving the Problem of "who did what to whom"

But things would not be as simple for an event such as a hugging. In a hugging, both participants are normally animate. So if a Grammie speaker referred to a hugging by stringing the verb and the two nouns together in a random order — *hug!, Clyde!, Lois!* — the hearer would not know whether Clark hugged Lois or Lois hugged Clark (unless this was somehow clear from the context). What's needed is a convention for how the semantic roles of a two-participant event correspond to the NPs in a sentence describing the event.

But this issue should be familiar from the last chapter. It's just what we expect from compositionality: language-specific grammatical conventions specify how the meaning of a phrase is derived from the meanings of the words or phrases that make up the phrase. The problem we face with sentences describing two-participant events is exactly the same sort of problem that must be solved for English noun + noun phrases of the type we learned about in this section. Given the unfamiliar phrase *bag box*, for example, a hearer has to be able to figure out whether this refers to a box that has something to do with bags or to a bag that has something to do with boxes. An English-speaking Hearer knows what to do because of a grammatical convention of English: it is the second noun in these phrases that is the head, that is, the noun that specifies the category for the phrase's meaning. (More precisely, the meaning of the phrase is a subcategory of the meaning of the second noun.)

Note that the solution for noun + noun phrases involves the notion of **syntactic roles**, particular positions within the phrase that play particular roles in the meaning of the phrase. The syntactic roles for noun + noun phrases (and adjective + noun phrases) are modifier and head.

For sentences the problem is the same. A compositional solution to the problem is one that says that the meaning of a sentence is some combination (specified in a grammatical convention) of the meanings of the verb and the other constituents of the sentence. Since the meaning of a sentence, at least as laid out in the theory of semantic schemas earlier in this chapter, involves semantic roles, we need a convention that relates particular sentence syntactic roles to particular semantic roles. We already have one sentence syntactic role for the NP in a one-participant sentence, the subject. We saw that this role is marked by its position or its form or both its position and its form in different languages. Most, if not all, modern languages use this same role within sentences that refer to multiple participants.

Transitive Sentences are for Referring to More than One Core Participant

We need one more syntactic role, then, for the other participant in sentences describing *do_to* events. The syntactic role that does this is the **direct object**. Like the subject, the direct object is thought to be a universal role, something that all languages have, though the way in which the direct object can be identified differs somewhat from language to language. A sentence with both a subject and a direct object is called a **transitive sentence**. Transitive sentences contrast with **intransitive sentences**, which do not have direct objects. The grammatical case for the direct object in most (but not all) languages is called the **accusative case** (abbreviated *acc*).

Direct Objects in English

Let's start with English. The English sentence that designates a hugging event in which Clyde is the agent (the hugger, that is) and Lois is the patient (the "huggee", that is) is of course the following.

1. *Clark hugged Lois.*

We know from the position of the NP *Clark* in this sentence that it is the subject. The other NP in the sentence, *Lois*, is the direct object. The normal position for the direct object in an English sentence is directly after the verb. (In the next section, we'll see one situation where the direct object is separated from the verb by another constituent.) Thus, like the subject, the direct object in English is usually identifiable by its position. Unlike the subject, the English direct object doesn't have a set of special forms for the personal pronouns; forms such as *me* and *him* are also used in other contexts. That is, English doesn't really have accusative pronouns. However, forms such as *me* and *him* are clearly *not* the subject. So consider the next sentence, which could describe the same event as sentence 1 does.


2. *He hugged her.*

In this sentence we know that the agent is the man not only because of the position of *he* before the verb but because the form of the pronoun is *he* and not *him*. By the same token, *her* is clearly the direct object, and not the subject, because of its position, but also because it is in the non-subject, objective form *her*.

English Word Order is Relatively Rigid

To summarize, English has two syntactic roles that Speakers use to refer to the core participants of a two-participant event. These roles, subject and direct object, are identifiable mainly on the basis of their position with respect to the verb. In sentences describing *do_to* events, that is, events with a clear-cut agent and patient, the subject refers to the agent and the direct object refers to the patient.

The relationship between the syntactic roles of a sentence and the semantic roles of the event or state that the sentence describes is an example of a **mapping**. The notion of mapping is a concept from mathematics that is very important in cognitive science. For a mapping there are two sets of elements, and each of the elements in one set *corresponds* to (or "maps onto") an element in the other. In our case each of the syntactic roles in the transitive sentence maps onto a semantic role in the situation. The figure below illustrates this. Because we are concerned now with the relationships between form (syntax) and meaning (semantics), there are now two large boxes. The green arrows denote these relationships. The one connecting the two large boxes denotes the fact that the sentence describes the event. The ones connecting the syntactic roles in the sentence with the semantic roles in the event denote the correspondences in the syntax-semantics mapping. The "<" symbols between the syntactic roles represent their usual order.



Direct Objects in Japanese and Spanish


Now let's look at subjects and direct objects in Japanese, which are a bit more complicated. One possible Japanese translation of sentence 1 above would be the following sentence.

3	Clark	wa	Lois	o	dakishimeta
	Clark	topic	Lois	acc	hugged
	'Clark hugged Lois.'				

Languages with Flexible Word Order May Have Case Markers to Help Hearers Identify the Syntactic Roles of Constituents

Notice first that the order of the sentence constituents is different in Japanese. The prototypical position for the verb in a Japanese sentence is at the end, so both the subject and the direct object precede it. The subject usually precedes the direct object, though this order is not so rigid. Second, notice that the Japanese sentence has two words that don't correspond to anything in the English sentence, *wa* and *o*. These words, sometimes called particles, specify the role of the NP they follow in the sentence. The particle *o* specifies that the previous NP is in the accusative case, that is, that it is the direct object of the sentence (though the detailed function of *o* is somewhat more complicated than this). A word like this is called a **case marker**. The particle *wa* has a different sort of function; it specifies that the NP it follows is the "topic" of the sentence, roughly what the sentence is "about". Notice that in this case *wa* follows the subject of the sentence. It turns out that Japanese also has a nominative case marker, *ga*, but when the subject is the "topic", the nominative case marker is replaced by the topic marker *wa*.

The figure below illustrates the syntax-semantics mapping for do_to sentences in Japanese. The dashed line separates the subject and direct object from the verb, which they normally precede. The "<" symbol between the subject and direct object is fuzzy because this ordering is only a tendency in Japanese.



A Japanese direct object can also be a "topic", though this is less common. In that case the accusative case marker *o* is replaced by the topic marker *wa*. Here is what our hugging example would look like with the direct object as topic; notice that the normal place for the topic is the first position in the sentence.

4	Lois	wa	Clark	ga	dakishimeta
	Lois	topic	Clark	nom	hugged
	'Clark hugged Lois.'				

You can think of the difference between sentences 3 and 4 as corresponding roughly to the following (somewhat awkward) English translations: 'as for Clark, he hugged Lois' (3) and 'as for Lois, Clark hugged her' (4).

What is important about all this for our purposes is that Japanese has special words, case markers, to indicate which NP is the subject and which the direct object. That is, even when the order of the subject and direct object deviates from the default order (subject first), a hearer can figure out which NP is which.


What is also important about Japanese is how it is *like* English. For the do_to schema the syntax-semantics mapping is the same: the subject maps onto the agent, and the direct object maps onto the patient.

Let's look at one more language before we consider other types of states and events. Spanish is similar to English in the default order of the constituents — subject, verb, direct object — but as already mentioned, there is much more freedom in Spanish to

deviate from the default. Spanish is also like Japanese in having an accusative case marker, though this is normally limited to direct objects referring to humans. Here is one Spanish translation of sentence 1.

5	<i>Clark</i>	<i>abrazó</i>	<i>a</i>	<i>Lois</i>
	Clark	hugged	acc	Lois
	'Clark hugged Lois.'			

The accusative case marker *a* appears before the NP *Lois*, marking it as the direct object of this Spanish sentence. The figure below illustrates the syntax-semantics mapping for Spanish *do_to* sentences. The accusative case marker is in a fuzzy box because it is only used in certain situations.



But depending on what is being emphasized and what is being treated as surprising information by the speaker, the constituents can be put in other, less common, orders. Here is a possibility that corresponds closely in meaning to the Japanese sentence 4.

6	<i>A</i>	<i>Lois</i>	<i>la</i>	<i>abrazó</i>	<i>Clark</i>
	acc	Lois	her	hugged	Clark
	'Clark hugged Lois.'				

Even with the normal positions of subject and direct object switched, the sentence is still interpretable because of the accusative case marker preceding *Lois*. To make it even clearer, Spanish speakers usually insert a redundant object pronoun when the order is the reverse of the usual one; in this case the pronoun is *la* 'her'.

Change of State and Causing of Change of State

Exercise 6.5.2

Say you learn about a new cooking technique along with the English verb for it, *drib*, and you hear the verb illustrated in sentences like the following.

- *Josh dribbed the artichokes*
- *I'm going to drib these beans.*

Now you hear the word used in a different sort of sentence.

- *The cauliflower is dribbing.*

How do you interpret this last sentence in terms of what you've already learned about the verb and the cooking technique?

For the relationship between states and changes of state, we saw that a language can simplify the task for the learner by making a generalization that recognizes what states and changes of state share. Being crazy and becoming crazy are similar in one way, and the learner only has to learn one new word for each new state and change of state. English makes a similar sort of generalization concerning changes of state and acts of an agent (that is, instances of *do_to*) that *result* in the changes of state. Consider the following two sentences concerned with a breaking event.


7. *The vase broke.*
8. *Jimmy broke the vase.*

The first sentence is intransitive; it mentions only the patient. The second sentence is transitive; it mentions both the patient and the agent. Notice that the patient appears in different syntactic roles in the intransitive and transitive sentences: the subject of the first sentence, the direct object of the second. This is an example of a generalization that English could, but does not make (though other languages do). That is, English fails to treat the patient in the same way in the two kinds of sentences. On the other hand, English does use the same verb for the two cases, in effect recognizing that the same kind of event is involved.

Syntax-semantics Mappings are Stored in the Lexicon with Verbs

So what does a speaker of English need to remember about the verb *break* in the lexicon? One possibility is that there are two separate verbs, one for intransitive sentences like sentence 7 and one for transitive sentences like 8. But this would miss the obvious generalization that is made in English; if these are two completely different verbs, it is very surprising that they are pronounced and written exactly the same. The other possibility is a single verb with a single meaning (whatever it is that breaking involves) and two different patterns for how the syntactic and semantic roles map onto one another.

Here is one way of writing the two syntax-semantics mappings for *break*. Each of the smaller boxes represents a different mapping, and only one of these appears in a given sentence with *break* as the verb. The upper box shows the mapping used in sentence 7; the lower box shows the mapping used in sentence 8.



But the English generalization goes beyond this. There is a whole set of verbs with this same property: *boil*, *freeze*, *cook*, *fry*, *bake*, *steam*, *soften*, *thicken*, *shatter*, *split*, *rip*, *tear*. For all of these verbs the same two sets of mappings illustrated in the figure above apply. So English speakers probably learn a general rule that allows them to produce and understand both transitive and intransitive sentences with verbs similar to these even if they have never heard the sentences before, as in the box at the beginning of this subsection.

This English example illustrates a more general point. While we've seen only similarities between languages so far in terms of the syntax-semantics mappings, these are conventions that do differ from one language to another. There will be a number of examples of these differences in the next section. But even within languages, particular verbs may have unpredictable mappings associated with them. For this reason, it is usually assumed that knowledge about how syntax and semantics map onto each other belongs in the lexicon, where it is associated with individual verbs or with clusters of similar verbs (such as *break* and *boil*). We'll look at some more syntax-semantics mappings for verbs in English and other languages in the next section.

This page titled [6.5: Direct Objects](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.6: Adjuncts

Prepositional and Postpositional Phrases

Exercise 6.6.1

How are the instrument and the beneficiary marked in the following sentences?

- Lois did a favor for Clark.
- Clark figured out the math for Perry with his new computer.
- Jimmy managed to unlock his car with a piece of wire.

As we saw in the section on semantic schemas, an instance of the *be* or *happen* schemas can have more than the single core participant, and an instance of the *do_to* schema can have more than the two core participants. There are additional, peripheral participants that are not central to the particular category of state or event but may be worth noticing or mentioning nonetheless. The semantic roles for these participants include instrument, beneficiary, sufferer, location, time, and manner.

For a given event, a relatively wide range of these roles is possible. For example, a transfer event can have any of the above roles. For this reason, when the Grammies realized they would want to talk about these participants, they decided that they would need more different ways of marking the roles than they had for the subject and direct object. They created a new category of words whose function was to specify the semantic role for these peripheral participants. In English we use such a set of words; we call them **prepositions** because they appear before (*pre-*) the NP referring to the participant. Here are some examples.

1. *The soup was cooking **on** the stove.*
2. *Lois tickled Clark **with** a feather.*
3. *Clark cooked some soup **for** Lois.*
4. *Lois sang **until** dawn.*

In sentence 1 the preposition *on* begins the phrase referring to the location of the cooking. In sentence 2 the preposition *with* begins the phrase referring to the instrument of the tickling. In sentence 3 the preposition *for* begins the phrase referring to the beneficiary of the cooking. And in sentence 4 the preposition *until* begins the phrase referring to the time of the singing. Semantically, we can see that the function of a preposition is to specify some relation between a state or event and a thing (the thing that fills the peripheral semantic role). A preposition may also function to specify the relation between two things, as it does for the relation between a cat and a table in the following sentence.

5. *The cat is under the table.*

Syntactically a preposition combines with an NP to form a larger phrase, called a **prepositional phrase**. In the last five examples, the prepositional phrases are *on the stove*, *with a feather*, *for Lois*, *until dawn*, and *under the table*.

Some Languages Have Prepositions; Others Have Postpositions. (A Few Have Both.)

Many other languages, including Spanish, Lingala, and American Sign Language, also have prepositions. But in other languages there are words with a similar function that *follow* the NPs that they are relating to the rest of the sentence. Japanese and Hindi are such languages. These words are called **postpositions**, and when we want to group prepositions and postpositions together, we call them **adpositions**. Here is the Japanese sentence corresponding to sentence 4. Note that the postposition *made* corresponds to the English preposition *until*.

6	<i>Lois</i>	<i>wa</i>	<i>yoake</i>	<i>made</i>	<i>utatta</i>
	Lois	topic	dawn	until	sang
'Lois sang until dawn.'					

Prepositional or postpositional phrases within a sentence are called **adjuncts** to distinguish them from subjects and direct objects. We'll meet some other kinds of adjuncts later in this section. In the rest of this section, we'll look at how adjuncts are used to refer to particular peripheral roles and how these peripheral roles can also be realized in other ways. In the process we'll discover some

interesting ways in which languages can differ from one another. One kind of difference involves a distinction that is made in one language but not in another. In other words, two different sentences in one language correspond to one ambiguous sentence in the other language. Another kind of difference involves alternate ways of conveying the same information. One language may express in adjuncts what another language expresses within the verb itself.

Talking About Move Events

Exercise 6.6.2

What words refer to the path in the following sentences?

- Lois took the dog out.
- Clark pulled the lid of the bottle off with pliers.
- The log floated down the river.

Remember that move events have a source, a goal, and path, any of which might be worthy of attention and mention. English typically uses the preposition *from* for the source and the preposition *to* for the goal. There are a number of other possibilities for the goal, however. Here is an example with *under*.

7. *The ball rolled under the table.*

English Sometimes Fails to Distinguish a Location from a Goal

But this sentence is in fact ambiguous, and it's an ambiguity we can express in terms of semantic roles. We could use the same sentence if the rolling took place under the table, that is, if the area under the table was the location rather than the goal of the rolling. (The location interpretation probably seems odd because it is hard to visualize what caused the rolling in this case.)

Often when a language exhibits ambiguity in a sentence, that sentence corresponds to different sentences in some other languages. In this case, for example, Japanese would convey sentence 7 differently, depending on whether the area under the table is the goal. Here are the two Japanese sentences corresponding to sentence 7, the first with the goal interpretation, the second with the location interpretation.

8	<i>booru</i>	<i>wa</i>	<i>teeburu</i>	<i>no</i>	<i>shita</i>	<i>e</i>	<i>korogatta</i>
	ball	topic	table	of	bottom	to	rolled
	'The ball rolled under the table (goal).'						
9	<i>booru</i>	<i>wa</i>	<i>teeburu</i>	<i>no</i>	<i>shita</i>	<i>de</i>	<i>korogatta</i>
	ball	topic	table	of	bottom	at	rolled
	'The ball rolled under the table (location).'						

The path in English is often expressed by one of a small set of words that behave somewhat like prepositions, except that they occur freely without a following noun phrase. These include *up*, *down*, *out*, and *off*, as well as some words that can also function as prepositions such as *in* and *around*. I will refer to these words as **directional adverbs**. Here are some examples.

10. *The boll rolled down the stairs.*
11. *The ball rolled down.*
12. *The balloon floated up the chimney.*
13. *Clark strutted around the room.*

In languages such as Spanish, on the other hand, it is more common to express the path as a part of the verb. That is, Spanish verbs for move events often include a specification of the direction of the movement. Here are some examples.

14	<i>la</i>	<i>pelota</i>	<i>bajó</i>
	the	ball	went:down
	'The ball went down.'		

15	<i>el</i>	<i>globo</i>	<i>subió</i>
	the	balloon	went:up
	'The balloon went up.'		

Languages Differ in Terms of Which Elements of Sentences Represent Which Information

The verbs in these sentences include the notions of falling or rising, so we can say that Spanish tends to **lexicalize** aspects of the path in verbs for move. Of course English has verbs such as *descend* and *rise* too, but these are not nearly as common as alternatives like *go up* and *go down* and are not normally learned by young children. The point is that the *basic* verbs used for motion in Spanish tend to express path, whereas the basic verbs for motion in English do not.

Instead English has a tendency to lexicalize manner in its verbs for move events. Examples in the above sentences include *roll* and *float*. Note that neither of these verbs specifies the path; the rolling or floating could be in any direction.

So what does Spanish do with manner in sentences corresponding to 10-13 above? It appears that Spanish speakers are less likely to refer to manner in sentences like this than English speakers are. When they do refer to manner, it tends to appear in an adjunct of a kind we have not seen yet, one that is a form of a separate verb expressing manner. Here is a Spanish translation of sentence 12. The adjunct *flotando* expresses the manner of the movement.

16	<i>el</i>	<i>globo</i>	<i>bajó</i>	<i>por</i>	<i>la</i>	<i>chimenea</i>	<i>flotando</i>
	the	balloon	went:down	by	the	chimney	floating
	'The balloon floated down the chimney.'						

For more on how languages can be grouped on the basis of what aspects of situations are lexicalized, see the work of Leonard Talmy, the linguist who pioneered work of this type.

Talking About Transfer Events, Indirect Objects

Exercise 6.6.3

What two ways does English have to refer to the recipient?

- *Clark sold Mary his house.*
- *Clark sold his house to Mary.*

Recall that transfer and information_transfer events are like move events in having a source and a goal, that is, the receiver. For the subcategory of events where the source of the transfer is also the agent, for example, in giving and telling, English has two common patterns. One uses the preposition *to* to mark the recipient. Here are two examples.

17. *Clark lent his laptop to Lois.*

18. *Lois told her story to her officemates.*

Verbs for Transfer Events Can Have Three Associated Core Syntactic Roles

The other possibility is to refer to the recipient in an NP that has no preposition indicating its role and appears right after the verb. This syntactic role is called the **indirect object**. It represents a third core syntactic role, along with the subject and direct object, in

English and many other languages. Notice that when there is an indirect object in an English sentence, it appears before the direct object. Here are sentences 17 and 18 reformulated using indirect objects.

19. *Clark lent Lois his laptop.*

20. *Lois told her officemates her story.*

Since sentences 17 and 19 and sentences 18 and 20 have different forms, however, we would expect their meanings to differ in some way. But the difference is very subtle and not well understood. And the choice of one pattern over the other depends on other factors not directly related to the meaning of the sentence, such as the length of the NP referring to the recipient. That is, sentence 19 would be much less acceptable if *Lois* were replaced by *the woman he met on the train to work last Friday*. In this case English speakers would prefer the pattern with the NP referring to recipient after the direct object and preceded by *to*.

Note that English makes no distinction between the form of the indirect and direct objects; that is, the same set of personal pronouns, the objective forms, is used for both. Spanish does make such a distinction, however, though only for the third person, where accusative case is distinguished from **dative case**, the form used for the indirect object. Thus there are two Spanish words corresponding to *him*: *lo*, the accusative form, and *le*, the dative form.

In English, the indirect object is also used for beneficiaries with some verbs. Here is an example.

20. *Lois knit Billy a sweater.*

Note that for beneficiaries the alternate form with a prepositional phrase adjunct instead of the indirect object uses the preposition *for* rather than *to*.

21. *Lois knit a sweater for Billy.*

In Chapter 8 we'll see how Lingala provides an elegant way to create new forms of verbs that take indirect objects for a wider variety of roles than is possible in English.

Talking About Experience

Exercise 6.6.4

In what ways is an experiencer like an agent? In what ways is an experiencer like a patient?

Modern languages tend to be fairly consistent in the syntax-semantics mappings for the *do_to* schema, with the agent represented by the subject and the patient represented by the direct object. What happens with other schemas with two or more core participants?

Let's try to get some insight into this by continuing our fictitious history of how language started. The Grammys first started using transitive sentences to refer to *do_to* events, events that are easy to observe and whose participants' roles are easy to understand. When they first felt the need to refer to human mental experience, it was not so clear how to proceed.

As we saw in the section on situation schemas, these states or events have two core participants, just as *do_to* does, but the roles they play are quite different. Rather than creating completely new syntactic roles for the experiencer and the experience theme, the Grammys realized that they could make do with the two core syntactic roles they already had, subject and direct object. Since the verb of a sentence could make it clear that the sentence was about an instance of experience rather than an instance of *do_to*, all that was needed was conventions for how the two semantic roles were to be referred to by the NPs in the two syntactic roles.

Languages Make Different Generalizations by Mapping Different Sets of Semantic Roles onto the Subject and the Direct Object

But the sensible thing would be to use the subject and direct object in a way that *resembled* their use in the familiar *do_to* sentences. That way each syntactic role would tend to have a consistent semantic interpretation. There seem to be two ways to go, however. In one way, the experiencer is like an agent; it is typically animate and intimately involved with the event or state. That is, when you see, believe, or hate something, you seem to have much more to do with what is going on than the thing that you see, believe, or hate does. On the other hand, the experiencer is like a patient in the sense that is affected by the state or event in a relatively passive way. Seeing or believing are processes that happen to you, not processes that you are in control of.

Since both of these associations make some sense, it is not surprising that the Grammys failed to agree on a single way to map the experiencer roles onto their syntactic roles. Some of them decided to use subjects for experiencers (as well as agents), while others decided to use other syntactic roles for experiencers. Still others treated some kinds of experience one way and other kinds another way. Modern languages also exhibit this sort of disagreement. Let's look at a few examples.

For the most part, English went with the option by which experiencer is like agent and theme is like patient. That is, it is the subject that typically refers to the experiencer and the direct object to the theme. Here are two examples.

22. *Lois sees Clark.*

23. *Lois likes Clark.*

Spanish agrees with English for sensory experiences like seeing and hearing: the Spanish subject refers to the experiencer, and the Spanish direct object refers to the theme. But look at a natural translation of sentence 23 into Spanish.

24	A	<i>Lois</i>	<i>le</i>	<i>gusta</i>	<i>Clark</i>
	To	Lois	her(dat)	is:pleasing	Clark
	'Lois likes Clark.'				

In this sentence Lois is referred to both by the NP *Lois* and the dative (indirect object) pronoun *le*. Though *Lois* is the first NP in the sentence, it is not the subject. In sentences like this, Spanish treats the experiencer like it treats a recipient. The subject of the sentence is *Clark*, referring to the theme.

Japanese differs even more from English. For sensory experiences, as in sentence 22 above, Japanese tends to put the experiencer in an NP marked with the postposition *ni*, roughly 'to'. A natural Japanese translation for sentence 22 is the following.

25	<i>Lois</i>	<i>ni</i>	<i>Clark</i>	<i>ga</i>	<i>mieru</i>
	Lois	to	Clark	nom	is:visible
	'Lois sees Clark.'				

The theme, Clark, is referred to by the subject, and the experiencer, Lois, appears in a phrase marked just as Japanese marks the recipient of a transfer. Japanese sentences for liking and hating states are more complicated, so I won't include them here.

Syntax-semantics Mapping Revisited

In this section we've seen that for some categories of situations, the syntax-semantics mapping is reasonably predictable. That is, for events with clear agents and patients, there is a strong tendency within and across languages to map the subject onto the agent and the direct object onto the patient. For some of the peripheral semantic roles, it is also possible to make generalizations about their syntactic realization that apply to sentences with many different verbs. For example, the English prepositions that are used for source, goal, location, and time can be used in almost any sentence where it is appropriate to mention these semantic roles. And the instrument can almost always be conveyed in English with the preposition *with*.

The Same Semantic Role Can be Realized as a Prepositional Phrase in One Sentence and the Subject in Another

However, even for these semantic roles, things are not always so simple. Some verbs even permit them to appear in one of the core syntactic roles. Here are some English examples.

26. *The stone broke the window.*

27. *The 1990s saw the end of the Cold War.*

28. *Perry splattered the floor with paint.*

In sentence 26, the subject, *the stone*, refers to the instrument apparently used by some unspecified agent. In sentence 27, the subject, *the 1990s*, refers to the time of the event that is realized as the direct object, *the end of the Cold War*. In sentence 28, the

direct object, *the floor*, refers to the goal, whereas the object of the preposition *with* refers not to the instrument of the splattering but rather the patient.

We have also seen that the syntax-semantics mappings are less predictable for some classes of situations. Within languages, experience states and events may be treated differently for different verbs, for example. In general, as noted in the last section, it's safest to assume that each verb is associated in the lexicon with its own syntax-semantics mapping(s).

Let's look at some more examples of these mappings. First, given the realization of the instrument in a sentence like 26, we see that there is a third mapping for verbs like *break*, in addition to the two discussed at the end of the last section. Here is a more complete set of mappings for *break* than appeared there. The optional instrument phrase is included in the last mapping to show how it differs from the realization of the instrument in the second mapping.

<i>break</i>	
subject	PATIENT
subject direct object	INSTRUMENT PATIENT
subject direct object (object of <i>with</i>)	AGENT PATIENT INSTRUMENT

Here are the two possible mappings for *give*, one of the verbs that can take indirect objects referring to recipients.

<i>give</i>	
subject direct object (object of <i>to</i>)	AGENT PATIENT RECIPIENT
subject indirect object direct object	AGENT RECIPIENT PATIENT

Finally, here are syntax-semantics mappings for English and Spanish verbs of liking.

<i>like</i>	
subject direct object	EXPERIENCER THEME
<i>gustar</i>	
subject indirect object	THEME EXPERIENCER

This page titled [6.6: Adjuncts](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.7: Sentence Functions

Your page has been created!

Remove this content and add your own.

Edit page

Click the Edit page button in your user bar. You will see a suggested structure for your content. Add your content and hit Save.

Tips:

[Drag and drop](#)

Drag one or more image files from your computer and drop them onto your browser window to add them to your page.

[Classifications](#)

Tags are used to link pages to one another along common themes. Tags are also used as markers for the dynamic organization of content in the CXone Expert framework.

[Working with templates](#)

CXone Expert [templates](#) help guide and organize your documentation, making it flow easier and more uniformly. Edit existing templates or create your own.

[Visit for all help topics.](#)

This page titled [6.7: Sentence Functions](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.8: Problems

Schemas and Semantic Roles

Identify the semantic role of each noun phrase **in bold** in the following sentences. Choose from the following set: AGENT, PATIENT, EXPERIENCER, THEME, RECIPIENT, SOURCE, GOAL, PATH, INSTRUMENT, LOCATION, TIME, CAUSER, BENEFICIARY. If you're not sure, you can say that the role is a blend of two roles.

Example 6.8.1:

The calendar fell off **the wall**.

the calendar: PATIENT, *the wall*: SOURCE.

1. **He** got fired. PAT
2. **The rock** was rolling down **the mountain**. PAT, PTH
3. **You** put **chile** in **this dish**? **I** don't taste **it**. AGT, PAT, GOL, EXP, THM
4. **A cache of weapons** was seized by **the FBI** on **Thursday**. PAT, AGT, TIM
5. **She** always amazes **me**. THM, EXP
6. **Bill** borrowed **a book** from **Sarah**, and **he** returned **it** to **her** right away. AGT/REC, PAT, SRC, AGT/SRC, PAT, REC
7. **She** said **only three words**. AGT, PAT
8. **She** got **her hair** cut at **Mabel's**. CSR, PAT, LOC
9. **Yikes!** **I** feel **something** crawling up my back. EXP, THM

Syntax-semantics Mappings

For each sentence or group of sentences for a given verb (shown in bold), write the mapping(s) of syntactic and semantic roles.

Example 6.8.2:

a. Clark **lent** Lois a book.

lend

a. subject: AGENT/SOURCE, direct object: PATIENT, indirect object: RECIPIENT

1. *disgust*
subject: THEME, direct object: EXPERIENCER
2. *admire*
subject: EXPERIENCER, direct object: THEME
3. *email*
subject: AGENT, SOURCE; direct object: RECIPIENT
- 4.
5. Jimmy **burned** the wood.

burn

a. subject: AGENT, direct object: PATIENT

This page titled [6.8: Problems](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

7: Grammatical Categories

[7.0: Introduction](#)

[7.1: Morphemes](#)

[7.2: Grammatical Categories and NPs](#)

[7.3: Grammatical Categories and Verbs](#)

[7.4: Problems](#)

This page titled [7: Grammatical Categories](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.0: Introduction

The things and the situations that language is about and the utterance contexts in which language occurs can be seen in terms of a small set of very general, abstract properties. For example, things are either objects or masses or they are abstractions which can be seen as more like objects or more like masses. Situations have a time when they were, are, or will be true. And utterance contexts involve various kinds of possible social relationships between the hearer and the speaker. Each language directly represents some of these abstract properties using words such as *the*, *some*, and *was* or meaningful parts of words, called morphemes, such as *-s*, *-ed*, and *-ing*. Each such form represents a grammatical category, a way of grouping things or situations or contexts on the basis of one of the abstract properties. These forms behave in a different way from the nouns, verbs, and adjectives we have looked at in previous chapters.

First, languages differ a lot in which grammatical categories they make use of. A natural translation of the English sentence *I saw a movie* into Japanese, *eiga o mimashita*, has no part that corresponds to either the English word *I* or the English word *a*. But the Japanese sentence does have a morpheme that tells the role of the movie in the seeing and a morpheme that conveys something about the social distance between the Speaker and Hearer. Neither of these is present in the English sentence. Second, grammatical categories are in a sense forced on the speakers of a language. In English, we need the *-s* on *pencils* in the phrase *three pencils* whether we like it or not; *three pencil* is ungrammatical even though it is perfectly understandable. Third, the linguistic forms that convey grammatical categories tend to look different than nouns, verbs, and adjectives. For example, some, such as the *-s* in *pencils*, can not even be pronounced in isolation. In this chapter we'll look at how grammatical categories "slice up" the world by dividing it into a set of very abstract semantic categories; what form they take in language; and how they vary across languages. In the process we'll be looking inside words again, this time not at phonological units but at the meaningful units that make up many words, for example, the *pencil* and *-s* in *pencils*.

This page titled [7.0: Introduction](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.1: Morphemes

The Lexical-grammatical Continuum

Let's compare some of the kinds of words we have encountered so far. A common noun like *apple* refers to a relatively specific category of things, with a narrow range of values on a number of dimensions (shape, size, consistency, taste, reflectivity, etc.). This category is one of thousands of categories of things that people know about and that all languages have words for (though not all languages agree on which specific categories of course). A verb like *run* refers to a relative specific category of move event with a characteristic set of semantic roles, a specified manner (in the most basic sense of the word, a characteristic pattern of movement of the legs resulting in a relatively rapid movement). This category is one of at least hundreds of categories of events and states that people know about and that all languages have words for.

Compare these words with a word like *he*. Very little is specified with this word — a gender, a person, a number, and a case (nominative) — and it can potentially refer to a very wide range of things, nearly half of the people in the world in fact, as well as some animals under the right circumstances. The word *he* belongs not to a set of hundreds or thousands of words but to a set of about 12 words, the personal pronouns of English. Likewise, consider the preposition *on*, which designates a spatial relation between two things (*the book is on the table*) or between an event and a thing (*Lois put the book on the table*). This relation is one of a very small set of spatial relations that English has words for; other words in this set include *in*, *at*, *over*, *under*, and *beside*.

The point is that words differ in terms of how abstract and general their meanings are. On the relatively specific, concrete end of this continuum, we have nouns, verbs, and adjectives, each designating a quite particular category of thing, situation, or attribute. There are not only many such words; the sets of these words are also relatively open-ended. Speakers can freely add new words to the "list". Within the last 100 years, English speakers have created nouns such as *nerd* and *grunge*; verbs such as *diss* and *zap*; and adjectives such as *icky* and *PC*.

On the relatively general, abstract end of the continuum, we have pronouns such as *he* and prepositions such as *on*. Each of these words has a very general meaning and so designates a very wide variety of things or situations. The words belong to small sets such as personal pronouns and prepositions which are not often extended in the history of the language. Imagine the strangeness of creating a new preposition for a particular spatial relation, say, *agrove* for being partly inside and partly outside of something (*the pencil is agrove the cup*). But there are even more abstract words. Take *the* or *a*, for example. Such words seem almost to have no meaning at all. (Of course they do have a meaning, but it is very abstract and difficult to describe.)

I will call this continuum that extends from words like *apple* on one end to words like *the* on the other the **lexical-grammatical continuum**. In this chapter and the next chapter, we'll be focusing on the grammatical end of the continuum. It is probably in this area of abstract meanings that languages differ most from one another. One language may pick up on an aspect of meaning that is ignored in the grammar of another language. In fact if there is place where the language you know influences the way you think or perceive the world, it is here.

Grammatical Morphemes

Exercise 7.1.1

What is the difference in meaning between *apple* and *apples*? What is the difference in form? What does this tell you about these two words?

Words Can Have an Internal Structure Much Like the Syntax of Phrases

Because words on the grammatical end of the lexical-grammatical continuum have such vague meanings, they do not contribute as much to the overall meaning of a sentence as nouns, verbs, and adjectives do. For this reason, they tend to get less emphasis from speakers. They tend to be short and to receive no stress. Over time, this may lead them to lose the character of words altogether and become "attached" in some sense to nouns, verbs, or adjectives. Here are a few English examples.

1. *pencils*
2. *walked*
3. *unnecessarily*

4. reconnecting

In each case there is a basic word on the lexical end of the continuum: *pencil*, *walk*, *necessary*, *connect*. In addition there are one or more elements attached to the beginnings or ends of these words to make longer, more complex words. These elements include *-s*, *-ed*, *un-*, *-ly*, *re-*, and *-ing*, where the hyphens indicate which side these elements attach to. In each case, these elements also contribute some meaning to the larger word. For example, in *reconnecting*, the *re-* indicates a repetition of the connecting. But in this section I'll have little more to say about the meanings of such elements. The focus for the moment is on their form.

In what sense are these elements not words? It turns out that "wordhood" is a matter of degree, like so many other concepts in linguistics. And like some other concepts, it is multi-dimensional; there are different criteria for being a word. First, a word should be pronounceable. In the most extreme cases, for example, *-s* in *pencils* (pronounced /z/) and *-ed* in *walked* (pronounced /t/), these elements that we're considering are not even legitimate English syllables. So on grounds of pronounceability, these make very bad words.

Second, a word should have a coherent form that does not depend much on its neighbors. In this sense the element that we write "ed" is not a good word because it is pronounced /t/ in words such as *walked*, /d/ in words such as *lived*, and /@d/ in words such as *needed*.

Third, a word should be relatively independent of the other words around it; that is, it should be possible to separate it from them with other words. In this sense *-ed* is not a good word since there is normally nothing that can separate it from the verb it follows; we can't say things like *walk away ed* or *walk alone ed*. By the same token, *re-* does not make a good word since it cannot be separated from the verb it precedes; it is impossible to say things like *re carefully connect* or *re don't connect*.

Since forms like *-s*, *-ed*, and *re-* are not words, we need another name for them. Each of these forms has a meaning (though as we'll see later in this chapter, the meaning may be so abstract that it is difficult to describe), and it can't be further subdivided into smaller pieces that have meaning. Such a unit is called a **morpheme**. Note that by this definition a full-fledged word like *cat* is also a morpheme because it has a meaning and cannot be broken into smaller meaningful units.

Morphemes near the lexical end of the lexical-grammatical continuum are called **lexical morphemes**; morphemes such as *the*, *-s*, and *re-* near the grammatical end of the continuum are called **grammatical morphemes**. Note that grammatical morphemes include forms that we can consider to be words like *the*, *a*, *and*, and *of* and others that make up parts of words like *-s*, *-ed*, *un-*, and *re-*.

In a word consisting of more than one morpheme, there is normally at least one lexical morpheme. Thus the word *walking* consists of two morphemes, a lexical morpheme, *walk*, and a grammatical morpheme, *-ing*. The word *cropduster* consists of three morphemes, two lexical morphemes, *crop* and *dust*, and a grammatical morpheme, *-er*. In this section we'll only be looking at words with a single lexical morpheme and one more grammatical morphemes. In these cases we'll call the lexical morpheme the **root** of the word. So in the word *walking*, the root is *walk*, and in the word *carefully*, the root is *care*.

Kinds of Morphological Combination

Exercise 7.1.2

How many morphemes do you think the word *feet* contains? If more than one, what are they?

A root combines with one or more grammatical morphemes in various ways. In this section, we'll look at the different possibilities that exist in the world's languages.

Affixation

Grammatical Morphemes Can Be Added Before, After, and Within Roots

The examples we've seen so far involve adding grammatical morphemes before or after the root. When they precede the root, they are called **prefixes**; when they follow it, they are called **suffixes**. We can also speak of the processes of adding these morphemes; these are called **prefixation** and **suffixation**. Prefixation and suffixation are the most common ways in which grammatical morphemes combine with roots in the world's languages. Note that a single word can include more than one suffix and more than one prefix. For example, the word *muddier* includes two suffixes, *-y* (spelled "i" in this word) and *-er*.

In English the root of a word with one or more prefixes or suffixes is usually a word in its own right. Thus the root of *walked*, *walk*, is a word, and the root of *taller*, *tall*, is a word. This is not always true in other languages. In Japanese, Spanish, Lingala, and Inuktitut for example, every verb must have a grammatical suffix of one sort or another; the root cannot occur by itself. The Japanese verb *yobu* means 'call + present'; that is, 'call' in the present time. It consists of the root *yob-* 'sing' and the grammatical morpheme *-u* 'present', but *yob-* cannot occur by itself as a word; in fact it is not even a pronounceable Japanese syllable.

There are also two less common ways to add a grammatical morpheme. One is a single morpheme that combines a part before the root and a part after the root. For example, the Amharic verb *alhedem* means 'he didn't go'. In this word the part that makes the negative, that is, that corresponds to English *not* consists of two parts, *al-* and *-m*. Such a morpheme is called a **circumfix**.

Another possibility is a morpheme that gets inserted within a root, breaking up the phonemes of the root. Such a morpheme is called an **infix**. In Tzeltal infixation can apply to some verbs for human actions to yield a form used for counting the actions. For example, *lotz* is a verb root meaning 'strike with the hand', and the word *lojtz* is used for counting blows of the hand. This word consists of the root *lotz* and the infix *-j-*, which is inserted right after the vowel in the root.

Suffixes, prefixes, circumfixes, and infixes are all types of **affixes**, morphemes that are added to a root.

Mutation

Additional Morphemes Don't Necessarily Mean Longer Words

Rather than add material, a grammatical morpheme can change some part of the root; this is called **mutation**. English examples include the past forms of some verbs. From the verb root *sing*, there is the past form *sang*; from the verb root *take* the past form *took*. In *sang*, the vowel in the root, /I/, has been changed to /æ/. In *took*, the vowel in the root, /e/, has been changed to /U/.

In sign languages it is relatively simple to produce separate morphemes simultaneously. In this sort of case, we can see a grammatical morpheme as modifying the lexical morpheme that it is superimposed on, a kind of mutation. For example, the basic sign for 'give' is shown below; it can be produced with either one or two hands.


However, the direction of the movement is modified to reflect who is the agent and who is the recipient of the giving. If the signer is the recipient, for example, the movement is toward rather than away from the signer. Such signs can be seen as a combination of the root sign for 'give' and a grammatical morpheme which "mutates" the root by modifying the direction of the movement.

Template

So far all morphemes have consisted of segments (sounds) which appear consecutively in a particular sequence (unless they are interrupted by an infix). But a morpheme may also be discontinuous; that is, its segments may normally be separated by segments from other morphemes. This process is best known from Semitic languages like Arabic, Hebrew, and Amharic. Here are some Amharic verb examples.

5. *mesber* 'to break'
6. *sebberē* 'he broke'
7. *yIsēbral* 'he breaks'
8. *sIber* 'break!'

What all of these words share is sequence of segments *sbr*. This sequence makes up the root of the verb that means 'break' in Amharic. These three consonants never actually appear in sequence, however. In fact, they are not even pronounceable as a sequence in Amharic. Instead they combine with another morpheme specifying the time of the breaking and various other meanings to form a pronounceable form, called the "stem" of the word, to which prefixes and or suffixes are added. This morpheme that combines with the root is a kind of template for making the stem. It consists of vowels and positions in the stem for the root consonants. In addition, it sometimes specifies that the root consonants are lengthened. Here is how the root *sbr* combines with the past morpheme to make the stem for the word 6 above, *sebber-*.



To make word 6 above, a suffix meaning 'he', $-\epsilon$ is added to the stem. Word 7, *yIsebral*, has both a prefix, *yI-*, and a suffix, *-al*. These are added to the stem for that word, *-sebr-*, which means 'break + present'.

Reduplication

Another common way to combine lexical and grammatical morphemes, though not common in the languages of Europe, is by copying some portion of the root. The copied part is placed before, after, or within the root, as if it were an affix. It differs from an affix in that its form depends entirely on the root. Here are some examples from Amharic.

9. *get* 'jewel'; *get'aget* 'jewelry'
10. *bIret* 'iron'; *bIretabIret* 'hardware'
11. *sebbere* 'he broke'; *sebabberε* 'he repeatedly broke'
12. *lek'k'emε* 'he picked up'; *lek'ak'k'emε* 'he picked up multiple things'

In the first two examples, a noun root (which is also a word in its own right) is given first. In the complex form given second, there are two morphemes. The grammatical morpheme takes the form of the complete reduplication of the noun separated by *-a-*. This morpheme extends the meaning of the root noun to a general category of merchandise including the category that the original noun designates.

The second two examples are more complicated. Recall from the discussion of template morphology above that an Amharic verb root can consist of three consonants. From the examples there, you know that the root for the verb meaning 'break' is *sbr*. In 11, *sebabberε* includes a copy of the second consonant of the root and an additional vowel, *-a-*, separating that consonant from the root consonant: *sebabberε*. Notice that the same pattern appears in the last example. The verb root for these words is *lk'm* 'pick up'. The reduplication morpheme in the last two examples means 'repeatedly' or 'with multiple objects'.

Solving Morphology Problems

Exercise 7.1.3

What would be required for a language learner to figure out that words like *apples*, *tigers*, and *caves* consist of two morphemes?

The combination of grammatical morphemes with a root is normally a **productive** process. That is, given knowledge of the morphemes and a rule for how to combine them, a Speaker can produce novel words and a Hearer can understand novel words. For example, as an English speaker, you could learn a new noun *vorg* referring to some category of object, and you could then produce the form of the noun to refer to more than one *vorg*, *vorgs*. You would not have to be told how to do this.

The **morphology** of a language is the set of conventions that govern how words in the language are made up out of morphemes. Morphological conventions concern both the forms of words and the meanings of words. Just as learning a language involves learning the lexical, phonological, and syntactic conventions that were discussed in previous chapters, it involves learning morphological conventions. And a linguist analyzing a language will also want to figure out the language's morphology, both in order to describe the language adequately and to discover what universal principles apply to morphology.

Languages Differ Greatly in Terms of the Number and Type of Grammatical Morphemes that Combine with Lexical Morphemes

To learn morphology, you need examples of words along with their meanings. If you are a second language learner or a linguist, in place of meanings, you may be given a translation of the words into a language you already know. The translation may be deceptive in one way. The number of morphemes in a word in one language will often not correspond to the number of morphemes in a translation of that word into another language. The Japanese word *tabemashita* means 'ate', that is, 'eat + past', but it also

includes a morpheme which conveys the formality of the situation where it is used, a morpheme that would not normally be translated at all in English. To make things simpler, in what follows I will give a gloss for each word that makes it clear how many morphemes the word contains.

Three other factors can also make morphology complicated to figure out. First, within a given category of words such as verbs, there may be subcategories with different morphological behavior. In Spanish, and other Romance languages, there are three different categories of verbs, each with its own suffixes. Second, the form of a morpheme may depend on what other morphemes it combines with. For example, in *needed*, the suffix is pronounced /@d/ whereas it is pronounced /t/ in *walked*. You'll learn more about this in a later section. Third, a single grammatical morpheme may combine several different meanings. This is common with Spanish verbs. For example, in the verb *cantó* 'he sang', the morpheme *ó* (the accent mark indicates stress on this vowel) means 'he + past'. In this section, all of the examples will avoid difficulties of these three types.

Let's consider an Amharic example. You already know from the section on template morphology above how the past form of at least one verb is produced. Here are some more examples of verbs in the past, together with their glosses in English. Where items are separated by a colon, they represent the meaning of one Amharic morpheme.

- 13. *mekkeren* 'he advised me'
- 14. *mekkeres* 'he advised you:feminine'
- 15. *mekkerε* 'he advised'
- 16. *mekkerkuh* 'I advised you:masculine'
- 17. *mekkerkuš* 'I advised you:feminine'
- 18. *seddebu* 'they insulted'
- 19. *seddebun* 'they insulted us'
- 20. *seddebku* 'I insulted'
- 21. *seddebes* 'he insulted you:feminine'
- 22. *seddeben* 'he insulted us'
- 23. *k'et't'eren* 'he hired me'
- 24. *k'et't'erεh* 'he hired you:masculine'
- 25. *k'et't'erkuh* 'I hired you:masculine'
- 26. *k'et't'eruš* 'they hired you:feminine'
- 27. *k'et't'erun* 'they hired me'

First, examine all of the glosses to get an idea of what the range of morphemes is. You can see that every word includes either three or four morphemes, that every word includes the past morpheme and a verb root, that every word includes a morpheme representing the subject of the verb, and that some of the words include a morpheme representing the direct object of the verb.

Next, look for pairs of words that differ by only one morpheme. One example is *seddebu* 'they insulted' and *seddebun* 'they insulted us'. The glosses indicate that these two words share the morphemes meaning 'they', 'insult', and 'past', and that the second word has an additional morpheme meaning 'us'. Examining the forms, try to figure out what is added or changed in the form to give the second; whatever this is should mean 'us'. It's easy to see that the second form has the suffix *-n*, so we can assign this suffix the meaning 'us'. At this point, we should note the position of this morpheme because we not only expect it to always occur in the same position, but it is also likely that other direct object morphemes will occur in this position.

Another pair differing by only one morpheme is *mekkeren* 'he advised me' and *mekkeres* 'he advised you:feminine'. These two words share the morphemes meaning 'he', 'advise', and 'past', and they differ in the fourth morpheme, which means 'me' in the first word and 'you:feminine' in the second word. Note that this is 'you:feminine' as direct object. From the two forms we see that they share *mekkerε*, so this must mean 'he advised'. The remaining parts of the words are *-n* and *-š*, so these must mean 'me' and 'you:feminine'.

Continuing in this way, and given what you already know about the past of Amharic verbs, you should be able to figure out the following.

- 28. *mkr* 'advise'
- 29. *sdb* 'insult'
- 30. *k't'r* 'hire'

- 31. $C1 \varepsilon C2C2 \varepsilon C3$ 'past'
- 32. $-\varepsilon$ 'he'
- 33. $-ku$ 'I'
- 34. $-u$ 'they'
- 35. $-\eta$ 'me'
- 36. $-n$ 'us'
- 37. $-h$ 'you:masculine'
- 38. $-\text{\textasciix}$ 'you:feminine'

You should also know that the order of the morphemes is the following.

verb_stem subject direct_object

Given this knowledge and a new verb root, you should be able to predict some words. For example, say you're told that the root for the verb meaning 'resemble' is *msl*. You would then guess that the word for 'I resembled you:masculine' is *messeleh*.

This page titled [7.1: Morphemes](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.2: Grammatical Categories and NPs

Grammatical Categories and Inflection

Exercise 7.2.1

Removing the grammatical morphemes in bold in the following sentences makes them ungrammatical, but how does it affect their interpretability? That is, would any information be lost?

- *Jimmy wrote the letters.*
- *Jimmy wrote two letters.*
- *Clark **un**buttoned his shirt.*
- *Lois **reread** the chapter.*

In the last section, we saw that morphemes can be divided into those with relatively specific meanings and belonging to large, open-ended classes — lexical morphemes — and those with very abstract meanings and belonging to small, closed classes — grammatical morphemes. In this section and the next, we'll look more closely at some of the meanings and functions that grammatical morphemes have. Grammatical morphemes are always associated with a particular lexical morpheme. They may be combined with the lexical morpheme to form a single word, as in *apples* or *walked*, or they may form a separate word that belongs to the same phrase as the lexical morpheme, as in *the apple* or *is walking*.

Grammatical morphemes have two basic kinds of functions distinguished from one another in terms of how the morphemes relate to the lexical morpheme that they combine with. One function, the subject of the next chapter, is the creation of a new concept based on the meaning of the lexical morpheme. For example, in *shorten*, the *-en* takes the meaning of the adjective *short* and turns it into a change of state along the dimension of length. In the process *-en* makes a verb out of the adjective. This function of grammatical morphemes is called **derivation**.

The other function of grammatical morphemes, the subject of the rest of this chapter, is similar to modification; the grammatical morpheme specifies some very abstract feature of the category that is the meaning of the lexical morpheme. In other words, its meaning is a very abstract **grammatical category**. For example, in *walked*, the *-ed* specifies that the walking took place before the time of speaking; it assigns the feature past to the event. In other words, past, contrasting with present and future, is a grammatical category in English. The combination of a grammatical morpheme with a lexical morpheme to form a word, as in *walked*, is called **inflection**. As we'll see, though, grammatical categories can also be defined by grammatical morphemes that are separate words.

Languages differ quite strikingly in terms of which grammatical categories are built into their morphology. In this section I'll describe a few of the kinds of grammatical categories that play a role in noun phrases. In the next section, I'll describe some grammatical categories that are marked on verbs.

Number

People not only have the capacity to recognize individual objects in their environment and categorize them as apples, stones, people, etc. They have the ability to recognize sets of objects that share a category, for example, sets of apples, stones, or people. Though an individual and a set seem to be very different things, the categorization process for the individual and for the elements of the set must be similar. This is reflected in an apparently universal property of human language: the same morpheme is used for individual objects belonging to a category and for sets of objects whose members belong to that category. In English, the morpheme *apple* is applied both to individual apples and to sets of apples.

People have a further ability; they can assign a cardinality to a set, that is, they can tell (or estimate) how many elements are in the set. And apparently all languages have systems of **numerals** such as *two* and *eight*. Each numeral is a label for a category of set, independent of what kinds of members the set has. For example, *eight* labels the category of sets consisting of eight elements.

Now let's imagine two tribes of Grammies. One uses common nouns like *apple* and *tiger* and numerals like *two* and *eight*, as well as adjectives like *many*, to talk about individuals and sets and finds that these forms suffice. They say things like *give me apple* whether they want one or several, and when it matters, they say things like *give me two apple* or *give me several apple*.

In another tribe, for one reason or another, a subgroup of members begins to explicitly mention whenever they are talking about a set rather than an individual. So they say things like *give me apple*, *some* whenever they want more than one and *give me apple* when they want exactly one. But they leave out the *some* when there is a numeral because the numeral makes it clear that more than one is intended. This practice catches on, and eventually two things happen. First, because the *somedoesn't* convey very much information, it gets pronounced more and more quickly and carelessly, and eventually all that's left of it is the *s* at the beginning. This *s* is pronounced as if it were part of the noun that it follows, and it even assimilates to the voicing of the last phone in the noun, so it is pronounced /z/ in *apples*. Second, the members of the tribe find it weird to say *apple* whenever they mean more than one, even when the context makes it clear that they do. So now they say things like *give me two apples*.

Even though this story is completely fictitious, it illustrates what has apparently happened in two kinds of modern languages. English is a language of the second type. It is ungrammatical in English to say *apple* when more than one apple is referred to. It is of course equally ungrammatical to say *apples* when only one apple is referred to. English grammar makes a two-way distinction in the way objects are referred to: individual objects and sets of objects are referred to differently. That is, English has the grammatical dimension number with two values or grammatical categories, singular and plural. English nouns are **inflected** for number, and number inflection is **obligatory**. Thus *three apple* and *lots of person* are ungrammatical in English.

The Grammar of a Language May "Force" its Speakers to Use Certain Morphemes in Certain Contexts, Even When They Seem to Contribute Nothing to the Meaning

Notice that in the case of the phrase *three apples*, the plural morpheme, *-s*, doesn't really carry any information; the numeral already makes it clear that more than apples is being referred to. That is, the grammatical morpheme is **redundant**. Redundancy is a frequent property of grammatical morphemes. Because they are obligatory, Speakers in a sense do not ask themselves whether they are necessary when producing sentences; they insert them in any case. It may seem odd that language would allow redundancy, but it is probably helpful to Hearers. Redundancy permits Hearers to understand the message even when they miss some part of it.

Japanese is a language of the other kind. Japanese does have a morpheme for plural, *-tachi*, but it can only be suffixed to nouns referring to people or animals, and it seems never to be obligatory. (The same morpheme is used for plural personal pronouns, meaning 'we' and 'you plural', and here it is obligatory.) Apparently number is not a dimension in Japanese grammar, except for the limited case of personal pronouns. Here are some examples that should make this clearer.

1	<i>ringo</i>		<i>ga</i>		<i>hoshii</i>
	apple		nom		want
	'I want some/an apple(s).'				

2	<i>ni-ko</i>	<i>no</i>	<i>ringo</i>	<i>ga</i>	<i>hoshii</i>
	two-class	of	apple	nom	want
	'I want two apples.'				

Singular and Plural are Part of the Grammar of English; They are Not Part of the Grammar of Japanese

In the first sentence, there is no indication of whether one apple or a group of apples is desired. It's perfectly grammatical in Japanese to leave this unspecified. The Hearer might be expected to figure out which is intended if it matters at all. In the second sentence it is clear that more than one apple is desired, but note that, unlike in English, there is no morpheme that explicitly indicates plurality. (The second sentence also contains a morpheme which I've indicated with *class*; more on this later.)

What does this mean about English and Japanese speakers? It certainly does *not* mean that Japanese speakers are incapable of understanding the difference between individuals and sets. Speakers of all languages not only understand this distinction but have ways of expressing it in their languages. What it means is that English grammar forces its speakers to make the distinction in places where other languages do not and in many cases cannot. In fact, in Japanese it's not only acceptable to leave the plural unmarked; it is impossible to mark plural on the noun for 'apple'. That noun, *ringo*, has no plural form.

Number Appears in English Grammar in Multiple Places

When a dimension such as number is part of the grammar of a language, it often turns up in more than one place. This is true for number in English. Consider the following sentences.

3. *An apple is on the table.*
4. *Some apples are on the table.*

Apple and *apples* are preceded by the words *an* and *some*. These words are called **indefinite articles**; both function roughly to say that the thing referred to is not already known to the hearer. But they differ in another way: *an* (or *a*) is used only before singular nouns, while *some* is used before plural nouns (and also before some singular nouns; more about this below). That is, these words also distinguish singular from plural. The verbs in the two sentences are also different. *Is* is appropriate only when the subject is singular, whereas *are* is used when the subject is plural. Again, the distinction between singular and plural matters somewhere in the grammar of the language. In Japanese, on the other hand, there is no distinction like that between *an* and *some* or between *is* and *are*.

Another thing to notice about sentences 3 and 4 is the degree of redundancy. Sentence 4 indicates in three places that multiple apples are being referred to, in the use of *some* rather than *an*, in the presence of the -s, and in the form of the verb *are*.

Countability

Exercise 7.2.2

In English we can say *lots of milk*, *lots of sand*, and *lots of salt*, but not normally *lots of milks*, *lots of sands*, and *lots of salts*. On the other hand, we can say *lots of girls*, *lots of trees*, and *lots of rivers*, but not normally *lots of girl*, *lots of tree*, and *lots of river*. What do you think is the difference between these two kinds of nouns?

There is another grammatical dimension with two values in English that is tied up with the use of plural and the distinction between *a(n)* and *some*. Consider these sentences.

5. *Some rice is on the table.*
6. *Two piles of rice are on the table.*

Notice that in sentence 5, *rice* is singular, and the verb is also in the singular form *is*. However, instead of *a*, the noun is preceded by *some*, the form used with a plural noun in sentence 4. In fact no matter how much there is on the table, we still won't say *some rices*. If the Speaker wants to mention the amount of rice, they have to use another noun such as *pile* or *bowl* or *cup*, putting *that* noun in the plural, as in sentence 6.

In the English Lexicon, Peas are like Beans and Potatoes; Rice is like Sugar and Vinegar

Apparently English has two kinds of nouns. One kind, **count nouns**, is used mainly for objects (and for abstract things that are construed as object-like). In the singular these nouns may be preceded by the article *a(n)*, and they are always pluralized when more than one of the objects is referred to. The other kind, **mass nouns**, is used mainly for masses (and for abstract things that are construed as mass-like). These nouns are always singular except in the special sense of 'multiple kinds' (for example, *wines* referring to different brands or varieties of wine), and they may be preceded by the article *some*. Of course there is a gray area between clear cases of objects and clear cases of masses, and in this area, a noun can go either way. Thus *rice*, as we have seen, is a mass noun. But *pea*, which designates something that, like rice, consists of small objects usually gathered together in a group, is a count noun. (In fact *pea*, in the form *pease*, used to be a mass noun like *rice*.)

So English has the dimension of countability built into its grammar. But note that it appears in the language in two places, in the grammatical forms that go with one or the other category (*a(n)* with singular, *some* with plural for count; *some* with singular and no plural for mass) and in the lexicon, where most nouns belong to one or the other type. That is, there is a strong tendency in English for the count grammatical patterns to go with certain nouns (such as *apple* and *house*) and the mass grammatical patterns to go with other nouns (such as *rice* and *milk*).

Unlike English Nouns, Spanish Nouns Don't Have *Inherent Countability*

Spanish differs from English in an interesting way. Like English, Spanish has both number and countability in its grammar. The count and mass grammatical patterns are similar to those in English, except that *some* usually corresponds to no article at all in Spanish. But Spanish differs in that its nouns do not strictly belong to one category or another. Apparently most nouns can be used with either mass or count morphology depending on what meaning is intended. Here are examples with the nouns *madera* 'wood' and *papel* 'paper'.

7. *madera* '(some) wood'
8. *una madera* 'a board'
9. *papel* '(some) paper'
10. *un papel* 'a sheet of paper'

Actually the difference between English and Spanish is a matter of degree. English also has some nouns that can appear with either mass or count morphology. For example, when the noun *chicken* designates the bird, it is treated as a count noun (*a chicken*, *some chickens*), but when it designates the meat of the bird, it is treated as a mass noun (*some chicken*). But Spanish goes a lot further in this respect; many more nouns can be treated in both ways.

The main point of this example is that a dimension like countability can be both a way in which things in the world are divided up conceptually by speakers and a way in which the words in the lexicon are divided up. Countability in English is both conceptual and lexical, whereas in Spanish the conceptual aspect predominates.

Spanish grammar makes another division that is even more lexically oriented than countability is in English. All Spanish nouns have a gender, belonging either to the masculine or the feminine grammatical category. A noun's gender affects several aspects of Spanish morphology, including the form of adjectives that modify the noun and the pronouns that replace a noun. Though gender seems to have had its beginnings as a categorization of the things in the world on some conceptual basis, and it is still true that nouns for female animals and people tend to be feminine while those for male animals and people tend to be masculine, for the most part the membership of nouns in one category or the other seems fairly arbitrary now. That is, Spanish gender is an example of a grammatical dimension which is much more lexical than it is conceptual.

Classifiers

We have seen three ways in which languages may divide the things that speakers talk about into two very general categories, on the basis of whether they are individuals or sets, on the basis of whether they are masses or objects, and on the basis of a single conceptual property (biological gender) that is extended more or less arbitrarily to cover all labeled categories of things. Another possibility, found in many languages, is a somewhat finer-grained grouping into a larger set of categories, each of which is still more general than the kind of category represented by a noun such as *apple*, *baby*, or *paper*. Each of these abstract categories is represented by a grammatical morpheme called a **classifier**. The most common basis for the classification of things appears to be shape, but it may also be based on orientation, animacy, function, or cardinality (for sets).

When Objects are Counted in Japanese, Their Shape Must Often Be Taken into Account

In many languages with classifiers, including Mandarin Chinese, Japanese, and Tzeltal, classifiers are used along with numerals when objects are counted. Many classifiers in these languages are shape-based. Here are some Japanese examples. The classifiers are indicated by the label class in the second line of each example.

11	<i>ni-ko</i>	<i>no</i>	<i>ringo</i>
	two:CLASS	of	apple
	'two apples'		
12	<i>ni-hon</i>	<i>no</i>	<i>enpitsu</i>
	two:CLASS	of	pencil
	'two pencils'		
13	<i>ni-mai</i>	<i>no</i>	<i>kami</i>
	two:CLASS	of	paper
	'two sheets of paper'		
14	<i>ni-ko</i>	<i>no</i>	<i>nasu</i>
	two:CLASS	of	eggplant
	'two (roundish) eggplants'		
15	<i>ni-hon</i>	<i>no</i>	<i>nasu</i>
	two:CLASS	of	eggplant
	'two (long, narrow) eggplants'	'two (long, narrow) eggplants'	'two (long, narrow) eggplants'

The three classifiers illustrated here, *-ko*, *-hon*, and *-mai*, represent categories based on the gross shape of the object referred to. *-ko* is used when the object has roughly the same extent in all three spatial dimensions, for example, when it has the shape of a sphere or a cube. *-hon* is used for long, narrow objects like pencils, legs, and cucumbers. *-mai* is used for flat objects like sheets of paper, boards, and pancakes. These morphemes are obligatory in the phrases above; leaving them out, or using one that doesn't agree with the shape of the referent, would result in an ungrammatical phrase.

As with other grammatical morphemes, classifiers are often redundant. This is true in sentences 11, 12, and 13, where the noun itself makes the gross shape clear. However, there are cases where more than one shape is consistent with a particular noun, for example, the noun *eggplant*. Eggplants can be roughly pear-shaped (the usual European variety) or long and narrow (the usual Asian variety). With a noun like this, classifiers can be informative, as they are in sentences 14 and 15.

Note that English has words which behave something like classifiers in Japanese. When we count masses in English, we must use words like *cups* or *pieces* after the numerals. In fact one way to think about the way languages like Japanese, Chinese, and Tzeltal work is to see them as treating all nouns as mass nouns. You can't count apples directly using the Japanese noun *ringo* because it really means something more like 'apple stuff' than 'apple individual'.

ASL also uses classifiers. They take the form of handshapes that signify general properties of things, much like those represented by classifiers in Japanese, Chinese, and Tzeltal. But rather than being combined with numerals in NPs, ASL classifiers are combined with verbs, so they are discussed in the next section along with verb inflection.

This page titled [7.2: Grammatical Categories and NPs](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.3: Grammatical Categories and Verbs

Categories of Verb Morphology

Exercise 7.3.1

What properties of the events described in the following sentences do the morphemes in bold tell us about?

- Jimmy **will** graduate in June.
- Jimmy **would** graduate if he studied.
- Jimmy **is** sleeping.

In the last section we saw how grammatical morphology can specify one or another abstract category for the things that nouns refer to. In this section, we'll look at how grammatical morphology can do the same for verbs, focusing on one particular kind of verb morphology, morphemes that indicate general properties of the participants in the event or state that the verb designates.

Just as things divide naturally into a small number of categories on the basis of dimensions such as number, countability, and shape, events and states also divide naturally into a small number of categories on the basis of several basic dimensions.

Time

The Grammys realized early on that when an event occurred or a state was true often mattered. An utterance like *Clark eat berries* wasn't much use if the hearer didn't know whether Clark had already eaten the berries, was eating them at that moment, or was going to eat them at some later time. The Grammys developed two kinds of expressions to help them talk about the time of an event or state, *absolute* and *relative* expressions. This is a distinction we've seen before, in the context of adjective meaning.

Absolute time expressions label specific points in time, such as *January 20, 1203*, or points within a repeating unit of time, such as *3:00 pm* (which labels a time within the day) and *Tuesday* (which labels a day within the week). The second type of expression may be used for repeating events or states (*I get up at 7:00*) or for a single event or state, in which case the hearer has to be able to figure out which unit of time the speaker has in mind. That is, *I got up at 7:00* is only meaningful if we know which day the speaker is talking about.

Expressions like *Yesterday* and *Ago* Express Times Relative to the Utterance Time

Relative time expressions label points in time relative to some other reference point. The most obvious reference point is the utterance time, which is one of the roles in the utterance context and is directly accessible to the hearer. Thus referring to time in this way is an example of a deictic use of language. For an event or state that is going on at the time of speaking, we have a word like *now*. For a past or future event or state, we can mention the length of time that has elapsed or will elapse between the time it occurred or will occur and the utterance time (*an hour ago*, *in an hour*), or we can simply say that it happened before the utterance time or will happen after the utterance time (*already*, *in the future*). There are other possible reference points for relative time reference. We can say things like *before that time* and *after the wedding*.

Just as number ended up grammatical in languages such as English, we might expect reference to the time of events and states to end up grammatical too. In fact, many, if not most, modern languages have a system for this, called **tense**, built into their grammar. For example, we distinguish *Clark fell asleep*, *Clark is falling asleep*, and *Clark is going to fall asleep*. Tense morphology divides events and states into the general grammatical categories past, present, and future; or a smaller set such as past and non-past; or a larger set, depending on the language.

As with other grammatical morphology, tense marking is normally obligatory in languages that have it, even when it is redundant. Both of the following English sentences have the past morpheme, even though that morpheme is redundant in the second example because the phrase *last night* makes it clear that the event happened before the utterance time.

1. *I slept ten hours.*
2. *I slept ten hours last night.*

Events May be Viewed "From Inside", as They are Going On, or "From Outside", Before They Begin or After They Finish

There are other ways of looking at the temporal properties of an event or state than when it occurred or was true. It could be viewed as ongoing or completed, for example. Consider the difference between these two English sentences.

3. *Clark was falling asleep.*
4. *Clark had fallen asleep.*

Both have an unspecified time in the past as a point of reference. In sentence 3 the event is seen as ongoing at that time, and in sentence 4 the event is seen as completed at that time.

The Speaker may also point out the repeated nature of an event or state. Consider the difference between these English sentences.

5. *Clark runs in the marathon.*
6. *Clark is running in the marathon.*

For both of these sentences, the point of reference is the utterance time ('now'). In sentence 5, the running is viewed as repeated around this reference time; in sentence 6 it is ongoing at the reference time.

The grammatical representation of duration, completion, and repetition of events and states is known as **aspect**. As with other grammatical morphology, aspect morphology is often obligatory. In English, for example, speakers have to commit themselves to the choice between ongoing, repeated, or completed for an event with present reference time. That is, it is impossible in English to talk about Clark running the marathon, as in sentences 5 and 6, without making such a commitment.

Possibility, Hypothesis, Desirability

Another set of properties that distinguishes some events and states from others is related to their truth: whether they are true or likely to be true, whether we are treating them as true just for the sake of argument, whether we would like them to be true. The grammatical representation of meanings like these is called **modality**. Here are two English examples where the verb morphology reflects these dimensions.

7. *If Jimmy spoke Spanish, he'd have a better chance with Lupe.*
8. *Perry suggested that Clark spend less time on computer games.*

In sentence 7, the Speaker knows that Jimmy doesn't speak Spanish; if he did or there were at least a possibility that he does, the verb would be *speaks* rather than *spoke*. And in the same sentence, *would* ('d) indicates the conditional nature of the state of "having a better chance"; it would be true if Jimmy spoke Spanish, but he doesn't, so it isn't. In sentence 8, *spend* is used rather than *spends*, indicating the tentative nature of the "spending less time"; this is only a suggestion, not yet reality.

Participants

Events and states are defined in part by their participants. The choice of a particular verb commits the Speaker not only to a category of state or event but to a set of semantic roles. But these semantic roles may often be filled by a variety of things. We can group events and states into a small set of abstract categories on the basis of some general properties of these participants. The next subsection focuses on verb morphology with this function.

Verb Agreement

Exercise 7.3.2

What makes the following sentences ungrammatical? What kind of rule can you specify for the verb morpheme -s?

- *Clark always arrive late.*
- *Clark's colleagues likes him a lot.*

In many languages verbs take inflectional morphemes that convey some information about one or more participants in the event or state that the sentence is about. One way to think about this is in terms of the **agreement** between the verb and those participants on a small number of abstract properties. On the one extreme are languages like Mandarin Chinese and Japanese, which have no morphology of this type (though sometimes the choice of a verb in Japanese is governed by some properties of the participants). In what follows, I'll briefly discuss verb agreement in four languages that have some form of it. Notice that since agreement morphology conveys abstract properties of participants, that is, things, this topic overlaps with the topic of the last section.

English

English is a language with limited verb agreement morphology, the vestiges of what was a full-blown agreement system in Old English. Consider these sentences.

9. *Clark plays golf.*
10. *Lois and Clark play tennis.*
11. *I play croquet.*
12. *Clark played 18 holes yesterday.*
13. *Clark likes team sports.*

In English -s is Plural When It Appears on Nouns But Singular When It Appears on Verbs

Notice that the form of the verb *play* differs in sentence 9 and 10. In sentence 9 the subject of the sentence, *Clark*, is 3rd person (that is, including neither the Speaker nor the Hearer) and singular, and the verb takes the suffix *-s* to indicate this. When the same verb is used with a subject that has any other combination of person and number, as in sentences 10 and 11, the verb takes no suffix. Notice also that an agreement suffix is only added to verbs in the simple present tense, that is, the tense category used in sentences 9, 10, and 11. Sentence 12 is in the simple past tense, and no distinction is made on the basis of person and number. Finally, notice that it is the participant in the *syntactic* role of subject, rather than any particular semantic role, that the verb agrees with. So in sentence 13, the verb again takes the *-s* even though the subject in this case refers to an experiencer rather than an agent, as in sentence 9.

With the verb *be*, there are three forms rather than two in the simple present, and rather than suffixes, completely unrelated forms are used: *am* (1st person singular), *is* (3rd person singular), and *are* (other person-number combinations). The verb *be* also has two forms in the simple past tense, *was* and *were*.

Thus English subject-verb agreement is limited both in terms of the number of different forms and the situations in which it must apply. However, it behaves just like the other examples of grammatical morphology we've been considering. It is often redundant, but it is obligatory even when it is. So in standard English dialects, at least, it is ungrammatical to say *Clark like Lois*, even though the missing *-s* would convey no new information.

So does the *-s* in *play* in sentences 9 and 13 mean anything? Yes, it means that the subject of that verb is 3rd person singular. In addition, because this suffix only occurs on verbs in the simple present tense, it also marks that tense category. Under most circumstances, this information would be obvious from the subject itself and from the context. But if the Hearer missed the subject for some reason, that *-s* could help sort things out. Also there are gray areas where Speakers may choose to use a verb in the 3rd person singular with a plural subject. Compare these two sentences.

14. *A hundred students are in this course.*
15. *A hundred students is more than this room can hold.*

In sentence 15, the subject is viewed as an individual quantity rather than a collection of individual things, so the verb is singular.

Spanish (and Japanese)

Spanish verbs agree with their subjects in all tenses, aspects, and modalities (TAM), and the number of different forms is at least four in each case, depending on the particular combination of TAM. The actual morphemes used vary with the TAM. In fact in most cases, the subject agreement morpheme also indicates the TAM. Let's look only at the past tense forms for one verb; remember that English makes no distinctions in this tense except for the verb *be*.

16. *habl-é* 'I spoke'
17. *habl-aste* 'you (singular, familiar) spoke'

18. *habl-ó* 'he/she/it/you (singular, polite) spoke'
19. *habl-amos* 'we spoke'
20. *habl-aron* 'you (plural)/they spoke'

A Spanish or Japanese Sentence Can Have a Subject Even When It Contains No Noun Phrase

The main thing to notice about these words is that they are also sentences. That is, Spanish is an example of a language that does not require **explicit** subjects. An explicit subject is a separate noun phrase functioning as subject. In these sentences there are no NPs, just verbs. But the verbs carry information about the subject, in this case, the agent of the speaking. In other words, these sentences do have a subject; it is just not expressed explicitly in an NP.

Of course Spanish sentences *can* have explicit subjects. If a 3rd person subject were not clear from the context, it would need to be spelled out, for example, *Lois habló* 'Lois spoke'. But what about cases where English would have a personal pronoun subject, such as *I spoke*, *you spoke*, and *he spoke*? Here Spanish can have an explicit subject as well, but it carries additional weight. *Yo hablé* or *hablé yo*, with the personal pronoun *yo* 'I', is more like an emphatic *I spoke* in English, for example, in a situation where the Hearer assumes somebody else was the one speaking.

But even without the pronoun that is required in the English sentence *I spoke*, the Hearer of the Spanish sentence can know who the subject is from the suffix on the verb. The extensive subject-verb agreement in Spanish probably has something to do with the fact that this language doesn't require explicit subjects.

Some Languages, like English, Have Syntactic Positions in a Sentence That Must be Filled; Others, like Japanese, Aren't So Rigid in This Way

But it is not quite that simple. Japanese is also a language that does not require explicit subjects, but Japanese has no subject-verb agreement. That is, the single verb *hanashita* 'spoke', with no indication of who (or what) the subject is, can function as a sentence on its own. Does this sentence have a subject? It does in the sense that the Speaker obviously had someone in mind who did the speaking, and if that person had been referred to in an NP, it would have been the subject of the verb *hanashita*.

This brings up another interesting difference between Japanese on the one hand and both English and Spanish on the other. Both English and Spanish have transitive verbs, verbs that have both a subject and a direct object, and for such verbs there must be an explicit direct object in both languages. That is, it is ungrammatical in English to say *Clark put*, even when it is clear what he put, and it is ungrammatical in Spanish to say the possibly corresponding "sentence", *Clark puso*. These verbs require an explicit direct object. In Japanese there are no such verbs. Just as subjects may be omitted if they are clear from context, direct objects of transitive verbs may be omitted if they are clear from context. Thus *oita* 'put (past)' is a perfectly good Japanese sentence. But just as it has an (implicit) subject, it has an (implicit) direct object.

Amharic

Like Spanish, Amharic has obligatory subject-verb agreement in all of the different TAM possibilities, though there are even more different agreement morphemes for a given TAM types because Amharic verbs have to distinguish feminine from masculine subjects in 2nd and 3rd person singular. And like Spanish, Amharic allows subjects to be omitted when they are clear from context. Further, Amharic uses personal pronoun subjects, even when they are clear from the subject-verb agreement morphemes, to put emphasis on the subject, again just as Spanish does.

A Verb in Some Languages May Convey Information About Multiple Participants in the Event or State

But Amharic goes beyond Spanish. As you have already seen from the Amharic examples in the section on morphemes, the language has direct-object-verb agreement. That is, a verb may have two agreement morphemes, one for the subject and one for the direct object. The circumstances in which the direct object morpheme is obligatory are complicated, but there are such circumstances. Here are a few examples in the past tense, similar to those you saw in the previous section. The root and past tense

morphemes are not shown separated; they combine by template morphology. For the glosses of the agreement morphemes, both the syntactic role and an English pronoun are given.

21	<i>messe</i> -	<i>ε</i> -	<i>h</i>
	resemble:past	subj=he	diobj=you:masc
	'He resembled you (masculine).'		

22	<i>messe</i> -	<i>n</i> -	<i>accew</i>
	resemble:past	subj=we	diobj=them
	'We resembled them.'		

23	<i>messe</i> -	<i>u</i> -	<i>n</i>
	resemble:past	subj=they	diobj=us
	'They resembled us.'		

Just as personal pronouns can be used as subjects in Spanish (and Amharic) when emphasis needs to be placed on them, personal pronouns can be used as direct objects in Amharic when emphasis needs to be placed on them, even when they are clear from the verb. So to emphasize the "us" of sentence 23, an Amharic Speaker could say the following.

24	<i>Inna</i> -	<i>n</i>	<i>messe</i> -	<i>u</i> -	<i>n</i>
	us	acc	resemble:past	subj=they	diobj=us
	'They resembled us.'				

In fact Amharic goes one step further. It even allows agreement with adjuncts, that is, objects of prepositions. There are only two possible "prepositions" that can appear within the verb. I'll just give examples of one of these, realized as the suffix *-bb*-followed by a suffix representing the object of the preposition. This morpheme agrees with a participant that is an instrument, a sufferer, a location, or a time. Here are two examples with this pair of morphemes used for a sufferer; in both cases the "pronoun" is 'us'. For *-bb*-, the gloss is just "B".

25	<i>t'effa</i> -	<i>acclhu</i> -	<i>bb</i> -	<i>n</i>
	disappear:past	subj=you:plur	B	us
	'We lost you guys.'			

26	<i>bella</i> -	<i>acclhu</i> -	<i>bb</i> -	<i>n</i>
	eat:past	subj=you:plur	B	us
	'You guys (went and) ate it on us. (You ate something to our disadvantage.)'			

Notice the way Amharic expresses loss to someone; sentence 25 is literally 'you guys disappeared on us'.

In summary, all Amharic verbs agree with their subjects, and some agree in addition with either their direct objects or an object of a preposition.

American Sign Language

The Grammars of Sign Languages May Be Just as Complex as Those of Spoken Languages

Finally let's consider agreement morphology on verbs in a sign language. We have already seen one example of this in the discussion of mutation morphology. ASL has a category of verbs that sign linguists call "directional verbs". These are verbs designating transfer events, or information_transfer events, or other events viewed as having a direction. These verbs have a basic handshape and a position on the body, but their direction has to agree with the source and the goal (often the recipient) of the event. The agreement is with what corresponds to person in ASL, the position in signing space of the participants. 1st and 2nd person have the position of the signer and the sign interpreter, and other participants are "placed" in signing space by the signer as they come up.

For example, to produce the sign for 'give' in ASL when the source/agent is neither the signer nor the sign interpreter and the recipient is the signer, the signer uses the basic handshape for 'give', moving one hand from the position of the giver in signing space to the signer's own chest. The direction would be the opposite if the roles were reversed.

Another form of agreement in ASL makes use of classifiers. Classifiers in ASL take the form of particular handshapes that represent general properties of things. For example, an index finger pointing upward represents a standing person, a cupped hand represents a container, and the extended thumb and first two finger represents a vehicle. One use of classifiers is as morphemes agreeing with the subjects of verbs designating move events and be_at states. In this case the agreement is the opposite of what happens with verbs of giving and telling. It is the handshape that represents the agreement morpheme and the movement of the hand(s) that represents the content of the verb. For example, to sign a sentence meaning 'the car is here', the signer would make the sign for 'car', then with the 'vehicle' classifier handshape sign 'be here', that is, move the hand downward in front of the body.

How is verb agreement in ASL like the verb agreement in the spoken languages we have considered? At least in many cases agreement in ASL is obligatory, as it is in spoken languages. It may also be redundant, as in the 'vehicle' example.

Agreement in ASL, in fact morphology in sign languages generally, is strikingly different from spoken language morphology in one way. It is invariably iconic; all of these examples we have seen "make sense". With respect to form alone, sign language grammatical morphology differs in another way from most spoken language grammatical morphology in that it occurs simultaneously with the root morpheme. Of course this derives from the potential in sign languages to maintain a particular handshape while a movement is executed.

One point of this section has been to show how much languages can vary in terms of what information gets represented on their verbs. It is on verbs that we see how different languages can get. Within our set of languages, we have seen a range of possibilities, but we still are not close to the extreme of some American Indian and Eskimo languages, like Inuktitut, where verbs frequently include more than ten morphemes. However, those words usually include morphemes that go beyond the functions we've discussed in this chapter. Such languages excel at creating new words from a small number of roots and extensive productive morphology. How this sort of process works is the topic of the next chapter.

This page titled [7.3: Grammatical Categories and Verbs](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.4: Problems

This problem concerns Lingala verbs. Each of the verbs in the lists below has three morphemes: a root; a morpheme indicating the person and number of the subject ('I', 'you:SINGULAR', 'he/she', 'we', 'you:PLURAL'); and a morpheme indicating some aspect of the time, the likelihood, or the desirability of the state or event. I'll refer to these three categories of morphemes as **root**, **subject**, and **tense-aspect-modality (TAM)**. The following values of TAM appear in the lists below (you don't need to bother with what they actually mean): SIMPLE PRESENT (SIMP PRES) PRESENT PERFECT (PRES PERF), FUTURE (FUT), SUBJUNCTIVE (SUBJ).

Tone, marked by the presence (high tone) or absence (low tone) of an accent mark on vowels, is important in the lexicon and the grammar of this language; don't ignore it! To write accent marks, you can put them after the vowels. Don't type accented characters in your word processing program, then save them as text and upload them to Annotate; they won't show up in Annotate. If you want accented characters to display in Annotate, you have to use the HTML codes for the characters. If you know these, feel free to use them.

In the first list of verbs, all are in the SIMPLE PRESENT form. You just need to figure out what the subject morphemes are and where they come in the words.

Word	Root	Subject	TAM
<i>nakoma</i>	write	I	SIMP PRES
<i>tokoma</i>	write	we	SIMP PRES
<i>tokóma</i>	arrive	we	SIMP PRES
<i>okóma</i>	arrive	you:SING	SIMP PRES
<i>osómba</i>	buy	you:SING	SIMP PRES
<i>asómba</i>	buy	he/she	SIMP PRES
<i>akabola</i>	divide	he/she	SIMP PRES
<i>bokabola</i>	divide	you:PLUR	SIMP PRES
<i>bondima</i>	agree	you:PLUR	SIMP PRES
<i>tondimá</i>	agree	we	SIMP PRES
<i>osepela</i>	enjoy	you:SING	SIMP PRES
<i>asepela</i>	enjoy	he/she	SIMP PRES
<i>natatabana</i>	be:surprised	I	SIMP PRES
<i>botatabana</i>	be:surprised	you:PLUR	SIMP PRES
<i>toloba</i>	speak	we	SIMP PRES
<i>nabóngola</i>	change	I	SIMP PRES
<i>oyébisa</i>	tell	you:SING	SIMP PRES
<i>atámbwisa</i>	drive	he/she	SIMP PRES
<i>botúna</i>	ask	you:PLUR	SIMP PRES
<i>napangwisa</i>	wipe	I	SIMP PRES
<i>atála</i>	look	he/she	SIMP PRES
<i>obóndela</i>	plead	you:SING	SIMP PRES

1. Where in the words does the subject morpheme appear?
2. What are the morphemes for the following subjects?
 - a. 'I'
 - b. 'you:SINGULAR'
 - c. 'he/she'
 - d. 'we'
 - e. 'you:PLURAL'

hint

I'll describe informally the process of how you might solve the problem.

Start by looking for some pairs whose meanings differ by only one morpheme. The first two words *nakomo* and *tokoma* differ only in that one has 'I' as subject, the other 'we'. The difference between the two forms is that the first begins with *na*, but the second begins with *to*. So tentatively at least we can assume the following:

- *na*- 'I'
- *to*- 'we'

Both of these morphemes come at the beginning of the word. Since morphemes belonging to the same set usually appear in the same place in the word, we can guess that the other subject morphemes will also appear at the beginning.

Before we go on, we might check to see if we are right so far by looking for other words with subject 'I' or 'we'. In fact all of the with subject 'I' begin with *na*, and all of those with subject 'we' begin with *to*.

Now we can either continue to look for pairs differing by only one morpheme, or we can try a different approach and look for all of the words containing a given unknown morpheme and see what they have in common. Let's try this second way and look for the morpheme meaning 'he/she'. The words containing that morpheme are *asómba*, *akabola*, *asepela*, *atámbwisa*, *atála*. These all share two things: they all begin with *a*, and they all end with *a*. But the *a* at the end could have nothing to do with 'he/she' because it is at the end of all of the words in the list (so probably has something to do with what all of the words share, that is, SIMP PRES). Based on what we already know, it is the beginning where we expect the subject morpheme to appear, so we can now conclude the following:

- *a*- 'he/she'

That leaves two more subject morphemes, 'you:SING' and 'you:PLUR'. We can do the same thing we did for 'he/she' and look for all of the words that share these morphemes. For 'you:SING' they are *okóma*, *osómba*, *osepela*, *oyébisa*, *obóndela*. These all begin with *o*. The words with subject 'you:PLUR' are all of the remaining words; they all begin with *bo*. So we can conclude

- *o*- 'you:SING'
- *bo*- 'you:PLUR'

Thus all of the subject morphemes are prefixes, appearing at the beginning of the words.

The following list includes verbs in different TAM forms. (Only forms with subjects 'I', 'you:SINGULAR', and 'he/she' are included; the other subject morphemes behave similarly.) You need to figure out what the verb roots and the TAM morphemes are. (Also refer to the words in the first list.)

Word	Root	Subject	TAM
<i>nakokoma</i>	write	I	FUT
<i>nakokóma</i>	arrive	I	FUT
<i>okosómba</i>	buy	you:SING	FUT
<i>akosómba</i>	buy	he/she	FUT
<i>nakotúna</i>	ask	I	FUT
<i>okoloba</i>	speak	you:SING	FUT
<i>akondima</i>	agree	he/she	FUT
<i>akoyébisa</i>	tell	he/she	FUT
<i>nakokabola</i>	divide	I	FUT
<i>okotatabana</i>	be:surprised	you:SING	FUT
<i>akotámbwisa</i>	drive	he/she	FUT
<i>nakopangwisa</i>	wipe	I	FUT
<i>nákoma</i>	write	I	SUBJ
<i>ókoma</i>	write	you:SING	
<i>ákoma</i>	write	he/she	
<i>násómba</i>	buy	I	SUBJ
<i>ókóma</i>	arrive	you:SING	SUBJ
<i>ákabola</i>	divide	he/she	SUBJ
<i>nátámbwisa</i>	drive	I	SUBJ
<i>óbóngola</i>	change	you:SING	SUBJ
<i>ásepela</i>	enjoy	he/she	SUBJ
<i>náyébisa</i>	tell	I	SUBJ
<i>nakomí</i>	write	I	PRES PERF
<i>okomí</i>	write	you:SING	
<i>akomí</i>	write	he/she	
<i>akómí</i>	arrive	he/she	PRES PERF
<i>nasómbí</i>	buy	I	PRES PERF
<i>ondimí</i>	agree	you:SING	PRES PERF
<i>obóngólí</i>	change	you:SING	PRES PERF
<i>nabóngólí</i>	change	I	PRES PERF
<i>atámbwísí</i>	drive	he/she	PRES PERF
<i>otámbwísí</i>	drive	you:SING	PRES PERF
<i>nayébísí</i>	tell	I	PRES PERF

<i>ayébísí</i>	tell	he/she	PRES PERF
<i>osepélí</i>	enjoy	you:SING	PRES PERF
<i>natatábání</i>	be:surprised	I	PRES PERF
<i>otatábání</i>	be:surprised	you:SING	PRES PERF
<i>okabólí</i>	divide	you:SING	PRES PERF
<i>akabólí</i>	divide	he/she	PRES PERF
<i>opangwísí</i>	wipe	you:SING	PRES PERF

3. What are the root morphemes with the following meanings?

- write
- arrive
- look
- agree
- speak
- buy
- change
- drive
- tell
- divide
- wipe
- be surprised

4. What are the following TAM morphemes?

- SIMPLE PRESENT
- FUTURE
- SUBJUNCTIVE
- PRESENT PERFECT

5. Given what you know about Lingala verb morphology, how would you say the following?

- look + you:PLUR + FUT
- wipe + we + SUBJ
- plead + I + PRES PERF

hint

Let's start as before looking for pairs of words that differ by only one morpheme. The first two words, *nakokoma* and *nakokóma*, differ only by the root morpheme, and their forms differ only by the tone on one syllable. But it is unlikely that the root is represented only by the tone on one syllable (since there will be many different roots). So maybe this is just a coincidence because of a similarity between these two particular roots. Further down, there three other words that differ from these only by the root, *nakotúna*, *nakokabola*, and *nakopangwisa*. Now we see that the part that distinguishes the forms is the part following *nako* and preceding the *a* at the end. So we can tentatively propose the following:

- *-kom-* 'write'
- *-kóm-* 'arrive'
- *-tún-* 'ask'
- *-kabol-* 'divide'
- *-pangwis-* 'wipe'

Now let's check different words with each of these root morphemes to see if they all share the forms that we're proposing. There are no others for 'ask', but for 'write' and 'arrive' the morphemes we proposed are found in all of the words. For 'divide'

and 'wipe', however, there are a few words where the tones differ: *okabólí*, *akabólí*, and *opangwísí*. The tones on the first syllables of the forms we proposed agree (all are low), but on the second syllable, they are sometimes low and sometimes high. Maybe the tone on the second syllable is not part of the root morpheme, but part of the TAM morpheme instead.

Let's continue to look for more roots as we have been, this time looking for all words that share a particular root. For example, all words with the meaning 'buy' have *sómb* right before the final vowel. In several cases, we again discover that roots that appear to have more than one syllable have different possible tone patterns: for 'change' we find both *bóngol* and *bóngól*, for 'drive' both *támbwis* and *támbwís*, for 'enjoy' both *sepel* and *sepél*, for 'be:surprised' both *tataban* and *tatábán*. Again the first syllables always agree, while the other syllables can be either high or low. So here are the additional roots we propose, with the understanding that the tone on any syllables after the first is not specified in the root. (We can figure out the roots that have only example by looking for "minimal pairs" or "near minimal pairs" with other roots.)

- *-ndim*- 'agree'
- *-sómb*- 'buy'
- *-sepel*- 'enjoy'
- *-bóngol*- 'change'
- *-támbwis*- 'drive'
- *-yébis*- 'tell'
- *-tataban*- 'be:surprised'
- *-tál*- 'look'
- *-lob*- 'speak'

Now let's try to figure out the TAM morphemes (the hardest part). Since we know whether the subject morphemes come, and we think we know what the root morphemes, the TAM morphemes should be whatever is left. One part that we haven't included in either the subject or root is the final vowel. Notice that this is *a* in all cases except when the TAM morpheme is PRES PERF. So we can propose that the following are at least parts of the TAM morphemes.

- *-a* 'SIMPLE PRESENT, FUTURE, SUBJUNCTIVE'
- *-í* 'PRESENT PERFECT'

Now let's look at some "minimal pairs". The SIMPLE PRESENT and SUBJUNCTIVE look very similar, so they might be good to start with. One pair is *nakoma* and *nákoma*, the first SIMPLE PRESENT, the second SUBJUNCTIVE. It looks like the SUBJUNCTIVE makes the tone on the subject prefix high. Let's look at the other SUBJUNCTIVE examples to see if this holds. Sure enough, all have high tone on the subject prefix. So now there are two ways we can say this. One is to say that the tone on the subject prefix is specified by the TAM morpheme: it is low for SIMPLE PRESENT, FUTURE, and PRESENT PERFECT and high for SUBJUNCTIVE. The other is to say that subject prefix has low tone but that this gets changed to high in the SUBJUNCTIVE (an example of mutation). Given the data here, either would be reasonable.

Now let's try to figure out the future. Comparing it with the SIMPLE PRESENT (there are several "minimal pairs"), we can conclude that the FUTURE has an additional syllable *-ko-* following the subject prefix. (This is also a prefix because it precedes the root.)

Now for the PRESENT PERFECT. So far it seems to consist of just the suffix *-í*. But if we compare the SIMPLE PRESENT/PRESENT PERFECT "minimal pairs" such as *okabola* / *okabólí*, *natatabana/natatábání*, *nabóngola/nabóngólí*, *atámbwisa/atámbwísí*, we see that the suffix is not the only difference. In fact, if we look carefully, we see that for the PRESENT PERFECT, the root syllables following the first one always have high tone, whereas these same syllables always have low tone for the other TAM morphemes. Again there are two ways to describe this. Either the tone of these syllables is part of the TAM morpheme in all cases: low for SIMPLE PRESENT, FUTURE, and SUBJUNCTIVE and high for PRESENT PERFECT. Or these syllables have low tone in the root, and PRESENT PERFECT makes the tone high. Notice that the first option is appealing because it has those root syllables agreeing in tone with the suffix, an assimilation-like process.

So here are the TAM morphemes we come up with (one of several possible descriptions).

- The tone of root syllables after the first agrees with the tone of the suffix.

- *-a* 'SIMPLE PRESENT'
- *-ko-* immediately after subject prefix, *-a* 'FUTURE'
- *-a*, high tone on subject prefix 'SUBJUNCTIVE'
- *-í* 'PRESENT PERFECT'

And now we should be able to put these together to make the novel words in 5.

- a. *bokotála* 'look + you:PLUR + FUT'
- b. *tópangwisa* 'wipe + we + SUBJ'
- c. *nabóndélí* 'plead + I + PRES PERF'

This page titled [7.4: Problems](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

8: Derivation

[8.0: Introduction](#)

[8.1: Derivational Morphology](#)

[8.2: Foregrounding and Backgrounding](#)

[8.3: Active and Passive Voice](#)

[8.4: More Verb Derivation](#)

[8.5: Problems](#)

This page titled [8: Derivation](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.0: Introduction

In the last chapter we saw how languages use grammatical morphology to create a very abstract kind of semantics, dividing the things, attributes, and situations that language is about into a small number of general categories. But grammatical morphology has another function, derivation, the creation of new words designating new concepts that are related to the meanings of existing lexical morphemes. Because this process is often generally applicable to whole categories of lexical morphemes, it is a good example of the **productivity** of language. Given a new adjective *zug* to designate some new attribute, an English speaker can create *unzug* to mean the attribute on the opposite end of some dimension from that attribute and *zugness* to mean the condition of having that attribute. As with the inflectional morphology described in the last chapter, languages also differ considerably in what possibilities they offer speakers for creating new words and new meanings using morphology. These differences lead to quite different ways of expressing similar meanings in different languages. In fact some languages may permit construals that are awkward or impossible in other languages.

One particular area of grammar where these differences are apparent is in the way the participants in events are represented in noun phrases in sentences. Many languages have productive verb morphology that allows particular participants to be foregrounded or backgrounded, giving these languages an unusual flexibility in this part of their grammar. In this chapter we'll examine this sort of flexibility in Lingala.

This page titled [8.0: Introduction](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.1: Derivational Morphology

Early on in their creation of words for the categories of things, attributes, states, and events around them, the Grammmies realized that there were often pairs of concepts that were associated by a particular abstract relation. For example, given a scalar attribute like the one designated by the word *wide*, there is the happen event involving a change of state in some object in the direction of that attribute and also the do_to event involving an agent who causes such a change of state. These three related concepts are exemplified in the following English sentences.

1. *wide road*
2. *The road widens at this point.*
3. *The workers are widening the road.*

Note the English uses, the same verb, *widen*, for both the happen and the do_to events, intransitive *widen* for the first, transitive *widen* for the second.

Further, the Grammmies saw that each of these abstract relations applied to many pairs of concepts. For example, the same relations that relate to the meaning of *wide* to the meaning of intransitive *widen* relate meaning of *dark* to the meaning of *darken* and the meaning of *dead* to the meaning of *die*.


We have seen several examples of how languages capitalize on generalizations such as this to save Speakers and Learners the trouble of learning separate words for each of the concepts they might want to refer to. For example, we saw how adjective + noun phrases avoided the use of a separate word for each combination of attribute dimension value and thing category. With the word *red*, a speaker could say *red apple*, *red pear*, *red rock*, and *red sky*, instead of learning a separate word for each of these. Because *red* can combined with any noun whose meaning is compatible with redness, the pattern *red* + noun is a productive one in English. In fact the pattern adjective + noun is a productive one, as we've seen.

The Grammmies saw that similar advantages could be gained by being systematic about how they created new words from old words on the basis of the abstract relations they'd discovered. That is, they realized that this was another place where language could make things easier for Learners and Speakers by being productive. For each of the abstract relations, they needed a morpheme to combine with a word, or a lexical root, to form the new word. This process is called **derivation**.

Derivation is common in the modern languages of the world. For example, many of them have productive ways of relating adjectives to change-of-state verbs. In English the verb is derived from the adjective by adding the suffix *-en* to the adjective. Examples are *weaken*, *shorten*, *lighten*, *blacken*, *sharpen*, *soften*, and *loosen*. Notice that this process is only somewhat productive; it doesn't apply to *long* (verb: *lengthen*), *big* (verb: *grow*), or *thin* (verb: *thin*). If we learned a new adjective, say, *zub*, we might not feel completely confident in making into a verb *zubben*. In some other languages, such as Amharic, it is usually the adjective that is derived from the verb. For example from the verb root *drk'* meaning 'get dry', the adjective *dIrk'* 'dry' is derived.

Another possibility, common in English, is for the two forms to be identical. For example, as we have seen, *widen* can describe a change in the width of something or a causing of such a change. In the last chapter, we treated these two meanings as two different syntax-semantics mappings. In any case, we have no basis for seeing one of them as derived from the other. Another example is provided by the relatively productive pattern in English by which a noun for an instrument can also be used as a verb designating the use of such an instrument. Examples are *hammer*, *saw*, *chisel*, *pin*, and *nail*. Though historically the nouns came first, again we can treat the nouns and verbs simply as related meanings of a single word.

When there is derivational morphology, a Speaker or Hearer of the language must know not only what the grammatical morpheme is and how it combines with the lexical morpheme but also the grammatical convention for how the meaning of the more complex word is derived from the meanings of the two components. Let's consider another English example, the addition of *-er* to a verb to produce an noun. The compositional convention would say something like this: the meaning of the complex word is a person who acts (routinely or one occasion) as the agent of the category of event (action) which the verb designates. Thus a *teacher* is a person who acts as the agent of a teaching event. The diagram below illustrates the relationships.



Like the derivation of verbs from adjectives (or adjectives from verbs) the agentnoun derivation of one sort or another is quite common in the world's languages. But languages also differ considerably in how much they make use of derivational morphology. Languages with rich morphology may allow a very wide range of new words to be derived from a single lexical root. In a sense these languages are making generalization that are not made in other languages, which must rely on separate, unrelated words or whole phrases to convey the different meanings. In the next I'll describe some of the possibilities for derivational morphology on Lingala verbs. Like other languages in the Bantu family, Lingala allows a number of different verbs to be derived from a single verb root.

This page titled [8.1: Derivational Morphology](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.2: Foregrounding and Backgrounding

Your page has been created!

Remove this content and add your own.

Edit page

Click the Edit page button in your user bar. You will see a suggested structure for your content. Add your content and hit Save.

Tips:

[Drag and drop](#)

Drag one or more image files from your computer and drop them onto your browser window to add them to your page.

[Classifications](#)

Tags are used to link pages to one another along common themes. Tags are also used as markers for the dynamic organization of content in the CXone Expert framework.

[Working with templates](#)

CXone Expert [templates](#) help guide and organize your documentation, making it flow easier and more uniformly. Edit existing templates or create your own.

[Visit for all help topics.](#)

This page titled [8.2: Foregrounding and Backgrounding](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.3: Active and Passive Voice

Your page has been created!

Remove this content and add your own.

Edit page

Click the Edit page button in your user bar. You will see a suggested structure for your content. Add your content and hit Save.

Tips:

[Drag and drop](#)

Drag one or more image files from your computer and drop them onto your browser window to add them to your page.

[Classifications](#)

Tags are used to link pages to one another along common themes. Tags are also used as markers for the dynamic organization of content in the CXone Expert framework.

[Working with templates](#)

CXone Expert [templates](#) help guide and organize your documentation, making it flow easier and more uniformly. Edit existing templates or create your own.

[Visit for all help topics.](#)

This page titled [8.3: Active and Passive Voice](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.4: More Verb Derivation

Your page has been created!

Remove this content and add your own.

Edit page

Click the Edit page button in your user bar. You will see a suggested structure for your content. Add your content and hit Save.

Tips:

Drag and drop

Drag one or more image files from your computer and drop them onto your browser window to add them to your page.

Classifications

Tags are used to link pages to one another along common themes. Tags are also used as markers for the dynamic organization of content in the CXone Expert framework.

Working with templates

CXone Expert [templates](#) help guide and organize your documentation, making it flow easier and more uniformly. Edit existing templates or create your own.

Visit for all help topics.

This page titled [8.4: More Verb Derivation](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.5: Problems

Your page has been created!

Remove this content and add your own.

Edit page

Click the Edit page button in your user bar. You will see a suggested structure for your content. Add your content and hit Save.

Tips:

[Drag and drop](#)

Drag one or more image files from your computer and drop them onto your browser window to add them to your page.

[Classifications](#)

Tags are used to link pages to one another along common themes. Tags are also used as markers for the dynamic organization of content in the CXone Expert framework.

[Working with templates](#)

CXone Expert [templates](#) help guide and organize your documentation, making it flow easier and more uniformly. Edit existing templates or create your own.

[Visit for all help topics.](#)

This page titled [8.5: Problems](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

9: Appendices

[9.0: Phonetic Symbols](#)

[9.1: Languages Cited in this Book](#)

[9.2: References](#)

This page titled [9: Appendices](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

9.0: Phonetic Symbols

The tables below give the phonetic symbols used in this book, which for the most part are the symbols in the International Phonetic Alphabet (IPA). Some of the symbols are alternates to IPA symbols that are commonly used to transcribe English (especially in the US) and some other languages. Where these differ from the standard IPA symbols, the IPA symbols are shown in parentheses in the tables.

Phonetic symbols are used in two ways, as described in this paragraph, to indicate the phonemes of a particular language and to indicate the relatively precise nature of a given sound. For example, in this book the symbol /e/ is used to represent one of the vowel phonemes of English. But, as we have seen, this phoneme is pronounced as a diphthong by most speakers, [eɪ] in General American. In the table below, the symbol [e] represents a "pure" vowel sound not found in General American (but found in accents of Scotland, Ireland, and northern England, where the phoneme /e/ is not a diphthong). In addition, for the sake of simplicity, we may use the same symbol to represent phonemes in different languages that differ phonetically. So in this book, when the phonetic details don't matter, I use the symbol /r/ to represent the first sound in the English word *red* (phonetically, a retroflex approximant) and the first sound in the Spanish word *rojo* (phonetically, an alveolar trill). Symbols that are used for a relatively wide variety of sounds in phonemic transcription in this book are marked with an asterisk(*) in the tables.

Recall also that the symbols are just a convenient way to *approximate* the sounds of human languages. Far more sounds than these are actually produced, so there is no way to describe all of the sounds perfectly using these symbols. I have not included all of the IPA symbols, only those for sounds that have been discussed in the book. Many of the feature values are also meant only as approximations. For example, the height of both [i] and [ɪ] is given as "high", but [ɪ] is actually slightly lower than [i], and the rounding of both [i] and [e] is given as "spread", but [i] is actually more spread than [e].

Vowels

Symbol	Feature			
	Height	Backness	Roundness	Tenseness
*i	high	front	spread	tense
ü(y)	high	front	round	tense
ɪ	high	front	spread	lax
*e	high mid	front	spread	tense
ɛ	low mid	front	spread	lax
æ	low	front	spread	lax
*u	high	back	round	tense
ʊ	high	back	spread	tense
ʊ	high	back	round	lax
*o	high mid	back	round	tense
ɔ	low mid	back	round	lax
ɑ	low	back	spread	lax
ɨ	high	center	spread	lax
ʌ	mid	center	spread	lax
ə	mid	center	spread	lax
*a	low	center	spread	lax

Consonants

Symbol	Feature		
	POA	MOA	Voice
p	bilabial	stop	vcls
b	bilabial	stop	voiced
m	bilabial	nasal	voiced
ɸ	bilabial	fricative	voiceless
β	bilabial	fricative	voiced
w	labiovelar	approximant	voiced
f	labiodental	fricative	vcls
v	labiodental	fricative	voiced
θ	dental	fricative	vcls
ð	dental	fricative	voiced
*t	alveolar	stop	vcls
*d	alveolar	stop	voiced
n	alveolar	nasal	voiced
s	alveolar	fricative	vcls
z	alveolar	fricative	voiced
*r	alveolar	trill	voiced
r	alveolar	tap	voiced
l	alveolar	lateral sonorant	voiced
ɭ	retroflex	approximant	voiced
č(ʧ)	postalveolar	affricate	vcls
ǰ(ɟʒ)	postalveolar	affricate	voiced
š(ʃ)	postalveolar	fricative	vcls
ž(ʒ)	postalveolar	fricative	voiced
ɲ	palatal	nasal	voiced
y(j)	palatal	approximant	voiced
k	velar	stop	vcls
ŋ	velar	nasal	voiced
g	velar	stop	voiced
x	velar	fricative	vcls
ɣ	velar	fricative	voiced
q	uvular	stop	vcls

Symbol		Feature	
ʔ	glottal	stop	vcls
h	glottal	fricative	vcls

These basic symbols are modified to indicate details, including additional features that are not simply accommodated by the feature dimensions given in the tables. The modifications used in this book are listed in the following table.

Symbol	Example	Meaning
ʔ	tʔ	glottalized, ejective
h	t ^h	aspirated
ɲ	ɲ	dental (as opposed to alveolar)
:	e:	long

This page titled [9.0: Phonetic Symbols](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

9.1: Languages Cited in this Book

In order to understand what is possible in human language, it is important to look at examples from languages that are quite different from your first language. In this book most of the examples will come from nine languages, selected to some extent because they are important in the modern world (English, Spanish, Chinese, Japanese) but mainly because they illustrate a wide range of possibilities. The word (or sign) for 'language' in each of these nine languages appears on the left side of the banner at the top of the Table of Contents page of this book. In the banners at the top of other pages, such as this one, you will see these examples as well as others from nine other languages (Hindi, Hungarian, Navajo, Persian, Russian, Samoan, Tamil, Thai, and Yoruba). As you see examples from these nine languages later on in the book, some of them may seem quite exotic to you, but keep in mind that your language probably also seems equally exotic to speakers of those languages. Each language is introduced briefly in this section. Some of the terms used to describe the languages will probably not be familiar to you if you have not read the relevant sections in the book.

Mandarin Chinese

In terms of both number of native speakers (about 875,000,000) and number of first and second language speakers (about 1,050,000,000), Mandarin Chinese is easily the most widely spoken language in the world. The great majority of these speakers live in the People's Republic of China (mainland China), but many others also live in the Republic of China (Taiwan) as well as other countries where Chinese speakers have settled. Mandarin Chinese is the official language of both mainland China and Taiwan. The word for 'Chinese' in Mandarin, transcribed into Roman using the *pinyin* system, is *zhōngguóhuà* (the marks over the vowel symbols indicate tones).

Mandarin Chinese is one of about 10 languages or dialects that are called "Chinese". Calling them "dialects" emphasizes that they are all spoken in China and that they all have a single written standard (closest to Mandarin). Calling them "languages" emphasizes the large differences in pronunciation, vocabulary, and grammar among them; I will follow this usage here. The Chinese languages belong to the Sino-Tibetan language family, which also includes Burmese, Tibetan, and about several hundred other languages spoken in China, Southeast Asia, and South Asia.

Like other Sino-Tibetan languages, Mandarin Chinese is a tone language; it uses pitch to distinguish words. In Mandarin, each syllable has one of four tones: high level, mid rising, low falling-rising, and high falling. Because Mandarin and other Chinese languages have so many one-syllable words and because the number of possible syllables is relatively small, many words are distinguished only by their tone. Mandarin syllable structure is relatively simple. The only consonants syllables can end in are /n/ and /ŋ/, and there are no consonant clusters at the beginnings of syllables. But Chinese pronunciation is still complicated for second-language learners because of the tones and because there are two separate sets of consonants similar to the English postalveolar consonants, that is, the final consonants in the words *reach*, *ridge*, and *rash*.

Mandarin and other Chinese languages have little or no bound morphology. Each word has a fixed form: verbs do not take prefixes or suffixes showing the tense or the person, number, or gender of the subject. Nouns do not take prefixes or suffixes showing their number or their case. This does not mean that Chinese has no grammar of course; what it means is that Chinese grammar is mainly concerned with how words are arranged to form meaningful sentences. For the second language learner, it also means that there are no grammatical paradigms to memorize as there are for languages such as Spanish and Amharic.

Chinese and its ancestor languages have been written for over 4000 years, giving them the longest continuous writing tradition of any group of languages in the world. The Chinese writing system uses a complex system of characters that is quite unlike any other modern writing system. Originally all of the characters were simplifications of pictures designed to represent the meanings of words directly. Later, characters were created that have parts representing elements of meaning or sound. Each character in the modern system represents a single meaningful syllable. To read Chinese, you need to know several thousand of these. If you have Chinese characters on your computer, you can see the Universal Declaration of Human Rights in Chinese [here](#).

English

In terms of number of native speakers (about 340,000,000), English is among the four most widely spoken languages in the world (English, Spanish, and Hindi are very close in number of native speakers). About 510,000,000 people speak it as a first or second language (second in the world). In terms of its role as a language of international communication, science, and business, English is

probably the most important language in the world. Most first-language speakers live in the United States, the United Kingdom, Canada, Australia, Ireland, New Zealand, and South Africa, but English is an official language of many other countries, especially in Africa and the Caribbean.

English is a member of the Germanic group (like German, Dutch, Swedish, Norwegian, Danish, and Icelandic) of the Indo-European language family. A good deal is known about the history of the language. The earliest period of the language, up to about 1100, is called **Old English**. Old English is so different from Modern English that we find it almost completely uninterpretable; it is best to treat it as a different language from Modern English. A number of dramatic changes in vocabulary and grammar were initiated by the conquest of England by French-speaking Normans in 1066. The period known as **Middle English** begins soon after that and continues roughly till the year 1500, as English long vowels were undergoing the changes known as the Great Vowel Shift. **Modern English** is what we speak now.

English has less bound morphology, that is, suffixes and prefixes on words, and more rigid word order than most other Indo-European languages. English has a large number of vowels, about twelve, depending on the dialect, and the vowel system is relatively unstable, changing frequently in the history of the language. The complex phonology makes English a relatively difficult language for adult second-language learners to learn to pronounce.

Each of the major countries in which English is spoken has its own (unofficial) standard dialect, though the standard dialects of the United States and Canada are very similar, and the standard dialects of Australia and New Zealand are also very similar. England and the United States also have a number of regional or social dialects which differ considerably from the standards. Examples in the United States include Southern (spoken in most of the Southeast), Northern Cities (spoken in cities such as Chicago, Detroit, and Milwaukee), New York, and African-American Vernacular English (AAVE). American English dialects differ from each other mainly in pronunciation, but the grammar of some varieties of AAVE is quite distinct.

Spanish

In terms of number of native speakers (about 350,000,000), Spanish is among the four most widely spoken languages in the world. About 420,000,000 people speak it as a first or second language (fourth in the world). Spanish is the official language of 21 countries, including Spain and 20 countries in the Western Hemisphere. Mexico is the country with the most Spanish speakers. Many first-language speakers of American Indian languages speak Spanish as a second language. In Spanish the language is called *español* or *castellano*.

Spanish is a member of the Romance group (like French, Italian, Portuguese, Catalan, and Romanian) of the Indo-European family. The Romance languages all descended from a language we call Vulgar Latin today. This was not the written Latin language you may have studied in school but the related language spoken by ordinary Romans at the beginning of the first millenium, including the soldiers of the Roman Empire who spread their language to many parts of Europe. All of the Romance languages have more bound morphology, that is, prefixes and suffixes, than English, especially on verbs. Spanish verbs indicate tense as well as the subjects of the verbs with suffixes. Like most other Indo-European languages, Spanish has grammatical gender: nouns are either masculine or feminine, and words that modify them must agree with them in gender. The dominant word order of Spanish is subject-verb-object, but Spanish is much more flexible than English in this regard, and it is not unusual for the verb to appear before the subject.

Spanish has a simpler phonological system than the other Romance languages, including only five vowels. This makes Spanish a relatively easy language for adult second-language learners to learn to pronounce.

The standard dialects in Spain and in the Americas differ in a number of ways, in much the same way as the standard English of England and the United States differs. Within Spain and within the Americas, there are also many regional dialects; some of these, such as Mexican and Argentine Spanish, are quite distinct and are treated as (informal) national standards.

Spanish is written using the Roman alphabet, included accented vowel characters to mark stress and distinguish some words from one another. Here is the Universal Declaration of Human Rights in Spanish.

Japanese

In terms of numbers of native speakers (about 125,000,000), Japanese is about the eighth most widely spoken language in the world. The great majority of the speakers live in Japan, where almost everyone speaks Japanese. In Japanese the language is called

nihongo.

There is disagreement on which languages Japanese is related to. It is only clear that its closest relatives are the Ryukyuan languages spoken in the very south of Japan. Some believe Japanese is related to Korean; some believe it and Korean are related to the Altaic language family (Turkish, Mongolian, etc.).

Japanese has extensive bound morphology on its verbs, almost all of its suffixes, though, unlike in Indo-European languages, the subject is not marked on the verb. Even though verbs do not agree with their subjects, subjects are often omitted when it is possible for the hearer to infer what they are from the context. Similarly the objects of transitive verbs may be omitted. Japanese is a verb-final language and shares many syntactic properties with other verb-final languages. Japanese is also known for the many devices for expressing level of formality in the language (both vocabulary and grammar) and for the differences between men's and women's speech.

Japanese has a relatively small set of phonemes, so second-language learners usually find it relatively easy to learn to pronounce, but it makes use of pitch to distinguish words from one another, and this system is often not mastered, or even attempted, by second-language learners.

Alongside Standard Japanese, which all children learn in school, there are many regional dialects, differing very greatly from one another (much more than the dialects of American English, for example).

Japanese is written using a combination of Chinese characters (*kanji*) and two related syllabic systems (*hiragana* and *katakana*). Despite this complexity, the Japanese are among the most literate people in the world. Here is the Universal Declaration of Human Rights in Japanese.

Amharic

In terms of numbers of native speakers (about 20,000,000), Amharic is approximately the fiftieth most widely spoken language in the world. Almost all of the native speakers live in Ethiopia, where they make up about 1/3 of the population. Because Amharic is the official language of Ethiopia, taught in all schools, it is also spoken as a second language by many other Ethiopians. In Amharic the language is called *amarinya*.

Amharic has seven vowels, and its consonants include a set of ejectives, unlike anything in a language such as English. For example, the consonants /t/, /d/, and /t'/ need to be distinguished. This makes it a somewhat difficult language for second-language speakers to pronounce.

Amharic belongs to the Semitic group (like Arabic and Hebrew) within the Afro-Asiatic family of languages. About 50 other Semitic and non-Semitic Afro-Asiatic languages are spoken in Ethiopia alongside Amharic.

Like other Semitic languages, Amharic has a very elaborate verb morphology. An Amharic verb root consists of a set of (usually three) consonants. Depending on the tense, and other grammatical features, the consonants may be separated by particular vowels and possibly geminated (doubled). A verb form normally also has one or more suffixes and possibly one or more prefixes as well, agreeing with the subject and sometimes the direct or indirect object of the verb. Complicating things further (at least for the adult second-language learner), there are at least ten different classes of verbs, each modifying its stem in a different way for the different forms. Like Japanese and many other languages, Amharic is a verb-final language. Amharic nouns are relatively simple by comparison, though they may take suffixes indicating possession ('my', 'his', etc.), plural, and a few other grammatical functions.

Unlike most African languages, Amharic has been a written language for many years, at least 500. It is written using a syllabic writing system that is unique to Ethiopian Semitic languages. Compared to other African languages, Amharic has a fairly sizable written literature. Here is the Universal Declaration of Human Rights in Amharic.

Lingala

Lingala is the first language of about 300,000 people in the two neighboring countries called Congo in central Africa, but is a second language for many more, at least 7,000,000 people, making it one of the most important languages in central Africa. In Lingala, the name of the language is *lingála* (the accent mark indicates a high tone on that syllable).

Lingala belongs to the Bantu group of languages, like most of the languages of central and southern Africa. Other important Bantu languages include Swahili, Zulu, and Shona. The Bantu languages belong in turn to the Niger-Kordofanian language family, which

includes most of the more than 2000 languages spoken in Africa. Lingala does not have a long history; it is an example of a **contact language**, a language which emerges out of the contact between people speaking different languages. Lingala began along the Congo River, first as a trade language, later, when Congo was colonized by Belgium, as a way for the colonists and their intermediaries to communicate with the local people. As usually happens in such situations, the resulting language borrowed words and grammatical features from a number of existing languages, in this case Bantu languages of the region. At first it was not the first language of anybody, but with the growth of cities and marriage between speakers of different first languages, it soon became a language learned by children as a first language. In recent years, the language has spread rapidly, partly because of its role in government, partly because it is the usual language in which much of the enormously popular music of the region is sung, especially the style of music known as *soukous*.

Lingala has relatively few phonemes and simple CV syllable structure. Like most other Bantu languages, it is a tone language, using pitch both to distinguish words from each other and to mark grammatical categories.

When a contact language emerges, the grammar (at least the morphology) of the source languages is normally simplified. This is true for Lingala too, though it preserves many of the features of Bantu morphology. Like other Bantu languages, Lingala divides nouns into about 14 classes, similar to the genders of European languages, each with its own characteristic prefix. Verbs have prefixes, suffixes, and particular tone patterns, to indicate the subject and the tense, and to derive a range of other forms from each verb. The basic word order is subject-verb-object.

Lingala is written in the Roman alphabet, but there is not yet an extensive literature. Most educated Congolese still use French for writing. Here is the Universal Declaration of Human Rights in Lingala.

Tzeltal

Tzeltal is spoken by about 215,000 people in Chiapas state in Mexico. Many of the people are bilingual in Spanish. Unlike many other American Indian languages, Tzeltal seems to have a chance of surviving, but this will depend on Mexican language policies, which in the past have not favored languages other than Spanish.

Tzeltal is a member of the Mayan family of languages, currently spoken by several million people in southern Mexico, Guatemala, and Belize. These people are the descendants of the the great Maya civilization that flourished in this region 1000 years ago.

Like Amharic, Tzeltal phonology includes a set of ejective consonants. Tzeltal syllables are relatively simple, except that clusters like /hp/ can occur at the beginnings and ends of words.

Tzeltal is known for its elaborate system of classifiers, morphemes which must follow numerals when objects or actions are counted. The choice of classifier usually depends on the shape or structure of the counted objects. One researcher has counted more than 500 classifiers in one Tzeltal dialect. Tzeltal is also known for the way it expresses spatial relations ('on', 'under', 'left of', etc.). It makes use of body part terms ('head', 'back', 'rump', etc.) and a set of spatial orientation verbs for close relations and absolute frame of reference (roughly 'east', 'west', etc.) for more distant relations. The most common word order is verb-object-subject.

The dialects of Tzeltal differ considerably from one another, and there is no commonly accepted standard dialect.

Ancient Mayan was written using a hieroglyphic system, but Tzeltal had no writing system until very recently. It is written today using the Roman alphabet. There are few materials written in it, however; Spanish is the usual written language for literate Tzeltals. Here is the Universal Declaration of Human Rights in Tzeltal.

Inuktitut

Inuktitut is spoken by about 14,000 people in northeastern Canada, mainly in the newly formed territory of Nunavut, for which it is one of the official languages. Inuktitut is one of about eight languages in the Eskimo group of languages within the Eskimo-Aleut family. Eskimo languages are spoken in a large, sparsely populated region extending from the far eastern end of Siberia in Russia, across Alaska and northern Canada to Greenland. Unlike most other indigenous languages of North America, Inuktitut seems to have a good chance of surviving.

Inuktitut has a relatively simple set of phonemes. For a given place of articulation, for example, velar, the language makes no distinction between voiceless (English /k/) and voiced (English /g/) stops. However, Inuktitut has a place of articulation not used at

all in English but also used in Tzeltal and other Mayan languages: **uvular** consonants are produced with the tongue making contact or approaching the top of the mouth even further back than for velar consonants.

Inuktitut and other Eskimo languages have some of the most elaborate morphology of any languages in the world. A single Inuktitut word often corresponds to an entire sentence in English. Here is an example of a word consisting of eight morphemes: *Pariliarumaniralausimanngittunga* 'I never said I wanted to go to Paris'. While Inuktitut does not distinguish gender anywhere in its grammar, it distinguishes three numbers, singular, dual, and plural. Each noun can have scores of forms because it must be marked for number and for case. Most verbs can have hundreds of forms because they must be marked for their subject and their direct object, if there is one, as well as for their tense. In some cases, verbs have a special set of suffixes when a question is being asked. In addition to the three persons that other languages of the world distinguish, Inuktitut has a "fourth person" for cases when two third person clauses within the same sentence have a different subject.

Inuktitut is usually written using a modification of a script that was originally designed by missionaries for the Cree language. In this writing system, each simple consonant-vowel syllable is represented by a single symbol. Here is the Universal Declaration of Human Rights in Inuktitut written in the syllabic script.

American Sign Language

American Sign Language (ASL) is the primary language of between 100,000 and 500,000 people out of the approximately 2,000,000 deaf people in the United States and Canada, but many other deaf people, and some hearing people, have some competence in the language.

ASL is one of hundreds of deaf sign languages in the world, none of them related to spoken languages, but some recent research has shown that sign languages may share some properties with one another even when they are not genetically related. British Sign Language is a separate language, not mutually intelligible with ASL. ASL developed in the schools for the deaf in the US begun in the early part of the 19th century. Because the teaching in these schools was influenced by early French methods of educating the deaf, many ASL signs derive originally from French Sign Language. However, ASL and French Sign Language are quite different languages today.

ASL is to be distinguished from Signed English, which uses English word order, replacing the English words with signs. The grammar of ASL is very different from that of English. Like Tzeltal and Japanese, but not English or Spanish, ASL makes use of classifiers, morphemes that are added to signs to agree with the shape or other physical property of some object. For example, expressions showing location ('be at') take a handshape that conveys something about the thing that is in the location (that is, whether it is a vehicle, an animal, etc.). As in other sign languages, there is a strong tendency for signs in ASL to be iconic, that is, to be motivated by their meanings rather than completely arbitrary. This iconicity extends into the grammar. For example, in ASL participants in a discourse are usually assigned places in the signing space and then pointed to later on when they are referred to, that is, where pronouns would appear in a spoken language.

There is currently no agreed on way of writing ASL, but ASL does have its own literature, including poetry.

This page titled [9.1: Languages Cited in this Book](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

9.2: References

Linguistic Society of America

Linguist List, mailing list and other resources for linguists

Ethnologue, database of all of the world's known languages (6,809!)

American Dialect Society, the academic association concerned with the study of the English language in North America

American dialect links: links to sites concerned with American English dialects (Note: Many of the sites are not at all scholarly, but still interesting sources of data. And some links are dead.)

Two sites concerned with endangered languages: Terralingua and the Foundation for Endangered Languages

ASL Browser, an online American Sign Language dictionary, complete with movies of the signs

Beginning Guide to ASL: links to lots of information about American Sign Language

Voice of America Language Sites, news from the VOA in about 50 different languages

Downloadable phonetic fonts from the Summer Institute of Linguistics

A paper by the famous sociolinguist William Labov on the dramatic changes going on in the vowels of many Americans

"Committing to an ontology", a paper on the shape bias and related issues by Eliana Colunga and Linda B. Smith of Indiana University.

Estuary English: very useful and informative site by a phonetician and expert on English accents about Estuary English, which some people believe is becoming the new standard accent of England

This page titled [9.2: References](#) is shared under a [GNU General Public License 3.0](#) license and was authored, remixed, and/or curated by [Michael Gasser](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

Index

A

Anticipatory Assimilation (Linguistics)

[4.2: Assimilation](#)

Assimilation (Linguistics)

[4.2: Assimilation](#)

C

constituents

[6.3: Constituency and Noun Phrases](#)

D

Deixis

[2.5: Deixis and Person](#)

M

metaphor

[2.4: Metaphor and Metonymy](#)

Metonymy

[2.4: Metaphor and Metonymy](#)

N

noun phrases

[6.3: Constituency and Noun Phrases](#)

P

Perseverative Assimilation (Linguistics)

[4.2: Assimilation](#)

phonemes

[3.1: Phonemes](#)

phonotactics

[3.7: Syllables](#)

R

retroflex

[3.5: English Consonants](#)

S

suprasegmentals

[3.7: Syllables](#)

syllables

[3.7: Syllables](#)

W

Word Forms (Linguistics)

[4: Word Forms - Processes](#)

Word Senses

[2.3: Word Senses and Taxonomies](#)

Glossary

Accent: the conventions of pronunciation associated with some speech community.

Appropriate: not only grammatical and meaningful but suitable for the context in which it occurs. Inappropriate utterances may be too casual, formal, direct, indirect; uninterpretable by the hearer; or in violation of the conventions of conversation,

Category: a grouping of individual things or situations by people on the basis of similarity. Individuals **belong** to categories, but a category should not be seen as the set of all of its individual members, or instances, but rather the abstract concept that they share. Categories are often thought of as having prototypes, that is, typical members; individuals can then be relatively "good" or "bad" members of a category based on how far they are from the prototype. Categories may also result in categorical perception.

Common noun: a word which labels a category of **thing** and shares morphological or syntactic properties with other such words. In English common nouns include words such as *racket*, *cheese*, and *road*. All English common nouns share several morphological and syntactic properties; for example, all can be modified by the article *the*.

Concept: a unit of cognitive experience, a way people have of abstracting over their experiences in the world.

Conceptual relation: a relation between two concepts that is based on some sort of similarity or association between them. Conceptual relations include specialization-generalization, part-whole, and physical similarity.

Constraint: a limitation on what is possible. For language, constraints come from the body, from cognition, and from the nature of human communication.

Convention: a behavior that is accepted by some group of people as appropriate for performing some function. For example, the community of English speakers agrees on the convention that the word *horse* refers to a particular kind of animal. Conventions are usually at least somewhat arbitrary; that is, they are in effect because they have been agreed on, not because they "make sense" in any way. For example, English speakers could just as well have settled on a totally different word form to refer to the category horse.

Dialect: the linguistic conventions shared by a speech community.

Deixis: use of language whose meaning depends at least in part on some element (role) of the utterance or discourse context. Deictic expressions in English include pronouns like *I* and *you*, adverbs such as *here* and *now*, and verbs such as *come* and *go*.

Dimension: a scale along which concepts can vary. Each individual or category for which the dimension is relevant has a particular value on the dimension. For example, color is a dimension which is relevant for physical objects and masses, and some possible color values are red, green, and blue. However, color is not a relevant dimension for actions; it would be unspecified for an instance of eating, say.

First person. See person.

Form: how a linguistic expression sounds (for spoken language) or looks (for written or signed language) or how it is produced.

Formality: a semantic dimension in some languages, distinguishing situations or relationships between utterance participants on the basis of the familiarity of the participants with one another, the social distance between the participants, the relative status of the participants, and the cultural significance of the interaction.

Function: (1) a use that people put language to, for example, asserting, commanding, or getting information. (2) the use that a particular word or pattern has within a stretch of language, for example, marking tense or marking the gender and number of the subject.

Gender: a grammatical dimension which divides nouns into a small set of classes (also called "genders"). The members of the classes normally tend to share some semantic properties. However, membership in one of the classes is often somewhat arbitrary and may be based on characteristics of the forms of nouns rather than their meanings. Therefore, gender is not usually considered to be a semantic dimension. In many languages, for example, Spanish, there is a correlation between biological gender and membership in the classes, which are called "feminine" and "masculine" for this reason. However, there seems to be little semantic motivation for the assignment of inanimate objects to the genders.

Generalize: to go beyond the available data, applying a word or a rule to a novel situation.

Generalization-specialization: the conceptual relation between a super-category and a sub-category. All of the properties of the super-category are shared by members of the sub-category, but the sub-category has properties not shared by all members of the super-category. For example, apple specializes fruit (or fruit generalizes apple). Different kinds of fruit are edible and sweet and contain seeds; these are all properties of apples. But apples have a core surrounding the seeds, which is not true for some kinds of fruit, such as strawberries.

Grammatical: describing a phrase or sentence which could be produced by a speaker of a given dialect or language and does not contain a speech error.

Hearer: the receiver of a linguistic act, the one who is expected to perceive and interpret the act, often referred to as the "addressee". For spoken language, a hearer; for signed language, an observer; for written language, a reader. (The word will be capitalized in this book to remind you that it's more general than "hearer".)

Individual: a particular thing or situation, as opposed to a set of things or situations, or a category of things or situations. The noun phrases *Clark*, *the linguistics class I'm taking*, and *the milk in that glass* refer to individual things. The sentence *there's a mouse under the bed* describes an individual situation.

Innate: genetically specified, available to the person or animal more or less automatically because it belongs to a particular species, not because of experience and learning. Thus in this sense we might call the liver innate; unless something is drastically wrong in our environment, we all get one. One view of language learning sees many of the properties of language as innate.

Instance: an individual belonging to some category.

Language: a set of dialects, grouped together usually on the basis of a relatively high degree of mutual intelligibility but also possibly because the speakers are members of a single political unit.

Language death: a situation arising in the history of a language when children stop learning it as a first language. Although it is possible to preserve recordings and linguistic descriptions of such a language, these can only be seen as a very incomplete record of the living language, which is now effectively irretrievable.

Language family: a group of languages that are genetically related, that is, having a single common ancestor language.

Lexical gap: a category without a word in a language, dialect, or idiolect in a domain in which most other categories have associated words. For example, most English speakers have no word for sibling, though they do have a word for parent.

Lexicon: the mental dictionary of a person, "listing" all of the words the person knows, their meanings, forms, and grammatical properties.

Mass: a thing with relatively stable properties but no clearly definable boundary, that is, not an object. Examples are some smoke, some sand, some oatmeal. Some things may be difficult to classify as being objects or masses, for example, a lump of clay or a piece of cheese.

Meaning: the concepts that linguistic words and grammatical patterns are about.

Metaphoric extension: extension of the meaning of a word on the basis of the similarity of the original meaning to the new meaning. For example, *wing* came to be used for a structure on an airplane by analogy with the similar structure on a bird.

Metonymic extension: extension of the meaning of a word on the basis of a strong association between the original meaning and the new meaning; in some sense the two "go together". For example, *wing* came to be used for a structure on an airplane by analogy with the similar structure on a bird.

Mutual intelligibility: the extent to which speakers of two or more different dialects or languages can understand each other.

Negative evidence: in learning, information about when a particular form does *not* apply. For example, in word learning, negative evidence could take the form of correcting a child's mistaken use of a word: "No, that's not a chicken; it's a duck." To be contrasted with positive evidence.

Number: a semantic dimension which distinguishes individual objects (**singular**) multiple objects (**plural**) and in some languages from two or three objects as well.

Object: something that is perceived as having a clear boundary around it and a set of stable properties such as size, color, shape, and consistency, for example, an apple, a bridge, or a snake.

Over-generalization, under-generalization: errors made by first or second language learners in which the meaning of a word in the target language is either generalized (for example, *dog* for sheep as well as dogs) or specialized (for example, *dog* only for terriers).

Personal pronoun: a word referring to an object in terms of whether it is (or includes) the Speaker or the Hearer.

Person: a grammatical category expressing whether a referent is (or includes) the Speaker (**first person**), the Hearer (**second person**), or neither the Speaker nor the Hearer (**third person**).

Plural. See number.

Positive evidence: in learning, information about when a particular form applies. For example, in word learning, positive evidence could take the form of the presentation of a referent together with an appropriate word. To be contrasted with negative evidence.

Prescription: specifying what sorts of behavior are desirable, acceptable, preferred. To be contrasted with description. Linguists are concerned with describing, not prescribing, language.

Proper noun: a word used to refer to an individual object, for example, *Clark, Indiana*. One kind of name. (Other names includes expressions like *the Grand Canyon* and *the White House*.)

Prototype: the central or typical member of a category; the basis on which individuals are evaluated as belonging to the category. One view of categories is that they *are* prototypes.

Regularity: a pattern that tends to recur in some data. For example, the co-occurrence of the word *apple* with an object with a particular shape, consistency, and taste is one sort of regularity, and the general tendency for the meanings of nouns referring to solid objects to be organized by shape is another sort of regularity.

Role: a concept defined with respect to the part it plays in a larger (surrounding, defining) concept. For example, the speaker role within an utterance, the pitcher role within a baseball team, the nose role within a face, the subject role within a sentence.

Second person. See person.

Semantic category: a category that is the meaning of a word.

Semantic dimension: a dimension which distinguishes things or situations from one another and which corresponds to distinctions in linguistic form. Examples are person, number, mass-count, and tense.

Semantic extension: addition of new sense to a word on the basis of some conceptual relation. For example, the noun *mouse* came to be used for a computer pointing device because of the physical similarity between the device and the animal.

Shape bias: the tendency for language to assume that the categories labeled by words for solid objects are organized by shape, that is, that shape matters more for these categories than color or texture.

Singular. See number.

Systematicity: the tendency for languages to be consistent in the dimensions that matter, the number of values along dimensions, and the applicability of particular rules or patterns to a variety of words. For example, if a particular dimension is relevant for consonants in one place of articulation, it is likely to be relevant in other places of articulation, and the possibility of referring to a transitive event with a passive construction tends to be applicable to all transitive verbs and to all sorts of arguments of those verbs.

Speaker: the producer of a linguistic act. For spoken language, a speaker; for signed language, a signer; for written language, a writer. (The word will be capitalized in this book to remind you that it's more general than "speaker".)

Speech error: a deviation from the conventions of a dialect or language caused by interference between words, a temporary memory lapse (perhaps due to fatigue or stress), or an ill-formed speech plan.

Standard dialect: within a language, the dialect that is accepted as being the most prestigious, often the only dialect that is written and the dialect that is used in education and the media.

Standardization: the process of settling on a standard, a set of conventions to be followed by everyone in a particular community.

Taxonomy: a set of concepts related to one another in a hierarchy by the generalization-specialization relation.

Thing: a set of co-occurring and relatively stable features, typically an **object** but also possibly a **mass**.

Third person. See person.

Utterance: an instance of language; a word, phrase, or sentence produced by somebody at a given time and place with somebody as Hearer

Utterance context: the Speaker, Hearer, location, and time associated with a particular instance of language.

Value. See dimension.

Word sense: one of the related meanings of a word. For example, the word *sponge* is used to refer both to a kind of sea animal and to an object or material used for cleaning, originally made from the internal skeleton of one of these animals.

Index

A

Anticipatory Assimilation (Linguistics)

[4.2: Assimilation](#)

Assimilation (Linguistics)

[4.2: Assimilation](#)

C

constituents

[6.3: Constituency and Noun Phrases](#)

D

Deixis

[2.5: Deixis and Person](#)

M

metaphor

[2.4: Metaphor and Metonymy](#)

Metonymy

[2.4: Metaphor and Metonymy](#)

N

noun phrases

[6.3: Constituency and Noun Phrases](#)

P

Perseverative Assimilation (Linguistics)

[4.2: Assimilation](#)

phonemes

[3.1: Phonemes](#)

phonotactics

[3.7: Syllables](#)

R

retroflex

[3.5: English Consonants](#)

S

suprasegmentals

[3.7: Syllables](#)

syllables

[3.7: Syllables](#)

W

Word Forms (Linguistics)

[4: Word Forms - Processes](#)

Word Senses

[2.3: Word Senses and Taxonomies](#)

Glossary

Sample Word 1 | Sample Definition 1

Detailed Licensing

Overview

Title: [How Language Works \(Gasser\)](#)

Webpages: 84

All licenses found:

- [GPL](#): 88.1% (74 pages)
- [Undeclared](#): 11.9% (10 pages)

By Page

- [How Language Works \(Gasser\) - GPL](#)
 - [Front Matter - Undeclared](#)
 - [TitlePage - Undeclared](#)
 - [InfoPage - Undeclared](#)
 - [Table of Contents - Undeclared](#)
 - [Licensing - Undeclared](#)
 - [1: Introduction - GPL](#)
 - [1.0: Introduction - GPL](#)
 - [1.1: What We Study - GPL](#)
 - [1.2: How We Study Language - GPL](#)
 - [1.3: What We Don't Do- Prescribing and Evaluating Language - GPL](#)
 - [1.4: Dialects and Languages - GPL](#)
 - [1.5: A Stand and Two Themes - GPL](#)
 - [1.6: Why Study Language - GPL](#)
 - [1.7: Problems - GPL](#)
 - [2: Word Meanings - GPL](#)
 - [2.0: Introduction to Word Meanings - GPL](#)
 - [2.1: Reference and Proper Nouns - GPL](#)
 - [2.2: Categories and Common Nouns - GPL](#)
 - [2.3: Word Senses and Taxonomies - GPL](#)
 - [2.4: Metaphor and Metonymy - GPL](#)
 - [2.5: Deixis and Person - GPL](#)
 - [2.6: Lexical Differences Among Languages - GPL](#)
 - [2.7: Learning Meaning - GPL](#)
 - [2.8: Problems - GPL](#)
 - [3: Word Forms - Units - GPL](#)
 - [3.0: Introduction - GPL](#)
 - [3.1: Phonemes - GPL](#)
 - [3.2: Iconicity - GPL](#)
 - [3.3: Vowels \(Part 1\) - GPL](#)
 - [3.4: Vowels \(Part 2\) - GPL](#)
 - [3.5: English Consonants - GPL](#)
 - [3.6: Consonants in Other Languages - GPL](#)
 - [3.7: Syllables - GPL](#)
 - [3.8: Problems - GPL](#)
 - [4: Word Forms - Processes - GPL](#)
 - [4.0: Introduction to Word Forms - GPL](#)
 - [4.1: Phonetic Contexts - GPL](#)
 - [4.2: Assimilation - GPL](#)
 - [4.3: Distribution of Phones - GPL](#)
 - [4.4: Learning Phonology - GPL](#)
 - [4.5: English Accents - GPL](#)
 - [4.6: Phonological Change - GPL](#)
 - [4.7: Phonology in the Wild - GPL](#)
 - [4.8: Problems - GPL](#)
 - [5: Composition - Combining Words - GPL](#)
 - [5.0: Introduction - GPL](#)
 - [5.1: Attributes and Attribution - GPL](#)
 - [5.2: Modification - GPL](#)
 - [5.3: Compositionality and Idiomaticity - GPL](#)
 - [5.4: Problems - GPL](#)
 - [6: Sentences - GPL](#)
 - [6.0: Introduction - GPL](#)
 - [6.1: States and Events - GPL](#)
 - [6.2: Situation Schemas and Semantic Roles - GPL](#)
 - [6.3: Constituency and Noun Phrases - GPL](#)
 - [6.4: Subjects - GPL](#)
 - [6.5: Direct Objects - GPL](#)
 - [6.6: Adjuncts - GPL](#)
 - [6.7: Sentence Functions - GPL](#)
 - [6.8: Problems - GPL](#)
 - [7: Grammatical Categories - GPL](#)
 - [7.0: Introduction - GPL](#)
 - [7.1: Morphemes - GPL](#)
 - [7.2: Grammatical Categories and NPs - GPL](#)
 - [7.3: Grammatical Categories and Verbs - GPL](#)
 - [7.4: Problems - GPL](#)
 - [8: Derivation - GPL](#)
 - [8.0: Introduction - GPL](#)
 - [8.1: Derivational Morphology - GPL](#)
 - [8.2: Foregrounding and Backgrounding - GPL](#)

- 8.3: Active and Passive Voice - *GPL*
- 8.4: More Verb Derivation - *GPL*
- 8.5: Problems - *GPL*
- 9: Appendices - *GPL*
 - 9.0: Phonetic Symbols - *GPL*
 - 9.1: Languages Cited in this Book - *GPL*
 - 9.2: References - *GPL*
- Back Matter - *Undeclared*
 - Index - *Undeclared*
 - Glossary - *GPL*
 - Index - *Undeclared*
 - Glossary - *Undeclared*
 - Detailed Licensing - *Undeclared*