D. Osherson + H. Lasnik eds., An Invitation to Cognitive Science MIT Press 1990 135

3

Phonology

Morris Halle

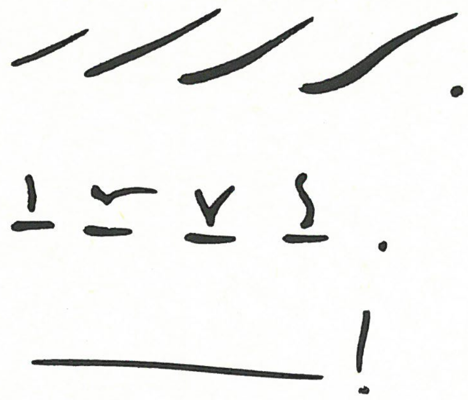

3.1 Speech as Strings of Discrete Sounds

A significant part of the knowledge that fluent speakers have of their language consists of the knowledge of its words. Normal fluent speakers of a language have little doubt that in producing an utterance they are producing a string of words. If pressed, they are likely to characterize words as temporally delimited sound sequences of which the letter sequences separated by blank spaces in our written or printed records of language are plausible representations. Unlike words on the printed page, however, the words in a spoken utterance are in reality not separated from one another (for instance, by little pauses). An utterance in which the speaker pauses briefly after each word sounds highly unnatural. Yet the fact that in speaking we do not separate words from one another affects our perception of utterances only rarely: in almost all cases we hear utterances in languages that we command (know) as sequences of words, and our understanding of an utterance is crucially based on our ability to segment the quasi-continuous acoustic signal into a sequence of discrete words.

The fact that words are not separated by pauses in an utterance does not give rise to serious conceptual problems. It is perfectly plausible that in learning a language we learn a large stock of words, which we then use to make up phrases and sentences. In pronouncing phrases and sentences, however, we no longer keep the words separate; rather, we pronounce one word after another as though the phrase consisted of just one very long "word." Knowing this fact enables hearers to segment the continuous utterance into its component words. In order to do this, however, hearers must know the language in which the utterance was framed. English speakers can readily segment into its component words an English sentence such as *adogneverplayswhenheisalone*, but they won't be able to do much with *pessinigdynehrajekdyžjesám*, the original Czech text by K. Capek of which the English sentence is a translation. The difficulty that English speakers encounter here is of course due to the fact that they lack knowledge of Czech and, most particularly, of Czech words.

It is reasonable to inquire at this point in what form speakers of a language store the words in their memory. (To simplify matters, we will ignore the fact that words have meanings, belong to specific lexical categories (such as noun, verb, adjective), and impose selectional restrictions on other words in the sequence (for instance, the verb *melts* selects the noun *sugar* but the verb *shaves* does not), and we will concentrate solely on words as phonic, auditory objects.)

A possible answer might be that when we learn a word, we memorize it as a purely acoustic event, as changes in the air pressure that are produced when the word is pronounced and that are sensed by our auditory system. There are reasons to doubt this answer. First, words have innumerable acoustic properties that speakers fail to remember. For example, we never remember the voice quality of the person from whom we learned a given word. Was the word spoken by a male voice or a female voice? Was it spoken slowly or rapidly, loudly or softly? And so on. All this information is obviously beside the point for the purpose of producing and understanding utterances in a particular language. So why should we remember it? Second, we are notoriously poor at identifying acoustic events. For example, few people can recognize more than a small number of sounds encountered in nature such as leaves rustling, the wind blowing, waves breaking, or one object hitting another, yet every normal speaker of a language has ready access to thousands of words.

Since there is nothing special about the sounds of speech from an acoustic point of view, and since we have a special propensity for memorizing words and no particular ability to memorize acoustic phenomena, it is plausible to suppose that we deal with words in a special way, radically different from our way of processing other acoustic signals that strike our ears. In fact, it has been assumed—more often tacitly than explicitly—that words are stored in memory as sequences of speech sounds—that is, as sequences of the sort of units that are at the base of alphabetic writing systems.

This proposal runs into immediate objections. First, words are obviously memorized by young children as well by others who lack all acquaintance with alphabetic writing. (The latter class is made up, on the one hand, by illiterates and, on the other hand, by people like the Chinese and the ancient Egyptians whose writing system is nonalphabetic.) One might wonder how speakers unacquainted with alphabetic writing would ever hit upon the idea of analyzing the quasi-continuous noises they hear into discrete sounds. Second, the fact that children learn the words of their native language with practically no overt teaching (of the relevant sort) implies that the knowledge children need in order to analyze words into sound sequences is acquired without benefit of teaching or even very extensive learning. This in turn raises the even more perplexing question of whether there can be knowledge that is not learned. (See chapters 8 and 9.) Third, since on this account memorizing words requires essentially the same processing of the speech signal as that involved in representing speech in an alphabetic writing system, we might wonder why great efforts have to be expended on teaching children to read and write, whereas they never need to be taught how to memorize the words of their native language. If these two processes are so similar, there should not be such a marked difference in the amount of training needed to acquire them.

Since the proposal that words are stored in memory as sequences of discrete speech sounds raises so many questions that lack obvious answers, it is necessary to remark at once that evidence in favor of the proposal is by no means lacking and that, as we will see, this evidence is quite persuasive.

Almost every language that has been studied supplements its basic stock of words via affixation. For example, in English we create agent nouns from verbs by suffixing *-er* to the verb stem (1a), we generate negative adjectives by prefixing *un-* (1b), and we make verbs by prefixing and/or by suffixing (1c):

(1)a. learn-er, work-er, teach-er, verbaliz-er, disestablishmentarianiz-er

b. un-clean, un-healthy, un-imaginable, un-original, un-otiose

c. em-power, en-rich, dark-en, hard-en, en-liv-en

We will use the term *morpheme* to refer to prefixes, suffixes, stems, and other meaning-bearing components of words. The processes of affixation do not always leave the component totally intact, as was the case in (1). In many instances affixation results in changes in the stems or the affixes or both. A simple illustration is provided by the formation of the feminine singular past tense and first person singular present tense forms of Russian verbs:

(2)a.		'crawl'	'can'	'bake'	'row'	'save'
	Past fem. sg.	polz-la	mog-la	pek-la	greb-la	spas-la
	Pres. 1 sg.	polz-u	mog-u	pek-u	greb-u	spas-u
b.		'stand'	'teach'	'sit'	'hold'	'bark'
	Past fem. sg.	stoya-la	uči-la	side-la	derža-la	laya-la
	Pres. 1 sg.	stoy-u	uč-u	siž-u	derž-u	lay-u
c.		'read'	'blow'	'live'	'know'	'sweat'
	Past fem. sg.	čita-la	du-la	ži-la	zna-la	pote-la
	Pres. 1 sg.	čitay-u	duy-u	živ-u	znay-u	potey-u

It is obvious from (2a) that the feminine singular past tense is signaled by the suffix *-la* and the first person singular present tense by the suffix *-u*. In (2a) the suffix has no effect on the stem; the examples in (2b, c) illustrate changes in the stem brought about by suffixation. The stems in (2b) end with a vowel in the past tense; yet this vowel is systematically eliminated in the present tense. (The *d* ~ *ž* alternation in *side-la, siž-u* is due to a special rule that will not concern us here.) As shown by their present tense forms, the stems in (2c) end with [y] or [v], but the stem-final consonant is deleted before the past tense suffix *-la*. A Russian-speaking child aged three or four knows how to form the past tense and present tense of these verbs, or of verbs very much like them. (In fact, presented with the first person present tense forms of nonsense verbs such as *nurey-u, butay-u* and asked to use them in the context appropriate for the feminine singular past tense (*Yesterday my mother* ——), a Russian-speaking child would undoubtedly spontaneously generate the appropriate *nure-la, buta-la*. For an analogous experiment with English-speaking children, see Berko 1958, discussed in chapter 8.)

It has been shown (Jakobson 1948; Lightner 1972) that the examples in (2) as well as a great many others are manifestations of the rules in (3):

(3)a. Delete a stem-final vowel before a vowel-initial suffix.

b. Delete a stem-final [y] or [v] before a consonant-initial suffix.

Part of the knowledge that a fluent speaker of Russian has therefore consists of the rules in (3). Research carried out by linguists since the beginning of scientific interest in language has shown that every other human language exhibits analogous rules. Such rules therefore represent an essential aspect of the knowledge that is universally required to produce and process human utterances.

What is important for our purposes about the rules in (3) is that in order to apply them, speakers must be able to analyze the words of their language into sequences of discrete sounds. In particular, a speaker must determine whether a given stem ends with a vowel, with [y], with [v], or with any other speech sound, and whether the suffix begins with a vowel or with a consonant. If we now assume that words are stored in memory as sequences of discrete speech sounds, we can readily explain the fact that all languages have rules of the type illustrated in (3). If we do not make this assumption, this fact remains a mystery.

3.2 Phonetics

If words are represented as sequences of discrete sounds, then the same must be true of utterances since, at least to a first approximation, utterances are sequences of words strung together one after another. Physically an utterance is manifested as an acoustic signal produced by a particular gymnastics executed by certain anatomical structures in the upper portions of the respiratory and digestive tracts, specifically the lower lip, the tongue, the soft palate (velum), and the larynx. The cavities bordered by these structures are commonly referred to as the *vocal tract*, and the acoustic signal that strikes the ears during speech is produced by changes in the geometry of the vocal tract. An X-ray motion picture recording the behavior of the vocal tract in the course of producing a particular utterance bears a striking resemblance to a stylized dance performed by dancers of great skill. If utterances are regarded as "dances" performed by the lower lip, tongue, soft palate, and other movable portions of the vocal tract, then one must also suppose that underlying each utterance ("dance") there is a "score" in some "choreographic" notation that instructs each "dancer" what to do and when. The different phonetic transcriptions of utterances are such "choreographic" notations, and the subsections that follow are devoted to a discussion of one such notation.

Figure 3.1 illustrates the six anatomical structures that are involved in the production of speech. Each of these "dancers" is capable of only limited behaviors. For example, the soft palate, which has the most restricted range of behavior, is capable only of being lowered or raised: when the soft palate is lowered, air can flow through the nasal cavities, resulting in the characteristic acoustic effect of nasalization; when the soft palate is raised, no air flows through the nasal cavities and these remain acoustically inert. The other anatomical structures involved in speech have a greater repertoire of behaviors, which we will examine in the following subsections.

The actions of the different anatomical structures involved in speaking are independent of one another. Thus, the soft palate can carry out its movements without regard for the movements simultaneously being car-

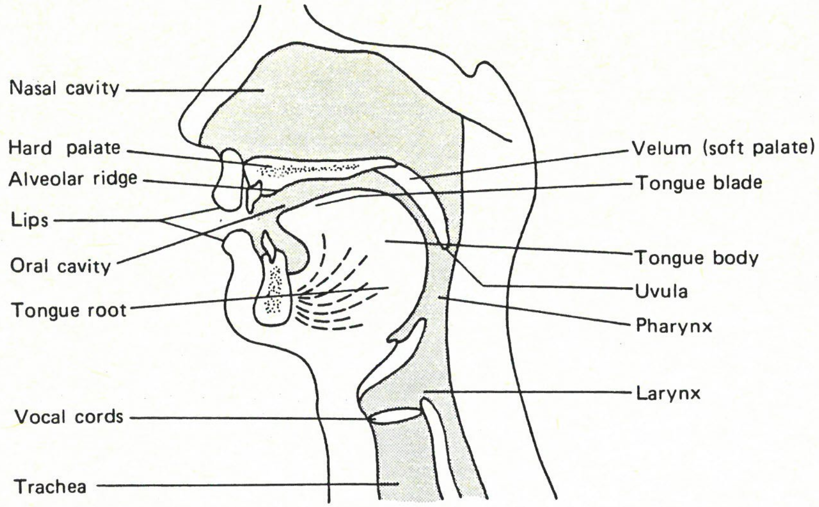

Figure 3.1
Midsagittal section through the human head and neck showing the articulators that make up the vocal tract. Adapted, by permission, from figure 4.5 of A. Akmajian, R. A. Demers, and R. M. Harnish (1984). *Linguistics: An introduction to language and communication.* Cambridge, MA: MIT Press.

ried out by the lower lip, the larynx, and/or the tongue. Viewed from the vantage point of articulatory independence, the tongue behaves not as a single entity but as three distinct agents: the blade, the body, and the root. We will use the term *articulator* to designate anatomical structures that function independently in the production of speech—namely, the six listed on the left in (4). The terms on the right are used to refer to these articulators in a number of recent studies and occasionally in this chapter as well:

(4) Larynx = LARYNGEAL
Velum = SOFT PALATE
Lower lip = LABIAL
Tongue blade = CORONAL
Tongue body = DORSAL
Tongue root = TR

3.2.1 The Production of the Vowels

What characterizes the articulation of vowels of all kinds is that in the center of its passage from the lungs up through the pharynx and the oral cavity the air flow encounters no obstacle sufficient to create turbulence or total blockage. The tongue body plays a central role in the production of

all vowels. (The other articulators also take part in vowel production, but their role is subsidiary.) The tongue body is capable of being moved up and down in a direction parallel to the spinal column as well as back and front in a direction perpendicular to the spinal column, and different positions of the tongue body serve to differentiate one vowel from another. The positioning of the tongue body (and of the other articulators, as well) is of course in each case the result of forces exerted by muscles. Since the acoustic signal emitted by the vocal tract depends directly on the positions and behaviors of the articulators and only indirectly on the contractions and relaxations of the muscles, the phonetic notation focuses on the behavior of the articulators rather than on that of the individual muscles. (For some discussion of the muscular behavior underlying the production of speech, see Halle 1983.)

Although the tongue body has extraordinary freedom of movement and can assume many different positions relative to the stationary parts of the vocal tract—that is, relative to the roof of the mouth and the back wall of the pharynx—human languages use only a limited number of these positions to differentiate the vowels. In fact, only three vertical and two horizontal positions of the tongue body are used systematically to generate distinct vowels. The English short vowels in *bit, bet, bat* exhibit the three vertical positions of the tongue body. As speakers of English can readily establish by self-observation, the tongue body is lowest in *bat* and highest in *bit*. In producing each of these three vowels, the tongue is in a *forward* position. Retracting the tongue to a *back* position while pronouncing the vowel in *bat* produces the vowel in *cot* as it pronounced in most varieties of American English (though not in that of eastern Massachusetts or in many of the dialects spoken in Britain). Similarly, retracting the tongue to a back position while pronouncing the vowel in *bet* produces the vowel in *but*. English does not make systematic use of the backed counterpart of the vowel in *bit*. This sound is encountered in only a few words (for example, the adverb *just* in such expressions as *just a minute*).

In the phonetic alphabet in widest use, that of the International Phonetic Association (IPA), the vowels just described are symbolized as follows:

(5) [ɪ] bit [ɯ] just
[ɛ] bet [ʌ] cut
[æ] bat [ɑ] cot

As noted earlier, only a limited number of articulator configurations play a role in language. Jakobson (1938) proposed—and linguists now widely (though not universally) accept—that for each independent articulator behavior, or *feature*, languages utilize exactly two configurations. Formally this fact is expressed by representing each feature with a coefficient—[xF]—where the coefficient x is understood as a variable ranging over the values + and −.

In the horizontal plane the tongue body assumes exactly two relevant positions. We capture this fact with the binary feature [back]: in [+back] vowels (see (6)) the tongue body is retracted, whereas in [−back] vowels it is advanced. The featural description of the vertical position of the tongue body, where three positions are actually distinguished, is somewhat less straightforward. Since we have only binary features at our disposal, we must use two distinct features: [high] and [low]. Since each of these two features can assume two values, they provide the means for distinguishing four vowel types. It appears, however, that no language utilizes the feature complex [+high, +low]. We implement this universal restriction by imposing on feature complexes a formal constraint that disallows the complex [+high, +low]. The vowels in (5) will therefore be represented by means of the features [back], [high], and [low] as shown in (6):

(6)

	[ɪ]	[ɛ]	[æ]	[ɯ]	[ʌ]	[ɑ]
[back]	−	−	−	+	+	+
[high]	+	−	−	+	−	−
[low]	−	−	+	−	−	+

The feature complexes in the columns of (6) are partial definitions of the speech sounds represented by the alphabetic symbols in the topmost line. We thus have two ways of representing the speech sounds: by means of alphabetic symbols or as feature complexes. The representation in terms of alphabetic symbols implies that speech sounds are atomic entities—in other words, units that are not to be analyzed further into their constituent properties. By contrast, the representation in terms of feature complexes implies that speech sounds are composite entities made up of features. In section 3.3 we will see evidence suggesting that the latter rather than the former representation more accurately reflects what goes on in speakers' minds.

The three features in (6) represent the different capabilities of the tongue body articulator in producing vowels. They do not, of course, exhaust the capabilities of the human vocal tract as a generator of vowels. Languages often utilize the lips to distinguish different classes of vowels. The vowels in (6) are all generated with spread, rather than rounded, lip. In English, lip rounding is not admitted in the nonback vowels. In the back vowels, on the other hand, there are contrasts between rounded and unrounded cognates. Thus, parallel to [ɯ] as in *just* (*a minute*), American English has [ʊ] as in *put*; and parallel to [ɔ] as in *caught*, most American dialects have [a] as in *cot*. In addition, American English has the rounded [u] and [o] as in *shoe* and *show*, and the unrounded [ʌ] as in *cut*. The remaining rounded

back vowels are of somewhat limited distributions in different dialects. For example, the rounded counterpart of the nonlow back vowel in *cut* is found in many British dialects in such words as *got, lock, Tom*, whereas the rounded counterpart of the low back vowel in *cot* is found primarily in the eastern Massachusetts dialect in the same words. By contrast with English, many French dialects exhibit rounding with both back and front vowels.

(7)

[−back, −round]			[−back, +round]			[+back, +round]		
[i]	bise	'north wind'	[y]	ruse	'ruse'	[u]	rouge	'red'
[e]	thé	'tea'	[ø]	creuse	'hollow'	[o]	sauge	'sage'
[ɛ]	thèse	'thesis'	[œ]	veuve	'widow'	[ɔ]	loge	'box'

In addition, French has a [+back, −round, +low] vowel: [a] in *âme* 'soul'.

The tongue root articulator plays a crucial role in the distinction between the so-called long (diphthongized) and short (plain) vowels of English illustrated in (8). We will follow the suggestion made by Halle and Stevens (1969) in assuming that these pairs of vowels contrast as [+ATR] (advanced tongue root) versus [−ATR].

(8)

[−back, −round]		[+back, +round]		[+back, −round]	
[+ATR]	[−ATR]	[+ATR]	[−ATR]	[+ATR]	[−ATR]
peel	bill	boom	bull	——	——
pale	bell	bone	——	——	come
——	pal	call	——	calm	bomb

In other treatments of English phonetics these distinctions are characterized by means of a contrast called *tense/lax* or *narrow/wide*.

The coronal or tongue blade articulator is active in English in signaling contrasts such as the one between the vowels in *cull–curl* and *bun–burn*.

The lowering and raising of the soft palate (velum) produce the contrast between nasal and oral vowels, a contrast systematically represented in languages such as French, Polish, and Portuguese, but not in English.

Laryngeal features affect the pitch of the vowels as well as properties of voice quality, sometimes referred to by terms such as *creaky voice* and *breathy voice*. The role of the laryngeal features in the production of consonants is discussed in the following section.

3.2.2 The Production of the Consonants

In the production of a vowel the air flowing from the lungs to the lips encounters no obstruction sufficient to create blockage or turbulence. By contrast, in the production of a consonant such an obstruction must always be present. To create this obstruction, one of the articulators makes full or virtual contact with the stationary part of the vocal tract (the roof of the mouth, the upper lip, or the rear wall of the pharynx). Of the six articulators, the lower lip, tongue blade, tongue body, and tongue root are capable of making this type of contact and therefore play a central role in the production of consonants. We will use the term *place articulators* to designate these four articulators.

The difference between the production of consonants and the production of vowels is perhaps manifested most strikingly in the pronunciation of syllables like English *kick* or *cook*, where most of the action involves the dorsal articulator. In the initial and final—consonantal—portions of the gesture by means of which these two words are produced the dorsal articulator makes contact with the roof of the mouth, whereas in the middle—vocalic—portion no such contact takes place.

This distinction between consonants and vowels is formally implemented by the feature [consonantal]: [+consonantal] sounds involve significant contact by a place articulator with a portion of the vocal tract, [−consonantal] sounds lack such contact.

There are differences among the consonants with respect to the type of contact involved in their production. The *stops* such as [p t k b d g] are produced with total blockage; they contrast with the *fricatives* such as [f v s z θ (as in *thin*) š (as in *shin*) ž (as in *usual*)], which are produced with a constriction narrow enough to generate turbulent air flow yet not narrow enough to block the flow of air. This distinction is formally implemented by the feature [continuant]: stops are [−continuant], fricatives are [+continuant].

Another distinction among consonants is implemented by a third feature, [strident], which contrasts [s z] (English *lease*, *ease*) and [θ ð] (*teeth*, *teethe*). Phonetically the former, noisier [+strident] consonants are produced by directing the air flow at right angles to a sharp obstacle, whereas their less noisy [−strident] counterparts are produced with an air flow that is parallel rather than perpendicular to the obstacle.

A fourth consonantal feature is connected with pressure buildup inside the vocal tract. In [+sonorant] consonants such as [m n l r] there is no significant pressure buildup inside the vocal tract; in [−sonorant] consonants such as [p t k b d g f s š v z ž] there is a noticeable increase in the pressure in the vocal tract.

These four features, which are known as the *stricture* features, participate in an important hierarchical relation: [continuant], [strident], and [sonorant] are used only for distinguishing among consonants and play no role whatsoever in the production of [−consonantal] sounds. The stricture features present yet another property that strikingly differentiates consonants from vowels. Each feature involved in the production of vowels—such as [back], [nasal], or [round]—is actualized by only one specific articulator: [back] by the tongue body, [nasal] by the soft palate, and [round] by the lower lip. By contrast, the stricture features [continuant], [strident], and [sonorant] are not articulator-bound; rather, they are realized by one of the four place articulators. As a result, when specifying a [+consonantal] sound, it is always necessary to indicate which of the four place articulators is the one that executes its stricture features.

These striking differences between vowels and consonants must not be allowed to obscure the obvious fact that for the production of both types of sounds, speakers have at their disposal only a single piece of anatomical machinery, the vocal tract with its six articulators. If the features discussed in section 3.2.1 correctly characterize the behaviors that each articulator is capable of, we should expect to encounter these features not only in the production of vowels but also in the production of consonants. And in fact we do.

The three features executed by the dorsal articulator, [back], [high], [low], distinguish the six classes of vowels in (6). But what role do they play in the production of consonants? Consider the difference in English between the [k] in *keel* and the [k] in *cool* (keeping in mind that here, and throughout this chapter, we will be concerned with the sounds rather than their—often idiosyncratic—orthographic representation). In the production of both words the tongue body touches the roof of the mouth, but contact occurs in different places: farther forward in the case of *keel* and farther back in the case of *cool*. This parallels precisely the position of the tongue body in the vowels of the two words. It is therefore said that the [k] in *keel* is [+high, −back], whereas the [k] in *cool* is [+high, +back]. Ladefoged and Maddieson (1986, 18ff.) indicate that in some languages the two different [k] sounds serve to distinguish otherwise identical words. Thus, the dorsal articulator serves to distinguish consonants in much the same way that it distinguishes vowels.

The parallelism between vowels and consonants is not restricted to the feature [back] but extends to the features [high] and [low]. In the vowels these two features implement three rather than four distinctions in the vertical position of the tongue body. Significantly, consonants too require three distinctions in the vertical position of the tongue. Thus, in addition to the *velar* [k]-type sounds discussed above, many languages have *uvular* consonants in which the obstruction produced by the tongue body is located at the level of the uvula (figure 3.1). Since the *velar* [k]-type sounds are [+high], these *uvular* consonants are [−high]. Among the languages that utilize the distinction between velar [+high] and uvular [−high] stops are many dialects of Arabic and other Semitic languages, as well as Quechua, one of the major languages spoken in Peru and adjacent areas, and Serer, a West African language (see Ladefoged 1964, 21–22).

We have seen that the feature [low] distinguishes two kinds of [−high] vowels. A parallel distinction is also found among the consonants, although it appears to be somewhat rare. A number of Arabic dialects have *pharyngeal* ([+back, −high, +low]) fricatives contrasting with uvular ([+back, −high, −low]) fricatives (Ladefoged and Maddieson 1986, 82).

Strikingly, no known language has consonantal counterparts to vowels that are [−back, −high]. The reason for this becomes obvious once we recall the fundamental distinction between vowels and consonants, namely, the requirement that the articulator make contact with a wall of the vocal tract in the production of consonants. Since the walls of the vocal tract are located above and in back of the tongue body, we can use the tongue body to produce consonantal contact only if it is [+high] and/or [+back].

In making contact, the coronal articulator can choose between two points along the hard palate and can vary distinctively the extent of the contact area. The place of coronal (tongue blade) contact is controlled by the feature [anterior]. *Coronal* consonants like [t d s z θ ð] produced with contact in front of the alveolar (teeth) ridge on the upper palate are [+anterior]; sounds like [š ž r] produced with contact behind the alveolar ridge are [−anterior]. The extent of the contact area is determined by the feature [distributed]: [+distributed] sounds like [θ ð š ž] are produced with a flat tongue that for some distance parallels the roof of the mouth; [−distributed] sounds like English [t d s z] are produced by touching or approaching the hard palate with the tongue tip. In English only [r] is [−distributed, −anterior], but many languages of India, for example, have a whole complement of [−distributed, −anterior] (*retroflex*) sounds, often symbolized by a dot under the letter representing the corresponding [+anterior] sound: [ṭ ḍ ṣ ẓ]. (9) illustrates the feature composition of the coronal consonants:

(9)

	[t d s z]	[θ ð]	[š ž]	[r ṣ ẓ ṭ ḍ]
[anterior]	+	+	−	−
[distributed]	−	+	+	−

In most languages only a single place articulator is involved in the production of the consonants, but consonants produced by the simultaneous operation of two place articulators are by no means unknown. The clicks of the southern African languages involve both the coronal and dorsal articulator. The widely spread *labiovelar* consonants [kp gb] encountered in the names of the African languages *Kpelle* and *Igbo* simultaneously involve the lips and the dorsal articulator. And there are even languages like Kinyarwanda (Sagey 1986) in which three place articulators participate in the production of a single consonant.

Unlike the four place articulators, the remaining two articulators—the soft palate and the larynx—participate freely in the production of consonants. When the soft palate is lowered, air is allowed to pass through the nose and the characteristic acoustic effect termed *nasal* is produced.

Sounds produced with the lowered soft palate are [+nasal]; sounds produced with the raised soft palate are [−nasal]. Since as a result of this lowering of the soft palate no pressure is built up in the vocal tract, [+nasal] consonants are always [+sonorant].

The primary role of the larynx in the production of consonants lies in the control of vocal cord vibrations (*voicing*) and aspiration. The presence versus absence of voicing distinguishes the two sets of consonants in (10):

(10) Voiced [b d g v z ð ž]
Voiceless [p t k f s θ š]

The vocal cords make up the edges of the slit (*glottis*) on the bottom of the vocal tract through which air from the lungs enters into the larynx and passes from there into the pharynx and the rest of the vocal tract. In order for air to flow upward from the lungs, the pressure in the lungs must exceed that in the cavities above. For small pressure drops across the glottis the vocal cords will or will not vibrate depending on their stiffness: if they are slack, they will vibrate; if they are stiff, they will not (Halle and Stevens 1971). At greater pressure drops, however, stiffness cannot prevent vocal cord vibration: in such cases greater stiffness results in more rapid vibrations, which are perceived as higher pitch, whereas lesser stiffness translates into slower vibrations and lower pitch.

Since air flows from the lungs upward during speech, the pressure of the air inside the lungs must be somewhat greater than that of the ambient (atmospheric) air. We have seen that when [−sonorant] sounds are produced, pressure builds up inside the vocal tract, but when [+sonorant] sounds are produced, no air is trapped in the vocal tract and no pressure is built up. This situation is illustrated schematically in figure 3.2. If we

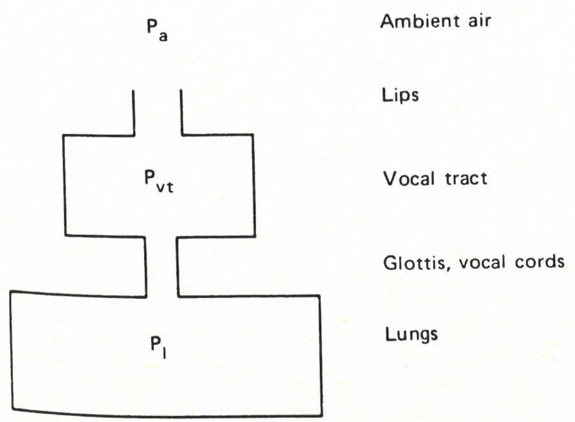

Figure 3.2
Schematic diagram of the air pressure relationships that obtain in the production of speech.

assume that during speech the lung pressure P_l does not vary appreciably and since the atmospheric pressure P_a is constant during a given utterance, then it follows that the pressure drop across the glottis $P_l - P_{vt}$ will be noticeably smaller in [−sonorant] sounds than in [+sonorant] sounds. We already know that when the pressure drop across the glottis is small, differences in vocal cord stiffness will determine whether or not they vibrate, whereas when the pressure drop across the glottis is large, differences in stiffness correlate with the rate at which the vocal cords vibrate. In other words, in nonsonorant sounds vocal cord stiffness controls the presence or absence of vibrations (voicing), whereas in sonorant sounds—and most especially in vowels—which normally are voiced, the stiffness of the cords determines the rate of vibration and hence the perceived pitch of the sound. Thus, the variation in stiffness of the vocal cords has perceptually two rather distinct consequences: presence versus absence of voicing in one class of speech sounds, and variations in pitch in another class of speech sounds. In view of this, it is worth noting that in many languages voiceless consonants are related to high-pitched vowels and voiced consonants to low-pitched vowels.

In addition to controlling the stiffness of the edges of the glottis, a speaker can control the size of the glottal opening, by spreading or constricting (pressing together) the vocal cords. When the vocal cords are spread, a special [h]-like sound—known to phoneticians as *aspiration*—is produced. In English, aspiration differentiates the voiceless stops in word-initial position from their cognates in position after [s], as illustrated by the pairs in (11):

(11) port – sport, till – still, core – score

In other languages, such as those of India, the aspiration distinction is much more widespread among the consonants. Thus, Hindi has the full complement of four stops allowed by the two binary features of [stiff vocal cords] and [spread glottis].

(12)	[p t k]	[pʰ tʰ kʰ]	[b d g]	[bʰ dʰ gʰ]
[stiff vocal cords]	+	+	−	−
[spread glottis]	−	+	−	+

The features for consonants and the articulators for these features are summarized in figure 3.3.

3.3 On the Psychological Reality of the Features

We have seen that speech sounds are not the ultimate constituents of language but instead complexes of features that are themselves structured in

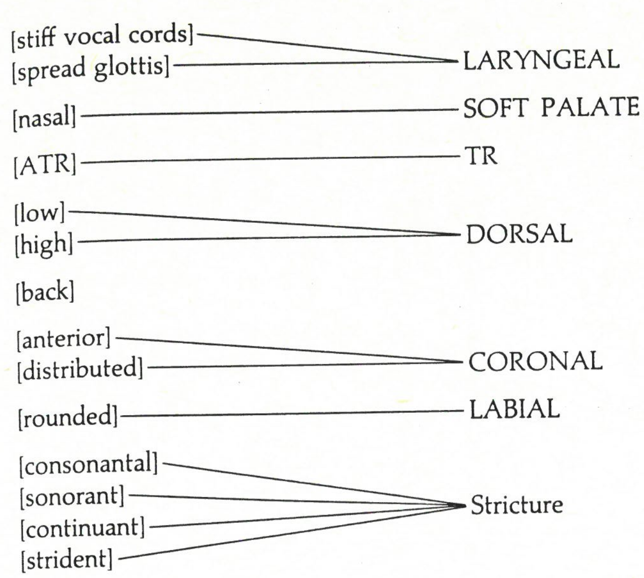

Figure 3.3
List of features and their articulators. The features in the *stricture* set are executed by one of the six articulators listed above, whose identity must be stipulated for each sound. For additional information on the internal organization of the features, see Sagey 1986.

a specific way. In this section we will see evidence that this elaborate structure not only is required by the way sounds are produced in the human vocal tract but also enters directly into the knowledge that speakers have of their language.

Three distinct suffixes are used to signal the regular plural of English nouns:

(13) [ɪz] places, adzes, porches, cabbages, ambushes, camouflages
[s] lips, lists, maniacs, telegraphs, hundredths
[z] clubs, herds, colleagues, phonemes, terns, fangs, holes, gears, pies, apostrophes, avenues, cellos, violas

The choice of suffix is governed by the last sound in the word, as stated in rule (14):

(14) [ɪz] if noun ends with [s z č ǰ š ž], otherwise
[s] if noun ends with [p t k f θ], otherwise
[z]

Examination of the groups of consonants figuring in (14) reveals that the members of each group share features that distinguish the group from all other sounds in the language. Translated into feature notation, the rule for the English plural suffiix reads as follows:

(15) [ɪz] if noun ends with [+coronal, +strident], otherwise
[s] if noun ends with [+stiff vocal cords] ([−voice]), otherwise
[z]

These two formulations predict different results with respect to nouns ending with sounds that are not part of the English sound system (L. Menn, personal communication). Rule (14) predicts that such words would invariably be given the suffix [z]. By contrast, (15) predicts that any one of the three suffixes could occur depending on the nature of the last sound in the word. A good test of these alternatives is provided by the German name *Bach*, which ends with the voiceless dorsal continuant, which is not found in English. If English speakers used rule (14), they would give [baxz], with [z], as the plural of this word; if they used rule (15), they would instead give [baxs], with [s]. Since the great majority indeed say [baxs], we must conclude that (15) is the rule employed.

This result implies that ordinary speakers have knowledge of the feature composition of the sounds of speech, including not only the sounds of their own language but also those of any other language. Thus, unless a more compelling interpretation of the results of our experiment is offered, we must accept the conclusion that *for the speaker words are made up of discrete speech sounds that are themselves complexes of features*. That the speaker might deny having this knowledge does not invalidate this conclusion; rather, it shows that we may know things without being aware of it. (See chapter 9.) The existence of such knowledge has been recognized almost since the beginnings of scientific concern with human knowledge. In the dialogues of Plato, for example, Socrates elicits many types of unconscious knowledge from his untutored interlocutors.

3.4 Rules and Rule Interaction

3.4.1 Plural Formation in Kasem

Affixation is one of the most common types of word modification (inflection and derivation) encountered among the languages of the world. A typical example is given by the singular and plural forms of Class C nouns in Kasem, a West African language, illustrated in (16). (For sources of data and additional discussion, see Halle 1978.)

(16)		'boy'	'granary'	'moon'	'shoulder'
	Singular	bakad-a	tul-a	čan-a	bakal-a
	Plural	bakad-i	tul-i	čan-i	bakal-i

As shown in (17), noun stems need not end in consonants: they may end in vowels as well. (Long vowels are represented by sequences of identical vowels.)

(17)		'pot'	'yam'	'chameleon'	'slave'
	Singular	kambi-a	pi-a	malaa	kabaa
	Plural	kambi	pi	male	kabe

If the forms in (17) arose via suffixation alone, we would expect the forms in (18) instead:

(18)					
	Singular	kambi-a	pi-a	malaa-a	kabaa-a
	Plural	kambi-i	pi-i	malaa-i	kabaa-i

Comparing the expected forms with the actual forms, we see that whenever the suffix vowel is identical with the stem vowel, one of the two vowels disappears. We will assume that the suffix is deleted. We therefore postulate rule (19), Truncation:

(19) *Truncation*

$V_1V_2 \rightarrow V_1$

Condition: V_2 is a word-final vowel and $V_1 = V_2$.

Truncation accounts for the singular forms in (17) as well as for the plural forms *kambi, pi*.

We still need to account for the plural forms [male] and [kabe] in place of the expected [malaa + i] and [kabaa + i]. These forms involve several rules. First, long vowels are shortened in prevocalic position by rule (20), Shortening:

(20) *Shortening*

VV → V before V

Kasem vowels are also subject to rule (21), Contraction, which converts [ai] into [e] and [au] into [o]. The effect of Contraction is to monophthongize the diphthongs [ai] and [au] and replace each one with the nonhigh cognate of its second element—that is, [ai] → [e], [au] → [o] but [aa] → [a]:

(21) *Contraction*

$$\underset{[+\text{low}]}{V_1}\ \underset{[+\text{high}]}{V_2} \rightarrow \underset{[-\text{high}]}{V_2}$$

As formalized in (21), Contraction will delete [a] only if it is short. In view of this, in order to obtain the correct outputs it is necessary that Shortening apply before Contraction and that Truncation apply before Shortening. In other words, the derivation must proceed as illustrated in (22):

(22)

	a.	b.
	malaa-a	malaa-i
Truncation	malaa	n.a.
Shortening	n.a.	mala-i
Contraction	n.a.	male

In applying the rules in (22), we follow the convention that rules apply in a linear order and that the input to a given rule R is the underlying sequence of sounds (the phonological analogue of the syntactic deep structure of chapter 2) as modified by all rules ordered before R. Thus, for example, if in derivation (22a) Shortening had applied before Truncation, we would have obtained the output *mala* with a word-final short [a] rather than the correct form with a long [aa].

This is not the only logically imaginable convention on rule application. For example, one might impose the convention that each rule applies to the underlying representation, rather than to the representation as modified by all rules ordered earlier than the rule in question. If this convention were imposed on the derivation (22b), Shortening and Contraction would produce, instead of [male], the output [malae]. Of course, we could generate the correct output [male] using this convention, but only at the cost of complicating the formulation of Shortening and Contraction. These complications would be gratuitous: they are forced upon us by the convention we are considering. We know that there is a simpler account of the facts, but this requires us to assume the convention on rule application exemplified in (22). Thus, unless we can establish the superiority of the alternative convention on rule application, we must adopt the convention of (22).

The forms in (23) show that in plural forms in Kasem, stem-final dorsal consonants [k g ŋ] are deleted, where [ŋ] represents the nasal cognate of [k g]. (The actual plural form of 'river' is [bwi], which is produced by the application of a rule of [u] desyllabification whose effect we will disregard here.)

(23)

	'river'	'room'	'song'	'leg'
Singular	bug-a	dig-a	laŋ-a	nag-a
Plural	bu-i	di	le	ne

We must therefore postulate the rule of Dorsal Consonant Deletion:

(24) *Dorsal Consonant Deletion*

V C V → V V
  |
[+DORSAL]

Condition: in the plural

In order to obtain the output [di] from [dig-i], it is necessary to order Dorsal Consonant Deletion before Truncation. The form [ne] from underlying [nag-i] shows that Dorsal Consonant Deletion must also precede Contraction. We have no evidence for the relative order of Shortening and Dorsal Consonant Deletion, but we know that both must precede Contraction.

Finally, consider these singular and plural forms from Kasem:

(25)

	'path'	'back'
Singular	kog-a	čog-a
Plural	kue	čue

(26)

	'sheep'
Singular	pi-a
Plural	pe

The underlying forms of the stems in (25) are [kaug] and [čaug], respectively. This leads us to postulate for the plural the underlying strings [kaug-i] and [čaug-i]. Since these are subject to Dorsal Consonant Deletion, we obtain [kau-i] and [čau-i], which would give us the incorrect [koi] and [coi] as outputs by the rules proposed so far. To obtain the correct outputs, we would need the representations [kua-i] and [čua-i], which can be derived from [kau-i] and [čau-i] by rule (27), Metathesis:

(27) *Metathesis*

$V_1V_2V_3 \rightarrow V_2V_1V_3$

Condition: V_2 and V_3 are distinct

The forms in (26) show that Metathesis must be restricted so as not to apply when the second vowel in the sequence is identical with the third. We then postulate [pia-a] and [pia-i] as the underlying representations for *pia* and *pe* and derive the correct outputs as follows:

(28)

	pia-a	pia-i
Dorsal Consonant Deletion	n.a.	n.a.
Metathesis	n.a.	pai-i
Truncation	pia	pai
Shortening	n.a.	n.a.
Contraction	n.a.	pe .

3.4.2 Velar Softening and [s]-Voicing in English

English has a large class of verbs composed of prefixes and stems of the sort illustrated in (29). (For additional discussion, see Halle and Mohanan 1985.)

(29)	im-pel	——	re-pel	com-pel	——	ex-pel
	in-fer	de-fer	re-fer	con-fer	pre-fer	——
	im-port	de-port	re-port	com-port	——	ex-port
	in-cur	——	re-cur	con-cur	——	——

Stem-initial [s] becomes [z] in cases where the prefix ends with a vowel and the stem [s] itself is followed by a vowel:

(30)	serve	con-serve	de-serve	re-serve	pre-serve
	sign	con-sign	de-sign	re-sign	——
	——	con-sent	——	re-sent	pre-sent

We will postulate that English is subject to the rule of [s]-Voicing, which is informally stated in (31):

(31) *[s]-Voicing*

[s] → [z] in the environment V ____ V

However, a number of stems do not undergo this rule. (Here again, recall that in the standard orthography of English a given letter does not represent a unique sound. In particular, in certain contexts (for instance, *call, gall*) the letters *c g* represent the sounds [k g], respectively, whereas in other contexts (for instance, *cider/peace, ginger/cage*) they represent the sounds [s ǰ]. In still other contexts (for instance, *delicious*) the letter *c* can also represent the sound [š]. Bear in mind that we are concerned with the sounds of the language rather than with their orthographic representations.)

(32)	cite	in-cite	ex-cite	——	re-cite	——
	cede	——	ex-ceed	con-cede	re-cede	——
	——	——	——	con-ceive	re-ceive	de-ceive

Irregularities of all sorts are commonly encountered in language. For instance, the inflection of the verb *be* in English and its equivalent in many other languages is totally unlike that of any other verb in the language and must be given by special rules applicable only to this verb. It might therefore appear at first sight that the irregularity in (32) is also to be accounted for by marking the stems *cite, cede,* and *ceive* as exceptions to [s]-Voicing. This proposal is likely to be incorrect, for it fails to bring out the fact that these very stems exhibit certain other apparent irregularities, all of which together have a single explanation.

A large part of the English vocabulary exhibits alternations between [k]–[s] and [g]–[ǰ] of the sort illustrated in (33):

(33)	electri[k]	–	electri[s]-ity	analo[g]	–	analo[ǰ]-ize
	Damas[k]us	–	Damas[s]-ene	collea[g]ue	–	colle[ǰ]-ial
	medi[k]	–	medi[s]-ine	tautolo[g]-ous	–	tautolo[ǰ]-y

This replacement of [k g] by [s ǰ], which traditionally has been referred to as Velar Softening, takes place before the syllable nuclei [ɪ ᴇ ay iy] (which are exemplified by the vowels in the words *ill, ell, aisle, eel*). The formal statement of this rule involves complexities that go beyond the scope of this chapter, but it can be stated informally in ordinary English as follows:

(34) *Velar Softening*

The dorsal stops [k] and [g] are replaced by [s] and [ǰ], respectively, before [ɪ ᴇ ay iy].

Since Velar Softening is part of the language, we now have a means of accounting for the absence of [s]-Voicing in the examples in (32). We postulate that in their underlying representations the stems in (32) begin with [k], which is turned into [s] by Velar Softening. If we assume further that Velar Softening is ordered after [s]-Voicing, we have explained the fact that the stems in (32) are not subject to [s]-Voicing. A modicum of support for this proposal comes from the fact that except for the verbs in (35), for which there is a special explanation, the stems of all verbs that are exceptions to [s]-Voicing have as their nucleus vowel [ɪ ᴇ ay iy]. If these stems were just irregular exceptions to [s]-Voicing, there would be no reason for them to be limited in this fashion.

The other class of exceptions to [s]-Voicing consists of stems that otherwise undergo the rule but fail to do so after certain prefixes such as *as-*:

(35)	sign	re-sign	as-sign
	——	re-sent	as-sent
	——	re-sume	as-sume

To account for this irregularity, we will assume that the prefix in the last column of (35) is subject to the special rule of Consonant Copy, which copies the initial consonant of the verb stem, applying only in verbs formed with this and a few other prefixes. As a consequence, verbs with these prefixes have phonological representations that are quite close to their representations in standard English orthography:

(36)	a-sign → a[s-s]ign	a-sent → a[s-s]ent	a-sume → a[s-s]ume
	a-fect → a[f-f]ect	a-cord → a[k-k]ord	a-tain → a[t-t]ain
	su-fuse → su[f-f]use	su-port → su[p-p]ort	su-round → su[r-r]ound

It is obvious that the representations in the first line of (36) cannot undergo [s]-Voicing, for [s]-Voicing applies only to a single [s] between vowels. The proposed representation thus accounts for the facts in (35). (Sequences of identical consonants in English are commonly pronounced like single consonants. Thus, [l] in the adverb *royally,* which is composed of the adjective *royal* and the adverbial suffix *-ly,* is pronounced exactly like the [l] in *royalist,* where the suffix added to *royal* begins with a vowel. To account formally for this fact, we must assume that English has a rule simplifying geminate (that is, double) consonants. It should be clear that this geminate simplification rule must be ordered after [s]-Voicing and Velar Softening; otherwise, the consonant created by Consonant Copy would be deleted, giving incorrect results.)

The forms in (37) provide additional evidence in favor of postulating the rule of Consonant Copy:

(37)	cede	pre-cede	con-cede	ac-[s]ede	suc-[s]eed
	——	——	con-gest	——	sug-[ǰ]est

Since the stem *cede* does not undergo [s]-Voicing in such verbs as *precede, proceed*, and *recede*, it must be represented underlyingly with a [k] that then undergoes Velar Softening, surfacing as [s]. We therefore postulate that the underlying representation for the verbs *accede* and *suggest* is as follows:

(38) a-[k]ede
su-[g]est

If we now postulate further that these verbs are subject to the rule of Consonant Copy, we obtain the representations in (39),

(39) a[k-k]ede
su[g-g]est

from which the required forms are readily derived by application of Velar Softening:

(40) su[g-ǰ]est
a[k-s]ede

3.5 On the Innateness of Linguistic Knowledge

We have examined the manner in which speech sounds are produced in the human vocal tract and the role that the sounds play in the expression of various regularities and rules. Two results of this investigation are particularly significant: (1) sounds are complexes of features that reflect phonetic capabilities of the independently movable portions of the human vocal tract, and (2) certain phonological regularities must be expressed by means of rules whose method of interaction is best characterized by applying them in a linear order. Since these are crucial components of the knowledge that speakers have of their language, we may ask how speakers acquire this highly recondite type of knowledge. Since we cannot plausibly suppose that children could acquire essential aspects of this knowledge—for instance, that rules must be applied in a linear order—in the ordinary course of growing up in a normal speaking family or other social unit, we have no alternative but to assume that the knowledge in question is innate, that is, available to humans (either at birth or at some later time) as part of their genetic endowment. (A similar argument concerning aspects of syntactic knowledge is presented in chapter 1.) On this view, the ability of humans to acquire the language of their community is more like their ability to learn to walk upright and rather unlike the ability that dogs, cats, and elephants have to acquire the same behavior—in those special cases where the latter have been trained to perform this feat. If the difference between humans and other species with respect to the ability to acquire bipedal gait is due to the different genetic endowment of the species in question, then the difference between humans and other species with respect to the ability to acquire language is also likely to be due to genetic factors. This hypothesis implies that most of what is needed for learning a language is already present in normal human beings and that to acquire command of a language speakers must (and can) obtain from a speech community information only about fairly restricted (peripheral) aspects of their language. Although we have a fair picture of what aspects of language can be learned and what aspects must be innate, much more information in this domain remains to be discovered.

Suggestions for Further Reading

Halle and Clements 1983 contains an elementary survey of most of the issues discussed in this chapter.

Good systematic introductions to phonetics may be found in Ladefoged 1975 and Lieberman 1977. Lieberman 1977 is especially to be recommended for its clear discussion of the relation between acoustic properties of speech and their articulatory implementation.

Extensive technical discussions of many issues in phonology are found in Chomsky and Halle 1968 and in Kenstowicz and Kisseberth 1979, although both books are now somewhat out of date. Van der Hulst and Smith 1982 contains more up-to-date treatments of many of the same issues but may be more difficult for the beginner because the papers in the collection represent heterogeneous theoretical positions.

Questions

3.1 English uses three distinct suffixes for the regular past tense forms of the verb: [ɪd], [t], [d]. Find the principle governing the choice of suffix by a particular verb. (See the discussion of English plurals in section 3.3.)

3.2 Give examples of English monosyllabic words containing a vowel that is

(i) [+high, −low, +back, +round]

(ii) [−high, −low, −back, −round, +ATR]

(iii) [−high, +low, +back]

3.3 What features are shared by the vowels in the following English words?

(i) full, fill

(ii) foal, cull

(iii) bother, father, moth

3.4 a. Give examples of English words ending with consonants that are

(i) [+nasal]

(ii) [+coronal] and [+strident]

b. State the other features of the nasal consonants in the words you've given in answer to question (ai).

3.5 What features are shared by the consonant sounds in the following English words? (In some of the examples a sequence of letters stands for a single sound.)

(i) pie, my, by

(ii) thigh, thy, sigh, xy(lophone)

(iii) key, tea, pea

3.6 Characterize in feature terms the differences among the [n] sounds in the following English words:

(i) tenth, tense, trench

3.7 Consider the following Latvian nominal declensions (three masculine, three feminine):

(i) Masculine	'father'	'swan'	'market'
Sg. Nom.	tæævs	gulbis	tirgus
Loc.	tæævaa	gulbii	tirguu
Acc.	tæævu	gulbi	tirgu
Dat.	tæævam	gulbim	tirgum
Gen.	tææva	gulbya	tirgus*
Pl. Nom.	tæævi	gulbyi	tirgi
Loc.	tæævuɔs	gulbyuɔs	tirguɔs
Acc.	tæævus	gulbyus	tirgus
Dat.	tææviæm	gulbyiæm	tirgiæm
Gen.	tæævu	gulbyu	tirgu

(ii) Feminine	'sister'	'land, earth'	'cow'
Sg. Nom.	maasa	zeme	guɔvs
Loc.	maasaa	zemee	guɔvii
Acc.	maasu	zemi	guɔvi
Dat.	maasay	zemey	guɔviy
Gen.	maasas	zemes	guɔvs*
Pl. Nom.	maasas	zemes	guɔvis
Loc.	maasaas	zemees	guɔviis
Acc.	maasas	zemes	guɔvis
Dat.	maasaam	zemeem	guɔviim
Gen.	maasu	zemyu	guɔvyu

*For the purposes of this exercise, ignore these two forms, which are exceptional.

a. Determine the underlying representation of each form, indicating morpheme divisions.

b. State the rules deriving the surface forms from the underlying forms.

3.8 Indefinite and definite adjectives in Latvian have separate declensions. These are given below for the root *lab-* 'good'.

(i)	M. Indef.	Fem. Indef.	M. Def.	Fem. Def.
Sg. Nom.	labs	laba	labays	labaa
Loc.	labaa	labaa	labayaa	labayaa
Acc.	labu	labu	labuɔ	labuɔ
Dat.	labam	labay	labayam	labayay
Gen.	laba	labas	labaa	labaas
Pl. Nom.	labi	labas	labiæ	labaas
Loc.	labuɔs	labaas	labayuɔs	labayaas
Acc.	labus	labas	labuɔs	labaas
Dat.	labiæm	labaam	labayiæm	labayaam
Gen.	labu	labu	labuɔ	labuɔ

In order to derive these forms, rules in addition to those postulated in question 3.7b will be required. State these rules and give the derivation of all the listed forms.

Questions 3.7 and 3.8 are adapted from M. Halle and G. N. Clements (1983). *Problem book in phonology*. Cambridge, MA: MIT Press.

References

Berko, J. (1958). The child's learning of English morphology. *Word* 14, 150–177.

Chomsky, N., and M. Halle (1968). *The sound pattern of English*. New York: Harper and Row.

Halle, M. (1978). Further thoughts on Kasem nominals. *Linguistic Analysis* 4, 167–185.

Halle, M. (1983). On distinctive features and their articulatory implementation. *Natural Language and Linguistic Theory* 1, 91–105.

Halle, M. (1985). Speculations about the representation of words in memory. In V. Fromkin, ed., *Phonetic linguistics*. Orlando, FL: Academic Press.

Halle, M., and G. N. Clements (1983). *Problem book in phonology*. Cambridge, MA: MIT Press.

Halle, M., and K. P. Mohanan (1985). Segmental phonology of Modern English. *Linguistic Inquiry* 16, 57–116.

Halle, M., and K. Stevens (1969). On the feature "advanced tongue root." In *RLE quarterly progress report* 94, MIT, Cambridge, MA.

Halle, M., and K. Stevens (1971). A note on laryngeal features. In *RLE quarterly progress report* 101, MIT, Cambridge, MA.

Jakobson, R. (1938). Observations sur le classement phonologique des consonnes. In *Selected writings*, vol. 1. The Hague and Berlin: Mouton.

Jakobson, R. (1948). The Russian conjugation. *Word* 4, 155–167.

Kenstowicz, M., and C. Kisseberth (1979). *Generative phonology: Description and theory*. New York: Academic Press.

Ladefoged, P. (1964). *A phonetic study of the West African languages*. Cambridge: Cambridge University Press.

Ladefoged, P. (1975). *A course in phonetics*. New York: Harcourt Brace Jovanovich.

Ladefoged, P., and I. Maddieson (1986). *Some of the sounds of the world's languages: Preliminary version*. UCLA Working Papers in Phonetics 64.

Lieberman, P. (1977). *Speech physiology and acoustic phonetics: An introduction*. New York: Collier-Macmillan.

Lightner, T. (1972). *Problems in the theory of phonology*. Edmonton, Alberta, Canada: Linguistic Research.

Sagey, E. W. (1986). The representation of features and relations in non-linear phonology. Doctoral dissertation, MIT, Cambridge, MA.

van der Hulst, H., and N. Smith, eds. (1982). *The structure of phonological representations*. Dordrecht, Holland: Foris.