FILE

4.2

The Value of Sounds: Phonemes

4.2.1 Allophones and Phonemes

In every language, certain sounds are considered by native speakers to be the "same" sound, even though they may be phonetically distinct. For example, native speakers of English consider the [l] in *lay* to be the same sound as that in *play,* even though the former is voiced and the latter is voiceless (see File 3.3 for an explanation of voicing). Or, if you ask a native speaker of English how many different sounds are represented by the underlined letters in the words *pin, bin,* and *spin,* they will probably say "two," grouping the aspirated [pʰ] of *pin* and unaspirated [p] of *spin* together. Though [pʰ] and [p] are phonetically different sounds, native English speakers typically overlook this difference and consider them to be the "same" sound.

One of the goals of this file is to understand more clearly the distinction between "same" and "different" sounds. To do this, we will discuss the terms **allophone** and **phoneme.** Since these concepts are the crux of phonological analysis, it is important that they be clearly understood. Perhaps the best way to explain these terms is through examples. On a separate piece of paper, transcribe the following five words:

top stop little kitten hunter

It is likely that you transcribed all of these words with a [t], like the following:

[tap] [stap] [lɪtl̩] [kɪtn̩] [hʌntɹ̩]

This is good, since it reflects something that is psychologically real to you. But, in fact, the physical reality (the acoustic phonetic fact) is that the 't' you transcribed in those five examples is pronounced slightly differently from one example to the next. To illustrate this, pronounce the five words again. Concentrate on what the 't' sounds like in each example, but be sure to say them as you normally would if you were talking to a friend.

What differences did you notice? Compare, for example, the /t/ of *top* to that of *stop.* You should be able to detect a short burst or puff of air after the /t/ in *top* that is absent in *stop.* That puff of air is called **aspiration,** which we will transcribe with a superscripted [ʰ]. So while a native speaker would consider the 't' sound in *top* and *stop* to be the same

sound, the 't' is actually pronounced differently in each word. This difference can be captured in the transcription, as in [tʰap] and [stap], respectively.

Now say the word *little* and *kitten*. We might say that the 't' in *little* sounds a lot "softer" than the one in *stop*, and is clearly voiced. For most speakers of American English (but not of British English), the 't' in words like *little* is pronounced as a flap, [ɾ], much like the *r* in Spanish in words like [paɾa] 'for' and [toɾo] 'bull'. English *kitten*, on the other hand, is pronounced with the same sound we hear in the expression "uh-oh," a glottal stop [ʔ]. So, we could transcribe *little* and *kitten* as [lɪɾl̩] and [kiʔn̩], respectively.

For some speakers of American English, in casual speech words like *hunter* are pronounced with no 't' at all, but rather as [hʌnr̩]. Try and say it this way and see if it sounds like something you've heard before. In any case, while you may have initially transcribed the five words above with a /t/, they may also be transcribed in a way that reflects the different pronunciations of that sound, as in the following:

[tʰap] [stap] [lɪɾl̩] [kʰɪʔn̩] [hʌnr̩]

To a native speaker, all of the words above have a 't' in them, at least at some psychological level. Proof of that lies in the fact that one may transcribe them all with a 't', at least until trained in transcription. Someone who lacks linguistic training would probably not hesitate to state that all the above words have a 't' and would need to be convinced that subtle differences, like aspiration, exist among them. In this sense, psychologically, the above words do have a 't'. On the other hand, we can observe that the 't' may be pronounced in several different ways.

Unlike a speaker of English, a native speaker of Hindi could not ignore the difference between aspirated and unaspirated sounds when speaking or hearing Hindi. To a speaker of Hindi, the aspirated sound [pʰ] is as different from unaspirated [p], as [pʰ] is from [b] to our ears. The difference between aspirated and unaspirated stops must be noticed by Hindi speakers because their language contains many words that are pronounced in nearly the same way, except that one word will have an aspirated stop where the other has an unaspirated stop. The data in (1) illustrate this.

(1)

Hindi	***Gloss***
[pʰəl]	'fruit'
[pəl]	'moment'
[bəl]	'strength'

A native speaker of English may not be aware of the difference between aspirated and unaspirated stops because aspiration will never make a difference in the meanings of English words. If we hear someone say [mæp] and [mæpʰ], we may recognize them as different pronunciations of the same word *map*, but not as different words. Because of the different ways in which [p] and [pʰ] affect meaning distinctions in English and Hindi, these sounds have different values in the phonological systems of the two languages. We say that these two sounds are **noncontrastive** in English, since interchanging the two does not result in a change of meaning. In Hindi, on the other hand, [p] and [pʰ] are **contrastive**, since replacing one sound with the other in a word can change the word's meaning. We will have more to say about this terminological distinction below.

Linguists attempt to characterize these different relations between sounds in language by grouping the sounds in a language's sound inventory into classes. Each class contains all of the sounds that a native speaker considers as the "same" sound. For example, [t] and [tʰ] in English would be members of the same class. But English [tʰ] and [d] are members of dif-

ferent classes because they are contrastive. That is, if you interchange one for the other in a word, you can cause a change in the word's meaning, e.g. [tʰaɪm] *time* vs. [daɪm] *dime*. On the other hand, speakers of Hindi would not classify [t] and [tʰ] as members of the same class because they perceive them as different. That is, they are contrastive.

A class of speech sounds that are judged by a native speaker to be the same sound is called a **phoneme**. Each member of a particular phoneme class is called an **allophone**, which corresponds to an actual phonetic segment produced by a speaker. That is, the various ways that a phoneme is pronounced are called allophones.

In this view, we can say that the 't' sounds in words like *stop, top, little,* and *kitten* all belong to a single class, which we will label by the symbol /t/, characterizing this particular phoneme class. By saying that *stop* and *top,* for example, each have the phoneme /t/, we are saying that the sounds [t] and [tʰ] are related.

In (2) we see how the phoneme /t/ is related to its allophones in English and how the Hindi phonemes /t/ and /tʰ/ are related to their allophones. In English, [t], [tʰ], [ɾ], and [ʔ] are allophones of the same phoneme, which we can label /t/. In this way, we can say that in English the phoneme /t/ has the allophones [t] as in [stap], [tʰ] as in [tʰap], [ɾ] as in [liɾl̩], and [ʔ] as in [kiʔn̩]. On the other hand, in Hindi, [t] and [tʰ] are allophones of different phonemes. Note that symbols representing phonemes are written between slashes; this distinguishes them from symbols representing phones, which are written between square brackets.

(2)

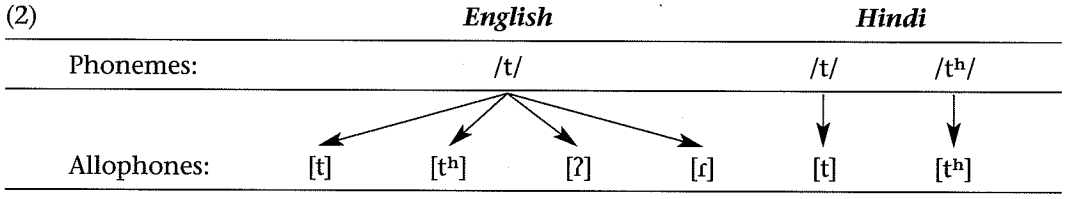

By providing a description like this, linguists attempt to show that the phonological system of a language has two levels. The more concrete level involves the physical reality of phonetic segments, the allophones, whereas phonemes are something more abstract. In fact, linguists sometimes describe phonemes as the form in which we store sounds in our memory. So, phonemes are abstract psychological concepts, and they are not directly observable in a stream of speech; only the allophones of a phoneme are.

The phoneme is a unit of linguistic structure that is just as significant to the native speaker as the word or the sentence. Native speakers reveal their knowledge of phonemes in a number of ways. When an English speaker makes a slip of the tongue and says [tʃi keɪn] for *key chain,* reversing [tʃ] and [k], he or she has demonstrated that [tʃ] functions mentally as a single unit of sound, just as [k] does. Recall from File 3.3 that this is not the only way to conceptualize [tʃ]: it is phonetically complex, consisting of [t] followed immediately by [ʃ]. Yet, since [tʃ] represents the pronunciation of a single phoneme /tʃ/ in English, no native speaker would make an error that would involve splitting up its phonetic components; you will never hear [ti kʃen] as a slip of the tongue (see File 9.9).

Knowledge of phonemes is also revealed in alphabetic spelling systems (see File 15.2). For example, English does not have separate letters for [pʰ] and [p]; they are both spelled with the letter *p*. Examples like this show that the English spelling system ignores the differences in pronunciation that don't result in meaning distinctions. For the most part, the English spelling system attempts to provide symbols for phonemes, not phonetic segments. In general, alphabetic writing systems tend to be phonemic rather than phonetic, though they achieve this goal with varying degrees of success.

4.2.2 Identifying Phonemes and Allophones: The Distribution of Speech Sounds

In order to determine whether sounds in a given language are allophones of a single phoneme or allophones of separate phonemes we need to consider the distribution of the sounds involved. The **distribution** of a phone is the set of phonetic environments in which it occurs. For example, nasalized vowels in English only occur in the environment of a nasal consonant. More precisely, a linguist would describe the distribution of English [ĩ], [õ], and so on, by stating that the nasalized vowels occur immediately preceding a nasal consonant. In this section, we will mainly be concerned with two types of distribution: contrastive distribution and complementary distribution, though a third distribution, free variation, will also be introduced.

Let us consider **contrastive distribution** first. Recall from above that a pair of phones is contrastive if interchanging the two can change the meaning of a word. This means that the sounds can occur in the same phonetic environment. It also means that the sounds are allophones of different phonemes. Two sounds are noncontrastive if replacing one phone with another does not result in a change of meaning.

Our earlier discussion of the patterning of [p] and [pʰ] in Hindi and English provides a good example of this difference. Recall that we said that in Hindi these two sounds could affect the meaning of a word based on examples like [pəl] *moment* and [pʰəl] *fruit*, where the two meanings are distinguished by the occurrence of [p] or [pʰ]. This means that the two sounds are contrastive in Hindi. In English, on the other hand, simply replacing [p] for [pʰ], or vice versa, will never affect a change in the meaning of a word; the sounds are noncontrastive in English.

We just determined whether or not [p] or [pʰ] are contrastive in Hindi and English by taking into account the distribution of sounds in each individual language. We did this by identifying a **minimal pair**. A minimal pair is defined as a pair of words with different meanings that are pronounced exactly the same way except for one sound that differs. When you find a minimal pair, you know that the two interchangeable sounds are contrastive and, thus, the sounds involved are allophones of different phonemes. If you try, you can think of many minimal pairs in English, or any other language you know well. For example, the minimal pair [tʰiːm] *team* and [tʰiːn] *teen* shows that [n] and [m] are allophones of separate phonemes in English since they can be used to contrast meaning. In Hindi, the words [pʰəl] 'fruit' and [bəl] 'strength' constitute a minimal pair, showing [pʰ] and [b] to be allophones of separate phonemes; [pʰəl] *fruit* and [pəl] *moment* also form a minimal pair in Hindi. But notice that there are no minimal pairs involving [pʰ] and [p] in English; these two sounds are never contrastive with respect to one another. Instead, they are allophones of the same phoneme, /p/.

Consider another example in which two languages make different distinctions using the same set of sounds. In English, it is possible to find minimal pairs in which [l] and [r] are contrasted, e.g., *leaf* [lif] vs. *reef* [rif]; *lack* [læk] vs. *rack* [ræk]. However, [l] and [r] are never contrastive in Korean. Consider the data below ([ɨ] represents a high central unrounded vowel).

Korean	***Gloss***
[param]	'wind'
[irɨm]	'name'
[pal]	'foot'
[mal]	'horse'

As these examples illustrate, minimal pairs can never be found for [r] and [l] in Korean because these two sounds do not appear in the same positions in words: [r] appears only

between two vowels, while [l] does not appear in this position. This observation about the distribution of [r] and [l] is not merely a property of these isolated examples but true of all Korean words containing these sounds. Observations of this sort play an important role in determining which sounds are considered to be the "same" by a native speaker; that is, sounds which are judged by a native speaker to be allophones of a single phoneme.

In addition to contrastive distribution, sounds may also be in **complementary distribution**. Sounds showing this type of distributional pattern are considered to be allophones of a single phoneme. To understand better what we mean by complementary distribution, think about what the term complementary means: two complementary parts of something make up a whole. For example, the set of people in your class at any given moment can be divided into the set of people who are wearing glasses and the set of people who are not. These two sets of people complement each other. They are mutually exclusive, but together they make up the whole class.

Now consider a linguistic example, namely, the distribution of the English sounds [p] and [pʰ].

spat	[spæt]	*pat*	[pʰæt]
spool	[spul]	*pool*	[pʰul]
speak	[spik]	*peek*	[pʰik]

As you can see in the English words just above, [p] and [pʰ] do not occur in the same phonetic environment. As a result, there are no minimal pairs involving a [p]-[pʰ] contrast. In fact, the phones are in complementary distribution: [p] occurs after [s] but never word-initially, while [pʰ] occurs word-initially but never after [s]. Since these sounds appear in different phonetic environments there can be no pair of words composed of identical strings of sounds except for [p] in one and [pʰ] in the other. As stated above, phones that are in complementary distribution are allophones of a single phoneme. In this case, [p] and [pʰ] are both allophones of the phoneme we can represent as /p/. Furthermore, the appearance of one allophone or another in a given context is predictable. For example, we can predict that the allophone [pʰ] (but never [p]) will appear in word-initial position. So even in words not listed above, we know that it will be [pʰ], rather than [p], that will occur at the beginning of a word. Similarly, we can predict that [p] (but never [pʰ]) will follow [s] in other words.

4.2.3 Free Variation

In some phonetic contexts more than one pronunciation of a given sound may be possible. Consider, for example, the following words containing [p] and [p˺] ([p˺] represents an unreleased voiceless bilabial stop).

leap	[lip]	*leap*	[lip˺]
soap	[sop]	*soap*	[sop˺]
troop	[trup]	*troop*	[trup˺]
happy	[hæpi]	—	*[hæp˺i]

It should be clear that these sounds both share some of the same phonetic environments; for example, they can both appear at the ends of words. Unlike the case of English [b] vs. [pʰ], or [m] vs. [n], however, there are no minimal pairs involving these sounds in the language. And note that although there are pairs of words in the list above that differ in only one sound, none of these words contrast in meaning. Thus, the choice between [p] and [p˺] in *leap, soap,* and *troop* does not make a difference in meaning; that is, the sounds are noncontrastive. Rather, they are interchangeable in word-final position. Sounds with this type

of patterning are considered to be in **free variation**. To a native speaker, sounds like [p] and [pˈ] that are in free variation are perceived as being the "same" sound. Thus, we can conclude that they are allophones of the same phoneme.

Summary

To summarize, a phone's distribution is the collection of phonetic environments in which the phone may appear; when linguists describe a phone's distribution they describe this collection. Relative to each other, two (or more) phones will be in contrastive distribution, complementary distribution, or in free variation. Phones in contrastive distribution may appear in minimal pairs and are allophones of different phonemes. Phones in free variation may also appear in the same phonetic environments but never cause a contrast in meaning; they are allophones of the same phoneme. In either of these two types of distribution, given a particular phonetic environment, one cannot predict which of the phones will occur. If the phones are in complementary distribution, their appearance in particular phonetic environments is predictable; they never appear in minimal pairs and they are allophones of the same phoneme.

Exercises

Look at the following Ukrainian words containing the sounds [s], [sʲ], [ʃ], and [ʃʲ]. The sounds [sʲ] and [ʃʲ] are palatalized variants of [s] and [ʃ]; palatalization sounds like a [j] sound right after (or on) the consonant and is very close to the [j] sound in the English word [bjuɾi] *beauty*. You might want to review these definitions before you begin: contrastive distribution, complementary distribution, and minimal pair. The words have been arranged to help you identify minimal pairs.

	[s]		[sʲ]		[ʃ]		[ʃʲ]	
1.	lɪs	'fox'	lɪsʲ	'sheen'	lɪʃ	'lest'		
2.	mɪska	'bowl'			mɪʃka	'little mouse'	mɪʃʲi	'mice'
3.	sapka	'little hoe'			ʃapka	'hat'		
4.	sɪla	'strength'			ʃɪla	'she sewed'	ʃʲistʲ	'six'
5.	sum	'sadness'			ʃum	'rustling'		
6.	sudɪ	'trials'	sʲudɪ	'hither'			koʃʲi	'baskets'
7.	sosna	'pine'	sʲomɪj	'seventh'	ʃostɪj	'sixth'		
8.	posadu	'job' (acc.)	posʲadu	'I will occupy'				

a. What minimal pairs can you identify in these words?

b. Is there a minimal triplet (like a minimal pair, but involving three sounds and three words)? What is it?

c. Which three of these four sounds are in contrastive distribution?

d. One of these sounds occurs only before a particular vowel. What is this sound, and what is the vowel? Which words indicate this?

e. Is the consonant you identified in (d) contrastive in Ukrainian, or not?

FILE

4.3

Phonological Rules

In File 4.2, we discussed the fact that phonemes and (allo)phones belong to different levels of structure in language—that is, phonemes are mental entities and phones are physical events. In this file we consider the connection between these two levels. The mapping between phonemic and phonetic elements is accomplished using **phonological rules** (recall that a rule of grammar expresses a pattern in a language). A speaker's knowledge of phonological rules allows him or her to "translate" phonemes into actual speech sounds; knowledge of these rules forms part of the speaker's linguistic competence. This change from the phonemic underlying form to the actual phonetic form of a word by means of phonological rules can be represented with a diagram:

phonemic form
⇓
rules
⇓
phonetic form

As an example, consider the English word *can* /kæn/. This word has a final /n/ sound in its phonemic form, and in fact it is frequently pronounced with a final [n]. If we listen carefully, however, we find that the final consonant of *can* (especially in casual speech) is often [m] or [ŋ]. The following examples illustrate this. (Here and throughout this file we use a fairly broad transcription style, recording phonetic detail only for the segments under discussion.)

I can ask	[aɪ kæn æsk]	(or [aɪ kn̩ æsk])
I can see	[aɪ kæn si]	(or [aɪ kn̩ si])
I can bake	[ai kæm beɪk]	(or [aɪ km̩ beɪk])
I can play	[aɪ kæm pleɪ]	(or [aɪ km̩ pleɪ])
I can go	[aɪ kæŋ goʊ]	(or [aɪ kŋ̩ goʊ])
I can come	[aɪ kæŋ kʌm]	(or [aɪ kŋ̩ kʌm])

As these transcriptions show, /n/ is pronounced as [m] when it precedes a labial consonant and as [ŋ] when it precedes a velar consonant. We can state this fact about English in a descriptive rule:

/n/ is pronounced as [m] before a labial consonant
[ŋ] before a velar consonant
[n] everywhere else.

(We will be adjusting this rule later on in this file.) Notice that a phonological rule has three parts: the sound(s) affected by the rule, the environment where the rule applies, and the result of the rule. Here /n/ is affected by the rule. The rule applies when /n/ is followed by a labial or velar consonant. The result of the application of the rule is that /n/ acquires the same place of articulation as the following consonant.

Now consider how the phonetic forms of some of the above examples are derived from the phonemic forms:

phonemic form:	/kæn æsk/	/kæn beɪk/	/kæn goʊ/
***apply rule*:**	kæn æsk	kæm beɪk	kæŋ goʊ
phonetic form:	[kæn æsk]	[kæm beɪk]	[kæŋ goʊ]

This illustrates what happens in speaking. In listening, a hearer reverses this process: he or she perceives the phonetic form of an utterance, then sends it "backwards" through the phonological rules, finally obtaining a phonemic form that matches a form stored in memory.

The rule illustrated above applies not only to /n/, but also to /t/ and /d/:

hat trick	[hæt trɪk]
hit batsman	[hɪp bætsmn̩]
night class	[naɪk klæs]
bad dream	[bæd drim]
head band	[hɛb bænd]
bad guy	[bæg gaɪ]

4.3.1 Natural Classes

Can we make one rule to state that /n/, /t/, and /d/ change place of articulation according to what sound follows? Is it random chance that these three sounds all seem to undergo the same phonological rule? To answer these questions, let's first take a look at the articulatory descriptions of these three sounds:

/t/	voiceless alveolar oral stop
/d/	voiced alveolar oral stop
/n/	voiced alveolar nasal stop

Not only are all three sounds alveolar stops, but they are the *only* alveolar stops in English. The description "voiceless alveolar oral stop" can only mean /t/. We can make the description more general by removing some of the properties:

/n/, /t/, /d/ alveolar stop

With respect to English, saying "alveolar stop" is the same as saying /n/, /t/, and /d/. These three sounds are all of the phonemes in English that are produced by stopping the flow of air at the alveolar ridge. Thus, they are the **natural class** of alveolar stops. A natural class is a group of sounds in a language that share one or more articulatory or auditory property, *to the exclusion of all other sounds in that language.* That is, in order for a group of

sounds to be a natural class, it must include all of the sounds that share a particular property or set of properties, and not include any sounds that don't.

4.3.2 Properties Used to Describe Natural Classes

All of the properties used in Files 3.3 and 3.4 to describe individual sounds can also be used to describe natural classes. For example, in the English vowels the monophthongs [i, u] and the first part of the diphthongs [eɪ] and [oʊ] are tense vowels, and there are no other tense vowels in English. Thus, these vowels are members of the natural class of tense vowels in English. Likewise, the consonants [k, g, ŋ] are all described as velar consonants, and they are the only velar consonants used in English; thus they constitute the natural class of velar consonants in English. Notice that we already referred to the natural class of velar consonants in the formulation of our rule at the beginning of this file. You'll recall that this rule affects the natural class of alveolar stops when followed by a member of either the natural class of velar consonants or the natural class of bilabial consonants. This shows that natural classes can be used to describe both the sounds affected by a rule *and* the environments where a rule applies.

In talking about groups of sounds, we must use a few properties in addition to those needed to describe individual sounds. For example, if you look at the consonant chart in (5) in File 3.3 you will notice that the only labiodental consonants in English are the fricatives [f] and [v], while the bilabial fricative slots are left empty. In many situations it is advantageous to refer to [f] and [v] together with [p, b, m, w] and [ʍ] as belonging to the same natural class. For this purpose we use the property *labial*.

Another property used to describe natural classes divides consonants into two groups, **obstruents** and **sonorants**. Obstruents are produced with an obstruction of the airflow. The sounds in this category are stops, fricatives, and affricates. Sonorants, on the other hand, are consonants produced with a relatively open passage for the airflow. Sonorant consonants include nasals, liquids, and glides. Thus, the class of labial obstruents in English is [p, f, b, v], while the class of labial sonorants is [m, w, ʍ]. The class of labial consonants is the union of both sets: [p, f, b, v, m, w, ʍ]. As we will see, being able to divide consonants into obstruents and sonorants is quite useful in stating phonological rules.

4.3.3 Classification of Phonological Rules

In addition to grouping sounds into natural classes, we can classify phonological rules according to the kind of process that they involve. Seven major kinds of processes are discussed here, along with examples from the phonology of English and other languages.

Assimilation

Rules of **assimilation** cause a sound to become more like a neighboring sound with respect to some phonetic property. In other words, the segment affected by the rule assimilates or takes on a property from a nearby (usually adjacent) sound. Rules of assimilation are very common in languages. The first rule we considered in this file falls into this category. We can call it alveolar stop assimilation because it applies to all alveolar stops (/t/, /d/, and /n/):

> **Alveolar stop assimilation** (English): Alveolar stops assimilate to the place of articulation of a following consonant.

Thus, when a sound having the properties alveolar and stop immediately precedes a labial consonant, this rule causes it to take on the property labial (thereby replacing its

(1) The major natural classes of English consonants

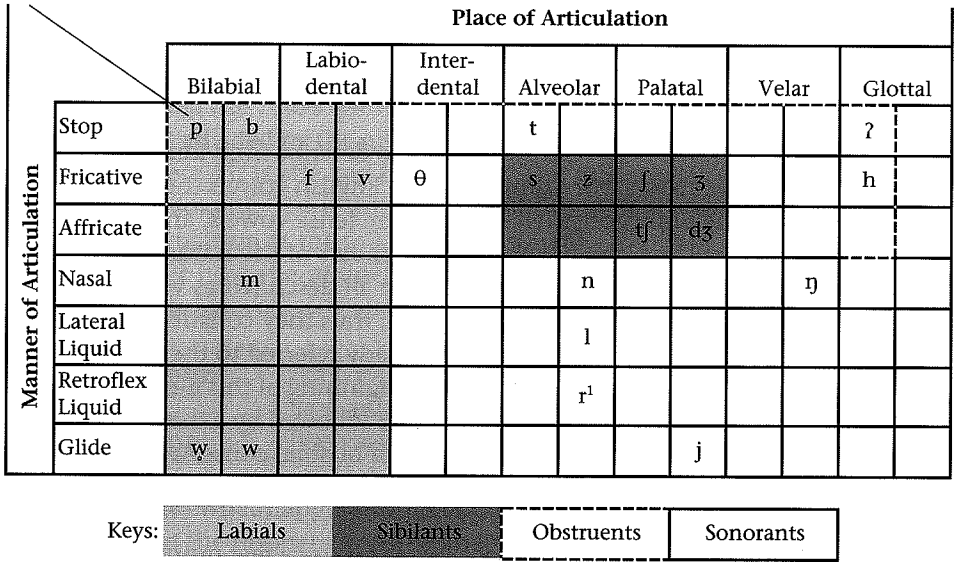

specification for alveolar). Similarly, this rule can apply to change the sound's place of articulation feature to dental when it precedes a dental consonant (examples such as *width* [wɪd̪θ] and *in this* [ɪn̪ ðɪs]), and so on for the other places of articulation.

Another common assimilation process is **palatalization**. Palatalization refers to a special type of assimilation in which a consonant becomes like a neighboring sound. For example, when American English speakers say *Did you eat?* rapidly, they very often pronounce it as [dɪdʒa it]. The sound [d] has been turned into a palatal sound [dʒ] because of the influence of the following palatal glide [j]. Vowels such as [i] and [e] also cause this change. The most common types of palatalization occur when alveolar, dental, and velar stops or fricatives appear before a front vowel. So the following are all common types of palatalization: [t] → [tʃ]; [d] → [dʒ]; [s] → [ʃ]; [k] → [tʃ]; [g] → [dʒ]. While there are variants on palatalization and other sounds can be palatalized, the main things to look for are (1) a sound becoming a palatal and/or (2) a sound change conditioned by a front vowel.

The rules of assimilation that we've discussed so far cause sounds to assimilate to *adjacent* sounds. This is a common way that assimilation occurs. However, long-distance assimilation also exists, and a relatively common type of long-distance assimilation is called **vowel harmony**. This typically causes all the vowels in a word to "harmonize" or agree in some property such as rounding or backness.

Finnish has a common type of vowel harmony rule, which can be stated as follows:

> **Vowel harmony** (Finnish): A back vowel becomes front when preceded by a front vowel in the same word.

By this rule, Finnish words have, with few exceptions, either all front vowels or all back vowels, but not both in the same word. We can see the vowel harmony rule in action when a suffix is added to the end of a word. In this case, the suffix vowel changes to match the quality of vowels in the word. For example, the suffix meaning 'in' has the form [-ssa] when added to a word where the last vowel is back, as in [talo] 'house', [talossa] 'in the house'. However, the suffix takes the form [-ssæ] when it attaches to a word with a final front vowel, as in [metsæ] 'forest', [metsæssæ] 'in the forest'. In cases like this, we can say that the vowel of the suffix harmonizes, or assimilates, to the preceding vowel.

[1] See footnote 1 in File 3.2.

b. Dissimilation

Unlike assimilation, which makes sounds more similar, rules of **dissimilation** cause two close or adjacent sounds to become less alike with respect to some property, by means of a change in one or both sounds. An example of dissimilation in Greek is the following:

> **Manner dissimilation** (Greek): A stop becomes a fricative when followed by another stop.

For example, in fast speech especially, the form /epta/ 'seven' can be pronounced as [efta], and /ktizma/ 'building' can be pronounced as [xtizma] ([x] is a voiceless velar fricative).

c. Insertion

Phonological rules of **insertion** cause a segment not present at the phonemic level to be added to the phonetic form of a word. An example of this kind of rule from English is voiceless stop insertion:

> **Voiceless stop insertion** (English): Between a nasal consonant and a voiceless fricative, a voiceless stop with the same place of articulation as the nasal is inserted.

Thus, for instance, the voiceless stop insertion rule may apply to the word *dance* /dæns/ → [dænts], *strength* /strɛŋθ/ → [strɛŋkθ], and *hamster* /hæmstɚ/ → [hæmpstɚ].

d. Deletion

Deletion rules eliminate a sound. Such rules apply more frequently to unstressed syllables and in casual speech. English examples include:

> **/h/-Deletion** (English): /h/ may be deleted in unstressed syllables.

The /h/-deletion rule would apply to a sentence such as *He handed her his hat* /hi hændəd hɚ hɪz hæt/ to yield [hi hændəd ɚ ɪz hæt]. Deletion is common in fast speech because it saves time and articulatory effort. Sounds like [h] that are not very perceptible are often the "victims" of deletion because speakers can save time and effort by deleting them without sacrificing much information. That is, the listener may not be relying on these sounds in order to understand what the speaker is saying.

e. Metathesis

Rules of **metathesis** change the order of sounds. In many instances, sounds metathesize in order to make words easier to pronounce or easier to understand. In Leti, an Austronesian language, consonants and vowels switch places when a word that ends in a consonant is combined with a word that starts with two consonants. The last two sounds in the first word trade places to avoid having three consonants in a row.

> **CV metathesis** (Leti): When three consecutive consonants occur, the first consonant trades places with the preceding vowel.

By this rule, /danat + kviali/ 'millipede' undergoes metathesis to become [dantakviali], and /ukar + ppalu/ 'index finger' becomes [ukrappalu]. On the other hand, /ukar + lavan/ 'thumb' does not undergo metathesis because there are not three consecutive consonants.

phonemic form:	/danat+kviali/	/ukar+ppalu/	/ukar+lavan/
***CV Metathesis*:**	dantakviali	ukrappalu	—
phonetic form:	[dantakviali]	[ukrappalu]	[ukarlavan]

f. Strengthening

Rules of **strengthening** (also called *fortition*) make sounds stronger. The rule of English aspiration, as stated below, provides an example:

> **Aspiration** (English): Voiceless stops become aspirated when they occur at the beginning of a stressed syllable.

The pronunciation of *tap* /tæp/ as [tʰæp] and *cat* /kæt/ as [kʰæt] illustrate the application of the English aspiration rule. Aspirated stops are considered to be stronger sounds than unaspirated stops because the duration of voicelessness is much longer.

g. Weakening

Rules of **weakening** (also called *lenition*) cause sounds to become weaker. The "flapping" rule of English is an example of weakening. [ɾ] is considered to be a weaker sound than [t] or [d] because it is shorter and it obstructs air less.

> **Flapping** (English): An alveolar oral stop is realized as [ɾ] when it occurs after a stressed vowel and before an unstressed vowel.

The pronunciation of *writer* /raɪtɹ̩/ as [raɪɾɹ̩] and *rider* /raɪdɹ̩/ as [raɪɾɹ̩] are examples of the application of this rule. Note that voicing assimilation is involved in the change of /t/ to [ɾ]: the /t/ takes on the "voicedness" of the vowels surrounding it.

4.3.4 Multiple Rule Application

To this point we have seen examples where one phonological rule applies. In reality there is often more than one change that occurs between a given phonemic form and a phonetic output. To illustrate this let's look at how plural nouns are formed in English. When you learned to write in English, you learned that the way to make a noun plural is to add an "s," which is usually pronounced [z]. There are actually three different phonetic forms of the English plural marker: [s], [z], and [əz], seen in the words *cats* [kæts], *dogs* [dagz], and *foxes* [faksəz]. We need only one phonemic form for the plural marker if we use two rules to derive the phonetic forms.

Try to pronounce [kætz] or [dags] in which the voicing quality of the final two consonants differ. You will probably find that it is difficult to produce a consonant cluster if one consonant is voiced and the other is voiceless. It is for this reason that the plural marker changes its voicing specification to match the sound it follows. We will use /-z/ as the phonemic form of the English plural marker, but notice that it doesn't matter whether we choose /-z/ or /-s/. It works out the same in the end.

> **Voicing assimilation** (English): /-z/ takes on the voicing specification of the preceding sound.

Now we can account for the different endings for *dogs* and *cats*, but what about the plurals of words like *fox, ditch, bush, orange, rouge,* and *maze*? What these words have in common is that they end in sounds that have a high-pitched hissing sound quality. These

sounds are called **sibilant** consonants. The English sibilants are [tʃ, ʃ, s, dʒ, ʒ, z]. Notice that the plural marker is also a sibilant. Because of the high-pitched hissing sound, it is very difficult to hear two sibilants that are next to each other. Try saying [fakss], [dɪtʃs], [brɪdʒz], etc., and you will get the idea. This is remedied by inserting a schwa between two sibilants.

Schwa insertion (English): Insert [ə] between two sibilants.

With these two rules, we can derive the plural for any English noun (except, of course, for special plurals like *oxen, octopi,* and *cherubim*).

phonemic form:	/kæt+z/	/dɑg+z/	/faks+z/	/brɪdʒ+z/
***schwa insertion*:**	—	—	faksəz	brɪdʒəz
voicing assimilation:	kæts	—	—	—
phonetic form:	[kæts]	[dagz]	[faksəz]	[brɪdʒəz]

Obligatory and Optional Rules

Notice that phonological rules may be **obligatory** or **optional**. Obligatory English rules include aspiration, vowel nasalization, vowel lengthening, and liquid and glide devoicing. Such a rule always applies in the speech of all speakers of a language or dialect having the rule, regardless of style or rate of speaking. The effects of obligatory rules are often very subtle and difficult to notice, but they are an important part of a native accent. For instance, it may be difficult to tell that a vowel is nasalized in English, but the application of vowel nasalization makes us sound like native speakers of English.

The existence of obligatory rules is what causes people to have foreign accents. It is easier to learn the rules of a new language than to "turn off" the obligatory rules of your native language. The very fact that we are often unaware of these rules causes us to apply them when they are not appropriate. When speakers of American English learn other languages, they often apply rules such as flapping and vowel reduction, even though the other language may not have these rules.

Optional phonological rules, on the other hand, may or may not apply in an individual's speech. Optional rules are responsible for variation in speech; for example, we can pronounce /kæn bi/ *can be* as either [kæm bi] or [kæn bi], depending on whether Alveolar Stop Assimilation is applied or not. The use of optional rules depends in part on rate and style of speech.

Conclusion

In this file we have covered seven types of phonological rules: assimilation, dissimilation, insertion, deletion, metathesis, strengthening, and weakening. These phonological rules operate on natural classes of sounds. We have also shown that a natural class is a group of all the sounds in a language that share some articulatory or auditory property(s) to the exclusion of all other sounds in that language. To describe natural classes we have used the properties *consonant, vowel, labial, sibilant, obstruent,* and *sonorant,* as well as properties used to describe individual consonants and vowels.

Exercises

1. List the members of the following natural classes of English sounds.
   a. alveolar obstruents
   b. voiced labial consonants

FILE

4.4

How to Solve Phonology Exercises

Because phonemes are important units of linguistic structure, linguists must have a general method for identifying them in all languages. But the task of determining what the phonemes of a language are and what allophones are assigned to them is not always straightforward. For one thing, the set of phonemes differs from language to language, and so a different analysis is required for each language. Moreover, phonemes are psychological units of linguistic structure and are not physically present in a stream of speech. As a result, it is not possible to identify the phonemes of a language simply by taking physical measurements on a sample of language. Nor is it always easy to identify phonemes by investigating a native speaker's intuitions, since the minute phonetic details on which decisions about phonemes are made are often precisely those which speakers are not accustomed to noticing.

To get around these problems, linguists have developed an objective procedure by which the phonemes of a language can be discovered through examination of a set of words written in phonetic transcription. This procedure is based on the following observations about patterns of sounds:

a. **Phonemes** make distinctions in meaning. If two sounds are members of separate phonemes, minimal pairs can almost always be found. For example, the minimal pair *led* and *red* is evidence that [l] and [r] are members of separate phonemes in English. But if two sounds are allophones of the same phoneme, minimal pairs differing only in those sounds will not exist. For example, [bʌʔn̩] and [bʌtʰn̩] are both possible pronunciations of the English word *button* (though [bʌtʰn̩] may sound a little stilted). This is because the sounds [ʔ] and [tʰ] are both allophones of the phoneme /t/. Thus, the meaning doesn't change.

b. The **allophones** of a phoneme are not a random collection of sounds but are a set of sounds that have the same psychological function. Accordingly, allophones of the same phoneme are systematically related to one another:

i. They share many phonetic properties.
ii. It is possible to predict which allophone will appear in a word on the basis of phonological rules.

By analyzing the patterns of sounds that are physically present, it is possible to draw conclusions about the psychological organization of a language, which is not directly observable.

4.4.1 How to Do a Phonemic Analysis

Although a phonemic analysis can be performed successfully on any language, it is easiest to begin with a problem based on English, since we already know in effect what the solution is. Look over the data below, which are given in a fairly narrow phonetic transcription:

pray	[pʰɻ̥eɪ]	fresh	[fɻ̥ɛʃ]
gray	[greɪ]	regain	[rigeɪn]
crab	[kʰɻ̥æb]	shriek	[ʃɻ̥ik]
par	[pʰar]	tar	[tʰar]
broker	[broʊkɻ̩]		

Beginning with the sounds [ɹ] and [ɻ̥], we attempt to answer the following question: are these sounds allophones of separate phonemes, or allophones of the same phoneme? (Of course, native speakers of English intuitively know that they are allophones of the same phoneme. However, the procedure for doing a phonemic analysis should produce the same answer without appealing to the intuitions of speakers.)

In order to answer this question, it is necessary to examine scientifically the **distribution** of sounds within these data. That is, for each sound in question we need to determine the set of phonetic environments in which it can occur. But just what do we mean by "environment"? For the time being, we can define the **environment** of a sound as the sounds that immediately precede and follow it within a word. For example, in the word [greɪ], [ɹ] is in the environment [g_e]; that is, [ɹ] is preceded by [g] and followed by [eɪ].

The best way to begin a phonemic analysis is to determine whether the sounds in question are contrastive. To do this, look first for minimal pairs. Suppose for a moment we were interested in the sounds [pʰ] and [tʰ] in the data above. These sounds do appear in a minimal pair: [pʰar] and [tʰar] have different meanings and differ phonetically by only a single sound in the same position. This tells us that [pʰ] and [tʰ] are in overlapping distribution and, more specifically, that they are in contrastive distribution, because the difference between them causes a difference in meaning. Therefore, they are allophones of different phonemes.

Returning to the problem at hand, namely, the status of [ɹ] vs. [ɻ̥], we see that there are no minimal pairs in the data that differ only by these two sounds. Since [ɹ] and [ɻ̥] are not in overlapping distribution in our data, we can assume that they are in complementary distribution. However, we must prove that this is so by making a generalization about where [ɹ] (but not [ɻ̥]) may appear, and vice versa. In order to do so we need to compare the phonetic environments of each of these sounds. The easiest way to do this is to make a list for each sound, as follows. (Note that "#" indicates a word boundary.)

[ɹ]	[ɻ̥]
[g_eɪ]	[pʰ_e]
[a_#]	[kʰ_æ]
[b_oʊ]	[f_ɛ]
[#_i]	[ʃ_i]

Once you have collected the list of phonetic environments for each sound, it is necessary to proceed as follows:

1. *Look at the environments to find natural classes.* [ɻ̥] is preceded by [pʰ], [kʰ], [f], and [ʃ], all of which are voiceless consonants. This generalization permits us to simplify the description of the environment for [ɻ̥]; instead of listing each sound separately, it is now possible to say:

[ɻ̥] appears after voiceless consonants.

Now look at the environments in which [r] appears. Are there any natural classes? Yes and no. Certainly [b] and [g] are voiced consonants, and [a] is also voiced, but the set that includes [b], [g], [a], the beginnings of words, and the ends of words does not form a natural class. Thus, the critical observation to make here is that there is no single natural class of environments in which [r] can be found.

We have looked at the sounds preceding [r] and [r̥], but what about the sounds that follow them? As you can see, only [r] may occur word-finally, but either [r] or [r̥] can occur before a vowel. Thus, the environments that condition the appearance of [r] or [r̥], i.e., the conditioning environments of these particular allophones, are their immediately preceding sounds.

2. *Look for complementary gaps in the environments.* So far, we have shown that [r̥] appears after voiceless consonants, while [r] appears in an apparently random set of environments. Yet, it is possible to make one more critical observation. [r] does not appear in the environments in which [r̥] appears, namely, after voiceless consonants. Moreover, [r̥] does not appear where [r] does; there is no [r̥] after voiced consonants or at the beginnings and ends of words. Since the environments of [r] and [r̥] have systematic and complementary gaps, we say that [r] and [r̥] are in complementary distribution. Therefore, they are allophones of the same phoneme.

3. *State a generalization about the distribution of each of these sounds.* In other words, write a rule that will make predictions about where each of the sounds can occur. Actually, we've done the hard part of this already by observing that [r̥] occurs following voiceless consonants. How should we state the distribution of [r]? We could try formulating our rule as follows:

[r̥] appears following voiceless consonants;
[r] appears following voiced consonants or vowels, or at the beginning or end of a word.

However, that's not a very succinct formulation of the rule. To simplify it, recall that wherever [r̥] occurs, [r] can't, because their possible environments form complementary sets. Therefore, we can revise our rule this way:

[r̥] appears following voiceless consonants;
[r] appears elsewhere.

4. *Determine the identity of the phoneme and its allophones.* This next step in writing the rule involves deciding what the phoneme to which these sounds belong should be. In order to do so, we need to decide which of the allophones is the **basic allophone** and which is the **restricted allophone**. We have determined that the conditioning environment for [r̥] consists of a single natural class of sounds. [r̥] is restricted to occurring only there, whereas [r] may appear anywhere else. Therefore, we can identify [r̥] as the restricted allophone and [r] as the basic one. It makes sense to name the phoneme after the basic allophone, since it is the one that can show up in a wider variety of contexts. Furthermore, the basic allophone is assumed to be the closest approximation of the mental "sound" that speakers store in memory. In choosing a name for the phoneme, we have made the leap from observable phonetic reality to unobservable psychological reality. (It is not always possible to choose one allophone as basic, however. In that case the phonology exercise's instructions will not tell you to do so, and any of the allophones would serve equally well as the name of the phoneme.)

We can improve on our rule once more. The arrows in the rule below mean "is pronounced as." We use slashes around symbols that represent phonemes and a single slash indicates the beginning of the environment specification:

/ r / → [r̥] / after voiceless consonants;
/ r / → [r] / elsewhere.

Now that we have formulated the necessary phonological rule, we can see which phonological process it involves (cf. File 4.3). In this rule a voiced phoneme changes into a voiceless sound when it follows another voiceless sound. In other words, /r/ becomes more like a preceding sound with respect to the feature of voicelessness. Therefore, we can conclude that the process of assimilation is involved in this phonological rule.

4.4.2 Some Potential Trouble Spots

The procedure outlined in the previous section will work for any language for which reliable phonetic transcriptions exist. However, beginners are often confused by certain questions.

For instance, if you discover that no minimal pairs exist for two sounds, is it possible to automatically conclude that they are allophones of the same phoneme? No. It is still necessary to show that the sounds are in complementary distribution, since allophones are predictable variant pronunciations of the same phoneme.

Consider what happens if you make a decision too soon. Using the data presented at the beginning of the previous section, suppose you wanted to know whether [g] and [ʃ] are allophones of the same phoneme. Since there are no minimal pairs differentiated by these sounds, it might seem reasonable to conclude that they are. (Of course, a speaker of English should have no trouble thinking of a minimal pair involving these two sounds. However, the data you will be given in phonology exercises will be sufficient for you to solve those problems.) But a careful examination of the data reveals that this is the wrong conclusion. Listing the data and the relevant environments, you find:

[g] appears in *gray* [greɪ], *regain* [rigeɪn]
generalization: [g] appears between vowels or at the beginning of a word;
[ʃ] appears in *fresh* [fr̥ɛʃ], *shriek* [ʃr̥ik]
generalization: [ʃ] appears at the beginning or end of a word.

As these data illustrate, [g] and [ʃ] are not in complementary distribution because their distributions overlap: either may occur at the beginning of a word. As a result, no phonological rule can be responsible for their distribution. In general, when *no generalization* can be made about where a group of sounds can occur, it is possible to conclude that they are members of separate phonemes. A conclusion based on such a demonstration is just as valid as showing that minimal pairs exist. This alternative way of showing that sounds are members of separate phonemes is useful because it's not always possible to find minimal pairs for all distinctive sounds. For example, there are no minimal pairs involving [ŋ] and [h] in English. But it is reasonable to assume that they belong to separate phonemes because they share few phonetic properties, and no phonological rule determines where they can occur.

The range of tests for identifying phonemes can be broadened somewhat by the use of **near-minimal pairs**. Recall that a minimal pair is a pair of words differing in meaning but phonetically identical except for one sound in the same position in each word. The definition of near-minimal pairs is the same, except that the words are *almost* identical except for the one sound. For example, *heard* [hr̥d] and *Bert* [br̥t] form a near-minimal pair involving [h] and [b]. We are justified in saying that [h] and [b] are allophones of separate phonemes because no conceivable phonological rule would permit only [h] at the beginnings of words ending in [d], and only [b] at the beginnings of words ending in [t].

One final point about minimal pairs: notice that we have not defined them as pairs of words that rhyme. It is not necessary for two words to rhyme in order to form a minimal pair. Consider English *state* [steɪt] and *steak* [steɪk], for example, or *boat* [boʊt] and *beat* [bit].

Nor is rhyming sufficient to qualify a pair of words as a minimal pair: *gray* [greɪ] and pray [pʰr̥eɪ] from the list of data above rhyme, but differ in two sounds. And to take another example, *glitter* and *litter* rhyme but do not form a minimal pair because they do not contain the same number of sounds.

Another question that often troubles beginners is this: when describing the environment in which a sound appears, how do you know where to look? In the problem we solved in the previous section, we considered only the sounds that preceded [r] and [r̥]. This is certainly not the only possibility. In fact, identifying conditioning environments is the most challenging part of doing a phonemic analysis.

Recall that in the previous section we temporarily defined the environment of a sound as the sounds immediately surrounding it. However, it is sometimes necessary to look beyond the sound's immediate environment. For instance, if you are examining the distribution of a vowel allophone, it is quite common that the conditioning environment involves a vowel in an adjacent syllable, even though consonants may intervene. It may also be necessary to consider preceding or following sounds even when they belong to another word that is adjacent in the stream of speech. However, it is best to start by examining the immediate environment of an allophone when you are trying to determine what its conditioning environment is.

Since there are many logically possible environments to consider, the task is made easier by eliminating all of those except the most plausible. This can be accomplished by using strategies like the following:

a. *Formulate hypotheses about the allophones.* Investigation of the world's languages has revealed that some sounds are more common than others (see File 4.7 for relevant discussion). For example:

- Voiced nasals and liquids are more common than voiceless ones.
- Oral vowels are more common than nasal vowels.
- Consonants of normal length are more common than long consonants.
- "Plain" consonants are more common than those with secondary articulations like velarization, palatalization, and labialization.

On the basis of these generalizations, it is possible to speculate that if a less common sound appears in a language, it is probably a restricted allophone. But these tendencies should be used only as a guide for forming hypotheses, not as a basis for jumping to conclusions, since some languages exhibit exceptions. For example, French has both nasal and oral vowel phonemes.

b. *Keep in mind that allophonic variation results from the application of phonological rules.* Also remember that rules usually involve some phonological process, such as assimilation or deletion. Once you have a hunch about which allophone is the restricted one, check the environment in which it appears for evidence that a phonological process has applied. This may involve looking in more than one place until you have discovered a reasonable candidate. In the problem in the previous section, we were guided by the knowledge that voicing differences in consonants are often caused by voicing assimilation, and that voicing assimilation frequently occurs in consonant clusters. Since /r/ is the second member of all of the clusters given, we concluded that the consonant preceding it constituted the conditioning environment. Even if it is not obvious that a phonological process has been at work, you should be able to write a phonological rule and, thus, state a generalization about where the allophones of the phoneme occur.

4.4.3 Flowchart for Discovering the Distribution of Sounds

The flowchart should help you to identify the type of distribution two (or more) sounds in a language have. The rectangular boxes ask you to do something or give you some information that your working through the flowchart has revealed. The diamond-shaped boxes pose a question. Try reading through the flowchart before you attempt to analyze the languages in the next file; it may help you to understand the relationship between the different types of distributions of sounds in a language.