I
Segments

1.1 Introduction

The fact that words, or more generally stretches of speech, can be divided up into individual segments, or speech-sounds, is familiar to speakers of languages. Thus speakers of English will generally agree that the word *bat* consists of the three sounds 'b', 'a' and 't'. They will further agree that the spelling system of English, i.e. its orthography, does not correspond in a one-to-one fashion to the 'sounds' of the language, so that a word such as *thatch*, although made up of six distinct orthographic symbols, contains only three, or perhaps four, sounds: 'th', 'a' and 'tch' (or perhaps 't' and 'ch'). This discrepancy means that phoneticians and phonologists require a system of transcription for the units of sound analogous to, but different from, that for the units of spelling. Various such systems have been proposed, and are familiar to the user of any dictionary giving the 'pronunciation' of the words of a language. In this book we will generally use the transcription system of the International Phonetic Association (IPA; see Appendix).

The transcription of the sounds of a word is not an entirely straightforward undertaking, and raises interesting theoretical questions in phonology. Thus the transcription of the English word *thatch* requires a decision (implicit or explicit) on the part of the compiler of the system as to whether the sequence *tch* represents two sounds, or **phonological segments** (specifically the two sounds found at the beginning of English *tore* /tɔː/ and *shore* /ʃɔː/),[1] or whether it is to be treated as a single sound, normally referred to as an **affricate**. In systems based on the IPA alphabet, the first option is taken, so that *chore* is represented phonemically as /tʃɔː/ and *thatch* as /θætʃ/, with *ch* or *tch* being

[1] In this book we will in general transcribe English words in the form in which they are realised in RP (Received Pronunciation), the prestige accent of British English. This is a matter of convenience; we are not thereby implying that RP has in any sense a privileged status in terms of its linguistic properties. We will, however, frequently consider other varieties where necessary; in particular we will have occasion to examine data from **rhotic** dialects, i.e. dialects in which postvocalic /r/ is pronounced. RP is non-rhotic, as evidenced by the realisations /tɔː/ and /ʃɔː/ for *tore* and *shore*; compare the pronunciations /toːr/ and /ʃoːr/ (or /toːɹ/ and /ʃoːɹ/) in a rhotic dialect such as Scots English.

represented as a sequence of /t/ and /ʃ/ (although the claim that /t/ and /ʃ/ are more closely related than a normal pair of segments can be indicated by the use of a ligature, as in /θæt͡ʃ/, or, more commonly, by combining the two symbols, as in /θæʧ/). In North American systems, however, such orthographic sequences are generally treated unambiguously as single segments, so that we find transcriptions such as /θæč/.

Notice that the concept of affricate illustrates not only that the relationship between sound and spelling is not entirely straightforward, but also, and perhaps more obviously of relevance for the phonologist, that the relationship between 'phonetic' and 'phonological' representation is also a matter of analysis. From a purely phonetic point of view, the nature of the relationship between the stop and the fricative in the final cluster of English *thatch* does not seem markedly different from that between the stop and the fricative in the final cluster of *hats*: in both cases we have a *phonetic* sequence of stop + fricative, [t̠ʃ] and [ts], respectively (we adopt the usual convention of giving phonetic representations in square brackets, and phonological ones between slant brackets; the line under [t] in [t̠ʃ] denotes retraction of the articulation, in this case to the postalveolar place of articulation of the [ʃ]). However, while the *tch* sequence is commonly treated as an affricate in phonological analysis, phonologists do not generally make a similar claim for the *ts* sequence of *hats*. On the other hand, the phonetically more or less identical cluster in German *Satz* [zats] 'sentence' *is* so treated.

The reasons for these differences (which we will not explore in any detail here) are thus phonological, rather than phonetic, although it is usually claimed that for something to be considered phonologically an affricate it must in any case have the phonetic property of **homorganicity**: i.e. the stop and the fricative must have the same place of articulation, so that [ts] (where both elements are alveolar) and [t̠ʃ] (where both elements are postalveolar) are both conceivable phonological affricates, while a sequence such as [ps] in English *cups* would not be. This claim is associated with the fact that it is just these homorganic sequences which may display a different distribution from 'normal' sequences of consonants. Affricates can generally occur *both* in syllable-initial position *and* in syllable-final position in a language, and thus violate the 'mirror-image' constraint on syllable structure.[2] This constraint states that a consonant cluster which can be syllable-initial in a language cannot be syllable-final, while the same cluster with its consonants in reverse order shows the opposite properties. English is typical in having initial /kl-/ and final /-lk/ (*class*, *sulk*), but not initial */lk-/ or final */-kl/ within a single

[2] We consider syllable structure in Chapter 3.

syllable. Contrast this with the distribution of affricates: /tʃ/ can be both initial and final in English (*chip* /tʃɪp/ and *pitch* /pɪtʃ/), as can /ts/ in German (*Ziel* /tsiːl/ 'goal' and *Satz*). On the other hand, the English sequence /ts/, like other stop + fricative sequences (e.g. /ps/, /ks/), occurs only in syllable-final position (and then almost exclusively as the result of morphological suffixation: e.g. *hats* = HAT + PLURAL).[3]

A full discussion of the status of affricates would take us much further. We return in §1.4 to the status of segments (or sequences) such as these, which exemplify the problem of dealing with what have been referred to as 'complex segments', and we will see that these phenomena have been the trigger for a great deal of interesting work in theories dealing with representation in phonology. Let us first, however, consider a rather more fundamental question regarding phonological representation: does the phonological segment have any internal structure? That is, is there anything which we can say about the way in which sounds behave by assuming some sort of internal structure which we could not say by having segments as the smallest phonological units?

1.2 Evidence for internal structure

It is not difficult to demonstrate that phonological segments in languages can be grouped together, in the sense that particular sets of segments may undergo what seems to be the same kind of **phonological process**. We are assuming here, fairly non-controversially, that it is reasonable to talk about phonological processes, in which a particular segment, or, more importantly here, a group of segments, is affected in some way. These may be either 'events' in the history of a language or relationships holding between the most abstract phonological representation of a segment or group of segments and its surface phonetic realisation.[4]

One such phonological process is that of **nasal place assimilation**, whereby a nasal consonant has the same place of articulation as a following obstruent (i.e. a stop, fricative or affricate). In English, for example, the effects of this process can be identified in various contexts, as in (1):[5]

[3] We indicate morphemes, i.e. minimal syntactic units, by the use of small capitals, as here.

[4] In the context of this book, however, we will beg the question of exactly what is meant by a surface 'phonetic' representation. For practical purposes, the 'surface' representations we consider will be fairly 'shallow' or 'concrete' *phonological* representations. Nevertheless, we will continue to refer to such representations as phonetic. More generally, as we noted in the Preface, we are assuming a model of phonology which is essentially derivational, in the tradition of Chomsky and Halle (1968). We do not adopt here the constraint-based model of Optimality Theory (see, e.g., McCarthy and Prince 1993; Prince and Smolensky 1993; Kager 1999). This is a matter of convenience, however, as we claim that much of what we have to say about the phonological representation of words is independent of whether we adopt a derivational or a constraint-based approach.

[5] The asterisks in (1c) denote that a sequence is ill formed.

(1) a. Edinburgh [ɛmbrə]
handbook [hæmbʊk]
b. unpopular [ʌmpɒpjələ]
unfair [ʌɱfɛə]
c. camber [kæmbə] *[kænbə] *[kæŋbə]
canter [kæntə] *[kæmtə] *[kæŋtə]
canker [kæŋkə] *[kæmkə] *[kænkə]

(1) shows examples of agreement in place of articulation between the nasal and the following obstruent. (1a, b) involve optional assimilations, particularly associated with fast-speech situations: realisations such as /ɛdɪnbʌrə/ and /ʌnpɒpjələ/, which do not show assimilation, also occur, of course. Those in (1b) can be analysed morphologically as involving a prefix ending underlyingly in the alveolar nasal /n/; e.g. UN + FAIR /ʌn + fɛə/. This analysis is supported by the fact that in such cases there are only two possible phonetic realisations of the nasal in the prefix: either as [n] or as the nasal which is homorganic with the following consonant. In addition, if there is no question of a possible assimilation, as in (2), where the following morpheme begins with a vowel or /h/, the only possible realisation is [n]:

(2) unequal [ʌniːkwəl]
unhappy [ʌnhæpɪ]

The forms in (1c) demonstrate a general constraint on English intervocalic clusters (at least those immediately following a stressed vowel within a single morpheme), which states that a sequence of nasal + stop must be homorganic. These differ from (1a, b), however, in that we are no longer dealing with cases in which, say, the labial nasal can be said to be *derived* from an alveolar nasal, as in [ɛmbrə] or [ʌmpɒpjələ] – there is no possibility of *camber* or *canker* occurring with /n/, as in *[kænbə] or *[kænkə], and there is no internal morphological structure which would lead us to suspect that these words have some kind of prefix CAN-.

Thus the process of nasal place assimilation is instantiated in various ways in English, and indeed in many other languages. However, our concerns here are not primarily with the status of the various different types of examples in the phonology of English; rather they focus on the characterisation of this type of process. In other words, how can we formalise the constraint represented in various ways by the data in (1)? Let us consider first (1a, b), in which we see that a cluster of /n/ followed by a stop may become homorganic in English. If the smallest available phonological units are complete segments, then we might represent the processes as in (3) (for the sake of simplicity, we ignore the case of nasals preceding /f/):

(3) a. /n/ → [m] / __ {/p/, /b/}
b. /n/ → [ŋ] / __ {/k/, /g/}

We use here a traditional **linear** type of notation for phonological rules:[6] the arrow denotes 'is realised as'; the underlying segment is given in slant brackets and its surface phonetic realisation in square brackets; the horizontal line denotes the environment in which the segment affected by the rule occurs, in this case preceding {/p/, /b/}; and the braces denote a set of segments. (3a), then, can be read as: 'Underlying /n/ is realised as phonetic [m] when it precedes either /p/ or /b/.'

There are various objections which can be raised with respect to the formulations of nasal place assimilation in (3). The common core of these objections is that the two parts do not look any more likely to be recurrent phonological rules than, say, any of the processes in (4), which are not likely to occur in any language:

(4) a. /n/ → [m] / __ {/k/, /g/}
b. /n/ → [ŋ] / __ {/p/, /b/}
c. /n/ → [m] / __ {/k/, /d/}
d. /n/ → [l] / __ {/t/, /d/}

Formally, the various rules in (4) are no more or less complex than those in (3), which express recurrent processes – surely an undesirable state of affairs. More particularly, the type of formulation in (3) and (4) is inadequate in two ways. In the first place, the formalism fails to relate the change characterised by a particular rule to the environment in which it occurs. Thus (4a), in which an alveolar nasal becomes labial in the environment of velar stops, is no more difficult to formulate than (3a), in which the same change takes place in the environment of labial stops. Yet (3a) is a natural process of assimilation, while (4a) is not. Secondly, the formalism does not show that the sets of consonants in the environments in (3a, b) are ones that we would expect to find triggering the same kind of change, whereas that in (4c), a set consisting of a voiceless velar stop and a voiced alveolar stop, would be most unlikely to be responsible for the change in (4c) (or, indeed, any other assimilation process). Again, though, (4c) is no more difficult to formulate than any of the other rules in (3) and (4).

This state of affairs clearly arises because we have neither isolated the phonetic properties which are shared by the set of segments involved in the process – nasality in the case of the input and the output (why should the output of (3a) be [m] rather than, say, [l]?); place of articulation in the

[6] See the Preface for a discussion of the difference between linear and non-linear approaches to phonological representation.

case of the output and the environment – nor incorporated them in our rule. In other words, we have failed to take account of the fact that it is the phonetic properties of segments which are responsible for their phonological behaviour, i.e. that phonological segments are not indivisible wholes, but are made up of properties, or, as they are usually referred to, **features**, which to a large extent correspond to the properties familiar from traditional phonetic description.

Furthermore, the fact that a change such as (4c) is an unlikely candidate for an assimilation rule shows that the class of segments triggering the process must share a particular property – in the case of (3a), for example, the property of labiality. A further examination of the phonologies of languages of the world would quickly show that a class of segments like this forms what is referred to as a **natural class**, i.e. a set of segments which *recurrently* participates as a class in phonological processes, such as the ones sketched above. Thus a set of segments which shares some phonetic property or combination of properties, to the exclusion of other sets of segments, forms a natural class.

Let us now identify a number of (ad hoc) phonological features which are relevant here, specifically [nasal], [labial], [alveolar] and [velar]. (Features are by convention enclosed in square brackets.)

We can use these features to write a general rule to characterise the assimilation processes illustrated by (3):

(5) a. $\begin{bmatrix}\text{nasal}\\ \text{alveolar}\end{bmatrix} \rightarrow \text{[labial] / \_\_ [labial]}$

b. $\begin{bmatrix}\text{nasal}\\ \text{alveolar}\end{bmatrix} \rightarrow \text{[velar] / \_\_ [velar]}$

However, we can formulate a rather more general statement about nasal place assimilation in English, which will also incorporate the data in (1c), in which there appears to be no reason to derive [m] and [ŋ] from an underlying /n/. This general statement about the class of nasals is given in (6):

(6) a. [nasal] → [labial] / __ [labial]
b. [nasal] → [alveolar] / __ [alveolar]
c. [nasal] → [velar] / __ [velar]

(6) successfully shows that the rule is a statement about a particular class of segments, nasals, characterised by a single feature which serves to distinguish the class from any other segments in the language. In other words, only nasals undergo the processes characterised by the rule, and no other segments in the language. Furthermore, it shows that the outputs and environments share a feature, namely the feature characterising place of articulation, which

makes just these processes more likely to occur than those in (4), for example. (6) is a non-arbitrary process, then.

Examples like these, which are typical of the way in which phonological processes operate in language, provide evidence for incorporating features in phonological description. It is with the nature of these features, and more particularly the question of whether they are organised in any way in the representation of segments, that we will be largely concerned in the remainder of this chapter.

However, at this point, let us note that the particular formulation in (6) will turn out to be far from adequate on a number of grounds, which do not, however, affect the validity of the points just made. Let us consider here just two of the problems.

(6) appears to consist of three sub-processes, whereas, as we have seen, nasal place assimilation is a single process in English. In traditional linear phonology, it is usual to 'collapse' rules like those in (6), all of which share the same input, to give (7):

(7) $\text{[nasal]} \rightarrow \begin{Bmatrix}\text{[labial] / \_\_ [labial]}\\ \text{[alveolar] / \_\_ [alveolar]}\\ \text{[velar] / \_\_ [velar]}\end{Bmatrix}$

The three expressions contained in braces are to be seen as alternatives; i.e. nasals are labial before labials, alveolar before alveolars and velar before velars. Thus the 'shared' part of the rule – the input – is mentioned only once.[7]

However, conventions such as that used in (7) still permit the collapse of unrelated rules, as well as rules which apparently belong together. Thus some languages have a rule whereby a nasal consonant becomes voiceless preceding a voiceless (aspirated) consonant. In some dialects of Icelandic, for example, *hempa* /hɛmp^{h}a/ 'cassock' is realised as [hɛm̥pa], with devoicing of the /m/. There seems to be no formal reason why the rule characterising this process cannot be collapsed with (7), especially as Icelandic also has nasal place assimilation processes:

(8) $\text{[nasal]} \rightarrow \begin{Bmatrix}\text{[labial] / \_\_ [labial]}\\ \text{[alveolar] / \_\_ [alveolar]}\\ \text{[velar] / \_\_ [velar]}\\ \text{[voiceless] / \_\_ [voiceless]}\end{Bmatrix}$

In other words, we have still failed to show that the features involved in the nasal assimilation process, i.e. [labial], [alveolar] and [velar], are related to

[7] A fuller formulation of the rule in question would also involve reference to other features; we ignore this here, as before.

each other in some way, i.e. that they characterise place of articulation, whereas [voiceless] is not related to any of the other three in this way.

A second problem is that, merely by incorporating features in our rules, rather than the segments of (3) and (4), we have not removed the possibility of formulating what are sometimes referred to as 'crazy rules'. Thus (9) is as easy to formulate as (7):

(9)
$$[\text{nasal}] \rightarrow \left\{ \begin{array}{l} [\text{labial}] \,/\, \_\_ \,[\text{alveolar}] \\ [\text{alveolar}] \,/\, \_\_ \,[\text{velar}] \\ [\text{velar}] \,/\, \_\_ \,[\text{labial}] \end{array} \right\}$$

Underlying these criticisms of the formal conventions of linear phonology is the belief that a phonological theory should be as restrictive as possible, in the sense that an ideal system should be able to represent only phonologically natural events and states, and should not be able to characterise unnatural events such as (4) or (9). This belief underpins many **non-linear** alternatives to the formulations above, alternatives which we will begin to consider in §1.4. For the moment, however, we turn in greater detail to the nature of the features which will be required in phonology.

1.3 Phonological features

The idea that segments are made up of phonological features has a long tradition, and received its first comprehensive formalisation in Jakobson *et al.* (1951). The most widely known system is that proposed by Chomsky and Halle (1968; henceforth *SPE*), which differs from the Jakobsonian model in a number of respects, most notably in that the later features are based entirely on articulatory parameters, whereas those of Jakobson *et al.* were defined primarily in terms of acoustic properties. A second important difference involves the fact that many of the Jakobsonian features were relevant to the description and characterisation of both vowels and consonants, while the *SPE* system used largely separate sets of features. Feature theory is not unique to linear approaches to phonology; indeed, much work within non-linear phonology adopts the set of features proposed in the linear framework of *SPE*. However, non-linear phonology typically differs from linear accounts of the segment in incorporating a greater degree of internal structure than a simple list of features, as we shall demonstrate later in this chapter.

As there is a great deal of discussion of individual features available in the literature (e.g. Kenstowicz and Kisseberth 1979; Lass 1984a: chs. 5–6; Keating 1988a; Clements and Hume 1995), we shall not attempt to provide a comprehensive account of the features which would be required to characterise the segments making up the phonological system of English, for example. Rather,

we shall introduce individual features as and when they become relevant, and only provide extensive discussion when necessary. Here the focus will be on how features interact in the representation of the segment, and in particular on the degree of structure required.

In the linear model of *SPE*, segments were viewed as consisting simply of an unordered list of binary features, which were established on grounds similar to those discussed above, i.e. the potential of a feature to define a natural class of segments. The features characterising a segment were organised into a **feature-matrix** in which the features were simply listed along with their value (either + or –) for the segment in question; thus the feature-matrix for the English vowel /iː/, for example, contains the following features, among others:

(10)
$$\begin{bmatrix} \text{+sonorant} \\ \text{−consonantal} \\ \text{+continuant} \\ \text{+voice} \\ \text{+high} \\ \text{−low} \\ \text{−back} \\ \text{−round} \end{bmatrix}$$

Within recent non-linear phonology, in which a more elaborate internal structure has been assigned to the segment, it has become customary to use a different type of formalism to represent the segment. We return in §§1.3.1 and 1.3.5 to the kind of motivation that can be adduced for suggesting a greater degree of structure than is embodied in (10); however, to facilitate comparison, we take the opportunity at this point of providing a 'non-linear' equivalent of (10), in which all of the features making up the segment are ASSOCIATED to a single segmental NODE, represented in (11) by '∘':

(11)

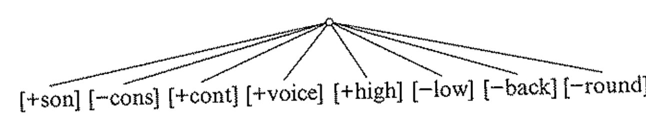

This node is generally referred to as the ROOT NODE – see §1.4.

In (11), as in (10), the features are unordered with respect to each other; any change in this ordering (vertical in the case of the feature-matrix in (10), horizontal in the case of the feature 'tree' in (11)) does not in this case yield anything different from the segment /iː/. We return in due course to the different claims made by the formalisms; in the meantime we devote a little space to the features themselves.

1.3.1 Major class features

The first two features in the matrix in (10) give the 'major class' to which the segment belongs, i.e. vowel; vowels are non-consonantal and, like liquids (i.e. *l* and *r* sounds) and nasals, they are sonorant. In the *SPE* model, sonorancy was defined in articulatory terms, as involving 'a vocal tract configuration in which spontaneous voicing is possible' (*SPE*: 302), but an acoustic definition is equally plausible: sonorant segments have relatively more periodic acoustic energy than non-sonorants (cf. Lass 1984a: 83). By characterising vowels, liquids and nasals as sharing the feature-value [+sonorant], of course, we are making the claim that they form a natural class (cf. §1.2), i.e. that there are phonological processes affecting just this group of segments, and no others. Equally, by assigning the value [–sonorant] to a particular group of segments (the class normally referred to as obstruents, made up of stops, fricatives and affricates), we are claiming that this group too should function as a class. It is not difficult to find processes to demonstrate this; thus the class of obstruents is typically the only class to display 'final devoicing' in many languages, as in various Scottish dialects of English, and Dutch, from which the examples in (12) are taken:

(12)	*singular*				*plural*		
a.	rib	'rib'	/rɪb/	[rɪp]	ribben	/rɪbən/	[rɪbə]
	bed	'bed'	/bɛd/	[bɛt]	bedden	/bɛdən/	[bɛdə]
b.	lip	'lip'	/lɪp/	[lɪp]	lippen	/lɪpən/	[lɪpə]
	kat	'cat'	/kɑt/	[kɑt]	katten	/kɑtən/	[kɑtə]
	nek	'neck'	/nɛk/	[nɛk]	nekken	/nɛkən/	[nɛkə]
c.	kam	'comb'	/kɑm/	[kɑm]	kammen	/kɑmən/	[kɑmə]
	man	'man'	/mɑn/	[mɑn]	mannen	/mɑnən/	[mɑnə]
	ring	'ring'	/rɪŋ/	[rɪŋ]	ringen	/rɪŋən/	[rɪŋə]
	nar	'fool'	/nɑr/	[nɑr]	narren	/nɑrən/	[nɑrə]
	bel	'bell'	/bɛl/	[bɛl]	bellen	/bɛlən/	[bɛlə]

The obstruents in the singular forms of (12a, b), which are syllable-final, must be voiceless, irrespective of whether they are voiced (12a) or voiceless (12b) in other contexts, such as in the plural forms, where they occur intervocalically. Because the obstruents in (12a) are voiced in other contexts, we assume that they are phonologically, i.e. underlyingly, voiced. In other words, we ascribe their voicelessness in (12a) to the environment in which they occur, i.e. syllable-final position.[8]

[8] Notice that if we had assumed that the obstruents in (12a) were underlyingly voiceless, rather than voiced, we would not have been able to predict whether they would surface intervocalically as voiced (as in *bedden*) or voiceless (as in *katten*). However, it should not be thought that a state of affairs in which an underlying voiceless obstruent becomes voiced intervocalically in a language is impossible; indeed, intervocalic voicing is a very common process.

The liquids and nasals in (12c), however, remain voiced in all contexts. Thus the rule of final devoicing in Dutch must make reference to the natural class of non-sonorant consonants, and can be formulated as (13):

(13) *Dutch Final Devoicing*

$[-\text{son}] \rightarrow [-\text{voice}] \ / \ \_\_ \ ]_\sigma$

(where we use $]_\sigma$ to denote 'end of syllable').

We can also find cases in Dutch in which [+sonorant] functions as a natural class. Dutch has a process of diminutive formation in which the diminutive suffix is added to a noun. The suffix has a number of different allomorphs, illustrated in (14):

(14)	*noun*				*diminutive*	
a.	nek	'neck'	/nɛk/	[nɛk]	nekje	[nɛkjə]
	pruik	'wig'	/prœyk/	[prœyk]	pruikje	[prœykjə]
b.	kam	'comb'	/kɑm/	[kɑm]	kammetje	[kɑmətjə]
c.	pruim	'plum'	/prœym/	[prœym]	pruimpje	[prœympjə]
	boon	'bean'	/boːn/	[boːn]	boontje	[boːntjə]
	haring	'herring'	/haːrɪŋ/	[haːrɪŋ]	harinkje	[haːrɪŋkjə]
	beer	'bear'	/beːr/	[beːr]	beertje	[beːrtjə]
	uil	'owl'	/œyl/	[œyl]	uiltje	[œyltjə]
d.	ui	'onion'	/œy/	[œy]	uitje	[œytjə]

The form of the allomorph of the diminutive suffix is predictable according to the phonological form of the noun to which it is attached. Crucially for our purposes here, it takes the form [[stop] + jə] only if the preceding segment is [+sonorant] (a consonant in (14c), a vowel in (d)). (Notice that in (c) the stop assimilates in place to the preceding consonant; the difference between the forms of the suffix in (b) and (c) is due to the nature of the segments preceding the final liquid or nasal.) Thus [+sonorant], just like [–sonorant], can function to identify a natural class of segments.[9]

This phenomenon also provides evidence that natural classes can be defined by a combination of two or more features. (14b), for example, is representative of the larger class in (15):

(15)	*noun*				*diminutive*	
	kam	'comb'	/kɑm/	[kɑm]	kammetje	[kɑmətjə]
	man	'man'	/mɑn/	[mɑn]	mannetje	[mɑnətjə]
	ring	'ring'	/rɪŋ/	[rɪŋ]	ringetje	[rɪŋətjə]
	nar	'fool'	/nɑr/	[nɑr]	narretje	[nɑrətjə]
	bel	'bell'	/bɛl/	[bɛl]	belletje	[bɛlətjə]

[9] In this discussion we are making no assumptions about the underlying form of the diminutive suffix in Dutch, which has been an issue of some debate (see, e.g., Ewen 1978; Trommelen 1983; van der Hulst 1984; Booij 1995 for discussion of diminutive formation). The validity of the particular argument here depends on the assumption that the allomorphs [tjə], [pjə] and [kjə] are derived, rather than underlying.

The class of segments which determine the choice of the [-ətjə] suffix, i.e. the class of nasals and liquids, is defined by the feature combination [+sonorant, +consonantal], together with the nature of the preceding vowel.

Similar evidence from natural class behaviour can be cited in the justification of the various features which we identify in what follows. However, we will only consider such evidence when it is of particular interest for the point we are making.

The next two features in (10), [continuant] and [voice], are used to make further distinctions among the various major classes (vowels, liquids, nasals, obstruents). [+continuant] sounds differ from [−continuant] sounds in not having a complete closure in the oral tract. In the obstruent category ([−sonorant]), fricatives (e.g. /f v ʃ ʒ x ɣ χ/) are [+continuant], in that the stricture of close approximation does not block the airstream entirely, while stops (/p b t d k g q/, etc.) are [−continuant]. Similarly, within the [+sonorant, +consonantal] category, nasals (/m n ŋ/) are [−continuant], in that there is again a complete obstruction of the airstream in the oral cavity (although air does of course escape through the nasal cavity, so that nasal stops can be prolonged), while liquids are [+continuant].[10]

As might be expected, [+voice] sounds are those produced with vibration of the vocal cords; [−voice] are those with no such vibration.

In this discussion, we have implicitly assumed a grouping of features ([consonantal] with [sonorant], which together define 'major classes'; [voice] with [continuant], involved in characterising 'manner of articulation') which is in no way reflected in the matrix in (10). Indeed, the internal structure of the feature-matrix seems quite irrelevant – for example, as we have seen, changing the order in which the features occur in the matrix does not yield a segment which is different in any way. There is, however, a great deal of evidence that grouping of this kind is phonologically relevant: the sets of segments characterised by combinations of particular values of the features within these 'groups' are typically – and recurrently – appealed to in phonological

[10] Notice, though, that it has been claimed that lateral liquids are [−continuant], because, although there is a stricture of open approximation at the sides of the tongue, they also display complete *central* closure. This claim is given weight by the fact that there appear to be processes in some languages in which the lateral liquid forms a natural class with the nasals, as opposed to the non-lateral liquid. Thus Ó Dochartaigh (1978) notes that in some dialects of Scottish Gaelic short vowels diphthongise before /l n/, but lengthen before /r/. Similarly, Clements (1989) demonstrates that in some rhotic dialects of English words like *prince* and *false* may be realised with an epenthetic or inserted [t], i.e. as [prɪn^{t}s] and [fɔːl^{t}s]. However, between an /r/ and an /s/, as in *nurse* ([nʌrs]), insertion is not possible, again showing that the lateral forms a natural class with the nasal, rather than with the non-lateral liquid. Behaviour like this would support the point of view that laterals may be [−continuant]. In other processes, however, they clearly form a natural class with the non-lateral liquids (e.g. /r/), and so appear to be [+continuant]. We will not be concerned here with how our feature system should capture this apparent anomaly.

processes. Consider again the features [consonantal] and [sonorant], which, it will be recalled, divide up the 'major classes' of segments as in (16):[11]

(16)

	O	N/L	V
[son]	−	+	+
[cons]	+	+	−

('O' is obstruent, 'N' nasal, 'L' liquid, 'V' vowel). The interaction of these two features is relevant to a number of phonological phenomena. In other words, various combinations of the two features define classes which occur frequently in phonological processes. In addition, the ordering of elements within a syllable is typically determined by these features, so that a vowel ([+sonorant, −consonantal]) forms the peak of a syllable, and an obstruent ([−sonorant, +consonantal]) the margin, with any liquids or nasals ([+sonorant, +consonantal]) being intermediate. Thus /prɪns/ is a well-formed English syllable, while */rpɪsn/ is not.

It is often claimed that the features [sonorant] and [consonantal] determine a **sonority hierarchy** (or sonority scale), and that this hierarchy is reflected in the behaviour of segments in the syllable: the higher the sonority of a segment, the closer it is to the peak of the syllable (see, e.g., Vennemann 1972; Hooper 1976; Kiparsky 1981; Clements 1990). Such hierarchies have a more widespread role, and sometimes involve the other two features already discussed, [continuant] and [voice]. Although we will not discuss this in detail at this point, [+continuant] segments are higher on the sonority hierarchy than [−continuant], and [+voice] segments are higher than [−voice]. This can be established with respect to processes such as the historical 'weakening' or lenition of stops to sonorants in intervocalic position, which involves the gradual assimilation of the features of the stop to those of the surrounding vowels, as illustrated by the development from Pre-Old English to Modern English of the word *own* (from Lass and Anderson 1975: 158):

(17) Pre-OE *[aagan] > OE [aaɣan] > ME [ɔɔwən] > lME [ɔɔn] > MdE /oon/ (/oʊn/, /əʊn/, etc.) 'own'[12]

Each of the changes in (17) represents a step along a lenition hierarchy, which for velars in intervocalic position involves the steps in (18):

[11] It will be observed that the combination [−sonorant, −consonantal] is also formally possible. Given the definition of [sonorant], it is difficult to see what class of segments might be assigned this representation; [−consonantal] segments (vowels) appear to be inherently sonorant. Chomsky and Halle (1968) in fact assign the combination to [ʔ h], but it is not clear from the phonological behaviour of this set of segments that they should be treated as non-consonantal. We consider an alternative account of [h] in §1.3.5; we will assume here that there are no [−sonorant, −consonantal] segments.

[12] OE = Old English; ME = Middle English; lME = late Middle English; MdE = Modern English.

(18) k > (x *or* g) > ɣ > w

On the basis of processes such as these, Lass and Anderson (1975: 150) establish the general lenition hierarchy in (19):[13]

(19)

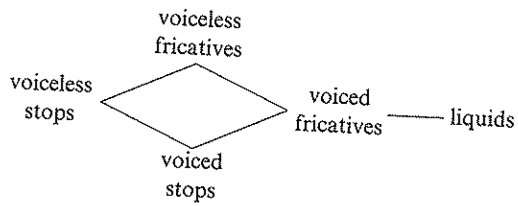

Notice that nasals typically do not participate in intervocalic lenition processes, for reasons that need not concern us at present; however, their status with respect to sonority within the class of sonorant consonants can be established on the basis of their behaviour in syllable structure: liquids ([+continuant]) are closer to the syllabic element than nasals ([–continuant]), as is evidenced by syllables such as English *kiln* and *barn* (/bɑrn/ in postvocalic *r*-pronouncing (i.e. rhotic) dialects of English), as opposed to the unacceptable syllables */kɪnl/ and */bɑnr/.

We have now shown that the four features considered so far together form a group with respect to which phonological regularities can be uncovered, as they, and they alone, distinguish the classes involved in sonority-based phenomena, as in (20):

(20)

	voiceless stops	voiced stops	voiceless frics	voiced frics	N	L	V
[son]	–	–	–	–	+	+	+
[cons]	+	+	+	+	+	+	–
[cont]	–	–	+	+	–	+	+
[voice]	–	+	–	+	+	+	+

(Here we ignore oppositions between voiced and voiceless sonorants, i.e. nasals, liquids and vowels.)

1.3.2 *Vowel features*

Let us now consider the remaining four features in (10), [high], [low], [back] and [round]. In *SPE*, these features, used primarily to distinguish the vowels of a language, are defined in terms of the position of the highest point of the tongue in the production of a vowel (for [high], [low] and [back]), and the

[13] Lass and Anderson in fact claim that lenition ultimately yields deletion, as illustrated by the development of *own* in (17), possibly via a vowel stage. For the moment, however, we confine the discussion to the consonant-types involved in such processes.

presence or absence of lip-rounding (for [round]). The definitions of the first three features refer to a 'neutral' position for the tongue (roughly the position for [ɛ]), such that [+high] sounds have their closest constriction higher than the neutral position, whereas [–high] sounds do not, and similarly for [+low] vs [–low] and [+back] vs [–back].

A system of this sort essentially treats the vowel features as interpretations of two axes, as in (21):

(21)

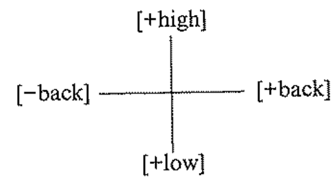

The claim inherent in the set of features given above is that languages typically make a three-way opposition on the vertical axis, but only a two-way opposition on the horizontal axis. Thus there is no separate feature [±front] beside [±back], but we do have a separate feature [±low] alongside [±high]. The only way of characterising a three-way opposition on the high–low axis within a binary feature framework is to postulate two features, giving the following possibilities:

(22)

	[–back]		[+back]	
	[–round]	[+round]	[–round]	[+round]
[+high, –low]	i	y	ɯ	u
[–high, –low]	e,ɛ	ø,œ	ɤ,ʌ	o,ɔ
[–high, +low]	a	ɶ	ɑ	ɒ

(For illustration we use here the Cardinal Vowel symbols (originally proposed by Daniel Jones; see e.g. Abercrombie 1967: ch. 10), rather than the vowels of any particular language.) Notice that the definitions of the features exclude the fourth logically possible combination of the two features defining the high–low axis, i.e. [+high, +low].

It is clear that other features will be required to characterise the vowel space, seeing that those discussed so far apparently fail to distinguish between various pairs of [–high, –low] vowels, for example. This is an area which has been the subject of major rethinking since the publication of *SPE*, and we devote some space here to a discussion of the issues involved.

1.3.3 *The vowel-height dimension and related issues*

There have been a number of proposals for distinguishing the various pairs of [–high, –low] vowels in (22). These proposals can be divided into three major groups: (i) those which distinguish the members of each pair by means of a binary feature [tense]; (ii) those which try to reflect the difference in height 'directly'; and (iii) those which introduce a feature [advanced tongue root] to distinguish the members of the various pairs.

The first proposal is found in *SPE*, in which it is argued that the difference is one of tense vs lax, with the tense vowel of each pair being 'executed with a greater deviation from the neutral or rest position' than its lax counterpart, so that 'the greater articulatory effort in the tense vowels is further manifested by their greater distinctiveness and the markedly longer duration during which the articulatory configuration remains stationary' (*SPE*: 324–5). Such a distinction is applicable to certain pairs of vowels which we have not yet considered, such as /iː/ vs /ɪ/ in RP English (e.g. *meal* /miːl/ vs *mill* /mɪl/), or in the German pairs given by Chomsky and Halle (e.g. *ihre* [īrə] 'her' vs *irre* [irə] 'err', or the similar *Huhne* [hūnə] 'chicken' vs *Hunne* [hunə] 'Hun'; notice that Chomsky and Halle distinguish tense vowels from their lax counterparts by means of a macron over the tense vowel).

This approach is also readily applicable to the [–high, –low] vowels in a system such as RP, in which the opposition between the members of each pair is again not just one of tongue-height (quality), but also of length (quantity). Thus the distinction between the two vowels in *beat* /iː/ and *bit* /ɪ/ is one which can be interpreted as tense vs lax, as is that between *mane* /eɪ/ and *men* /ɛ/, as well as in *coat* /əʊ/ vs *cot* /ɒ/.[14] The crucial claim that is being made here, then, is that the *type* of phonological opposition between, say, the two [–high, –low] vowels ([e] and [ɛ]) is different from that holding between, say, the high vowel [i] and the higher of the two mid vowels [e], or between the lower of the two mid vowels [ɛ] and the low vowel [a].

As we have seen, this account seems appropriate for a system like RP, but it has encountered criticism from those who believe that there are vowel systems in which it is reasonable to speak of (at least) *four* distinct vowel heights. The front-vowel system in (23), for example, is that of some dialects of Scots English:

(23)	*beat*	[bit]
	bit	[bɪt]
	bait	[bet]
	bet	[bɛt]
	bat	[bat]

[14] We are here following Chomsky and Halle's position that the phonological distinction between these pairs is one of quality, not length. Notice too that the diphthongs of *main* and *coat* are tense, just like the monophthong of *beat*.

Here the various vowels are apparently distinguished *only* by vowel height, with no apparent difference in length, and any appeal to the notion of tense vs lax, as defined by Chomsky and Halle, seems inappropriate.

The existence of systems like this has led some phonologists to propose systems which reflect the height dimension more directly. Thus Wang (1968) replaces the feature [low] by [mid], which allows the expression of four heights, rather than the three of *SPE*:

(24) $\begin{bmatrix}\text{+high}\\ \text{–mid}\end{bmatrix}$ $\begin{bmatrix}\text{+high}\\ \text{+mid}\end{bmatrix}$ $\begin{bmatrix}\text{–high}\\ \text{+mid}\end{bmatrix}$ $\begin{bmatrix}\text{–high}\\ \text{–mid}\end{bmatrix}$

/i/ /e/ /ɛ/ /æ/

Such a formulation certainly allows the expression of four heights, but notice that there is still a fundamental problem associated with the expression of vowel height by means of binary features: the fact that we have to use (at least) two binary features to express what appears to be a *single* phonetic dimension. That is, a sequence such as /i/–/e/–/ɛ/–/æ/ can be seen as a set of points on a single scale, and this has led some phonologists, e.g. Ladefoged (1975) and Williamson (1977), to abandon binary features for the expression of vowel height, and to introduce a **multivalued scalar** feature, as in (25):

(25)	/i/	[4 high]
	/e/	[3 high]
	/ɛ/	[2 high]
	/æ/	[1 high]

We do not pursue this approach at this point (but see our discussion of features in §2.1). Rather, we now consider the third type of proposal that has been made in this area, the introduction of a feature [advanced tongue root] (henceforth [ATR]).

It has been observed that one of the articulatory correlates of the difference between a pair of vowels such as [e] and [ɛ] typically involves the position of the tongue root: for [e] the tongue root is further forward, while for [ɛ] it is further retracted. A similar relationship holds between [i] and [ɪ]. Thus the difference between the two vowels does not relate exclusively to the relative height of the *body* of the tongue, as is suggested by (25), but involves additional phonetic parameters.

Clearly, the choice amongst the three alternatives just outlined – we refer to them here as the tense/lax, height and ATR approaches – depends on whether we can show that one of them more successfully predicts what actually happens in phonological systems and processes than the others. That is, if we find processes which show that the relationship between [i] and [e] is phonologically the same as that between [e] and [ɛ], this would provide evidence in favour of a multivalued feature [high], as in (24).

In fact, we believe that all three approaches are required in phonological theory. That is to say, we believe that vowel systems may be organised along any one of the three lines suggested by the approaches just discussed, so that the nature of the phonetic parameters which play a role in a particular sound-system is reflected in the phonology of the language in question.

Let us illustrate this with a further consideration of [tense] and [ATR]. We have already seen that the feature [tense] plays a role in the phonology of RP, for example. Thus /ɪ/ in *bit* and /ʊ/ in *look* are the [−tense] counterparts of /iː/ in *beat* and /uː/ in *Luke*, with which they are otherwise identical in terms of their feature make-up. The fact that the [−tense] vowels of RP form a class (which includes, besides /ɪ ʊ/, also /ɛ æ ʌ ɒ ə/) is shown by the fact that just this set of vowels cannot occur in final position in a stressed syllable, while the set of [+tense] vowels can (cf. /biː/ *bee* and */bɪ/, for example).[15] Similarly, they can occur before /ŋ/, while the [+tense] vowels cannot (e.g. *bang* /bæŋ/, but **boong* /buːŋ/).[16] Thus a feature such as [tense] is required in the analysis of systems such as RP. Crucially, at least with respect to the oppositions between /iː/ and /ɪ/ and between /uː/ and /ʊ/, the vowel system is organised in terms of 'central' vs 'peripheral' vowels. (26) gives the representation of a ten-vowel system such as this, consisting of a peripheral ([+tense]) set /i u e o ɑ/ and a central ([−tense]) set /ɪ ʊ ɛ ɔ ə/. Notice that for the low vowels peripherality is often manifested as greater pharyngeal constriction, so that peripheral /ɑ/ is considerably retracted:

(26)

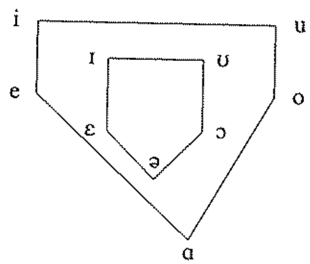

[15] It is also possible to deal with these restrictions in terms of vowel-length, rather than the apparently qualitative distinction of tense vs lax. We discuss vowel-length in Chapter 3.

[16] The RP vowel /ʊ/ does not in fact occur before /ŋ/ in the native vocabulary of English. However, in loanwords such as *Jung* from German, we find /ʊ/, not /uː/ (cf. Collins and Mees 1996: 97), which further demonstrates the validity of the analysis.
The status of /ə/ is more problematical. It again fails to occur before /ŋ/, and shows other phonological behaviour which suggests that it forms a set of its own in some respects. However, given that it patterns with the lax vowels in not taking stress in final position in a syllable, we feel justified in categorising it in this set.

Such an analysis is proposed for Classical Latin by Allen (1973: 132), who observes: 'the tenseness is . . . responsible for the long vowels occupying a larger, more "centrifugal" perimeter of articulations'. He represents the system of Latin as in (27), where the long tense vowels are represented with a macron (¯), the short lax vowels with a breve (˘):

(27)

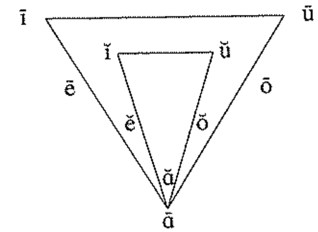

Other vowel systems, however, are organised quite differently, even though they may contain more or less the same vowels as (26). In particular, many languages divide the set of vowels not into a tense and a lax subset as in (26), but rather into two subsets according to tongue-root position, i.e. into one [+ATR] set and one [−ATR] set, as in (28):

(28)

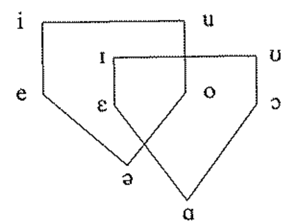

Here the vowels are grouped into /i u e o ə/ ([+ATR]) and /ɪ ʊ ɛ ɔ ɑ/ ([−ATR]).[17] The evidence that there are ten-vowel systems which are organised in this way comes from processes involving **vowel harmony**, i.e. processes in which all the vowels within a particular domain, often the word, must have the same value for a particular phonological feature. One such system is that of the Asante dialect of Akan, a language spoken in Ghana, in which all the vowels within a single word must have the same value for [ATR] (see Stewart 1967, 1983; Clements 1981). Stewart (1967: 186) gives the following examples:

[17] The relation between the symbols used for the non-low back vowels in (28) is to some extent arbitrary, although conventional. Thus /ʊ/ is used in (28) for a vowel which is apparently more peripheral on the front–back dimension than /u/, whereas in (26) this relationship is reversed. The opposite choice would be equally arbitrary, however, in that the same discrepancy would hold, but on the high–low dimension, rather than the front–back.

(29)

[–ATR]		[+ATR]	
/wʊbɛnʊmʔ/	'you will drink it'	/wubenumʔ/	'you will suck it'
/ɔːbɛtʊʔ/	'it is going to lay'	/oːbetu/	'he is going to pull it out'
/mɪkjɪrɛ/	'I show'	/mitie/	'I listen'

The forms in (29), involving the non-low vowels of Akan,[18] show the harmony process in operation: all the vowels in the left-hand column are [–ATR], while those in the right-hand column are [+ATR]. Forms with a mixture of [–ATR] and [+ATR] vowels, such as */wʊbenumʔ/, are ill formed.

We do not at this point pursue the question of how such harmony processes are to be analysed, an issue which we return to in some detail in §1.4.2 in relation to Turkish; however, the phenomena just outlined provide us with sufficient reason to claim that systems such as (26) and (28) are both to be found in languages of the world, and hence that the tense/lax and ATR features must *both* form part of our feature system.

We can also find evidence to support the point of view that the vowels of languages may be organised in terms of relative height, as suggested by the Scots English data in (23). Such evidence can be adduced from processes in which some or all of the vowels of a language move one 'step' up or down. For example, the effect of the English Great Vowel Shift was to move non-high long vowels up one step, giving the changes in (30) (from Lass 1987: 130), which shows various Middle English and early Modern English forms (*c.* 1600):

(30)

	ME		1600
beet	eː	>	iː
beat	ɛː	>	eː
mate	aː	>	ɛː
boot	oː	>	uː
boat	ɔː	>	oː

There are also processes which make appeal to the notion 'one step lower'. Lindau (1978: 545) observes that in Scanian, a Swedish dialect spoken in Malmö, there is a process in which the diphthongisation of a long vowel yields a first element which is one step lower than the original monophthong:

[18] As is typical in [ATR] harmony systems, the low vowels of Akan behave rather differently from the non-low ones, for reasons which need not concern us at this point.

(31)

/iː/ → [ei]	/yː/ → [øy]	/uː/ → [eu]
/eː/ → [ɛe]	/øː/ → [œø]	/oː/ → [ɛo]
/ɛː/ → [æɛ]		

The range of processes surveyed in this section suggests that vowel systems can be organised along different phonetic and phonological parameters, and hence that our feature system must be rich enough to be able to describe all of the parameters found to play a role in the organisation of vowel systems.[19]

1.3.4 Consonantal features

The fact that the set of features we have now isolated forms an intuitively obvious group hardly needs confirmation: there are clearly many phonological processes which make reference to particular subsets of vowels.

Much the same observation applies to the set of features defining articulatory place for consonants, which were not included in the feature-matrix for the vowel in (10). It is particularly clear that this set forms a group, as will be seen by a further consideration of (7). One of the respects in which (7) is inadequate is its failure to represent the notion 'agreement in place of articulation' in any coherent way. In other words, it fails to show that the features [labial], [alveolar] and [velar] form a group, just as the matrix in (10) fails to identify the various groups of features we have distinguished.

There are formidable difficulties in formalising these observations, in particular with the formalisation of the notion 'nasals agree with a following consonant in all the features forming the group characterising place of articulation', and we postpone discussion of this until §1.4. This problem is exacerbated by the fact that there is little agreement on what should constitute the set of features in this group. An *SPE*-type system would characterise, say, English /t/ by means of a matrix something like:

(32)
$$\begin{bmatrix} -\text{sonorant} \\ +\text{consonantal} \\ -\text{continuant} \\ -\text{voice} \\ +\text{anterior} \\ +\text{coronal} \\ -\text{high} \\ -\text{low} \\ -\text{back} \\ -\text{round} \end{bmatrix}$$

[19] In §2.5 we consider an approach which characterises scalar processes like these in terms of dependency relations between the features involved.

The two features [anterior] and [coronal] distinguish the major places of articulation as in (33):[20]

(33)

	labial	alveolar	post-alveolar	palatal	velar	uvular	pharyngeal
	p b	t d		c ɟ	k g	q ɢ	
	ɸ β	s z	ʃ ʒ	ç ʝ	x ɣ	χ ʁ	ħ ʕ
[ant]	+	+	–	–	–	–	–
[cor]	–	+	+	–	–	–	–

[+anterior] sounds are produced with a stricture in front of the postalveolar region; [+coronal] sounds are produced with the blade of the tongue raised above the 'neutral' position. We will not consider the motivation for the definitions here, other than noting that natural classes can be defined along the expected lines.

The tongue-body and tongue-root consonants (i.e. the [–anterior, –coronal] set) can be distinguished from each other by the use of the 'vowel' features [high], [low] and [back], as shown in (34):

(34)

	labial	alveolar	post-alveolar	palatal	velar	uvular	pharyngeal
	p b	t d		c ɟ	k g	q ɢ	
	ɸ β	s z	ʃ ʒ	ç ʝ	x ɣ	χ ʁ	ħ ʕ
[ant]	+	+	–	–	–	–	–
[cor]	–	+	+	–	–	–	–
[high]	–	–	+	+	+	–	–
[low]	–	–	–	–	–	–	+
[back]	–	–	–	–	+	+	+

The representation of place of articulation is an area which has given rise to a number of alternative proposals in recent years, and we shall outline one here which has gained wide currency, and which differs in various respects from the original *SPE* proposal.

McCarthy (1988: 99) characterises the *SPE* proposal as embodying 'place of articulation theory', in which the feature [anterior] is defined in terms of the passive articulator (in front of the postalveolar region). He contrasts this with 'articulator theory', in which the distinctions between the segments are made in terms of 'the active articulator making the constricting gesture rather than by place of articulation'. In articulator theory, the major places of articulation are distinguished by the features [labial], [coronal], [dorsal] and [radical], as in (35):

[20] For reference, in the following discussion we will provide symbols for representatives of each of the categories given, drawn from the stop and fricative series.

(35)

	labial	alveolar	post-alveolar	palatal	velar	uvular	pharyngeal
	p b	t d		c ɟ	k g	q ɢ	
	ɸ β	s z	ʃ ʒ	ç ʝ	x ɣ	χ ʁ	ħ ʕ
[lab]	+	–	–	–	–	–	–
[cor]	–	+	+	+	–	–	–
[dors]	–	–	–	+	+	+	–
[rad]	–	–	–	–	–	+	+

In line with McCarthy's distinction between articulator and place of articulation theory, we can say that [labial] consonants are produced with the lips, [coronal] consonants with the blade of the tongue, [dorsal] consonants with the tongue body (dorsum), and [radical] consonants with the tongue root.

What advantages does (35) have over the system given in (33) and (34)? Notice first of all that [anterior] is no longer used to identify any of the major places of articulation. One reason for this is that in the *SPE* treatment, the class of labials and alveolars is predicted to form a natural class, characterised as [+anterior]. As Yip (1989: 350) points out, this is not a recurrent class in the phonologies of languages; we do not find phonological processes affecting the set of coronal and non-coronal anteriors (e.g. /p t̪ t/). Abandoning [anterior] as the characterisation of a major place of articulation means that this anomaly is removed (although, as we shall see presently, [anterior] is still required in this system to subcategorise the class of coronals).

The introduction of a feature [labial] allows the class of labial consonants to be given a 'positive' characterisation, rather than their somewhat opaque definition in *SPE*: 'consonants produced with a stricture in front of the palato-alveolar region whose production does not involve raising of the blade of the tongue', i.e. [+anterior, –coronal].

As we have already noted, tongue-body consonants are characterised as [dorsal]. However, notice that in (35) the segments characterised as being [coronal] have been extended so as to include palatals (e.g. /ç ɟ/) as well as postalveolars (e.g. /ʃ ʒ/). Whether palatals should be characterised as [coronal] or as [dorsal], or perhaps as both, is a topic of some debate. Although, as we have seen, Chomsky and Halle characterise them as [–coronal], in accordance with their definition of [+coronal] as involving raising of the blade of the tongue, Hall (1997: 6) observes that 'the vast majority of phonologists have concluded that palatal sounds are [+coronal] because they pattern in many languages with the alveolars', and notes that this causes Halle and Stevens (1979: 346) to reformulate the definition of [+coronal] as involving 'the raising of the frontal (i.e. tip, blade, and/or central) part of the tongue so as to make contact with the palate'. We will assume here that palatals are indeed [+coronal].[21]

Articulator theory, then, allows a positive characterisation of each of the major places of articulation. However, it is clear that in the coronal region, many more oppositions are required than we have considered up to now. Phonological distinctions in this region are found which involve a number of parameters which we have not yet considered in either of the theories which we have been discussing.

The feature [anterior] is generally retained in articulator theory, but is restricted to those segments which are [coronal]. That is, [coronal] segments may be [+anterior] or [–anterior], but segments which are not coronal simply have no specification for the feature [anterior].[22] The situation for coronals is shown in (36) (we add dental and retroflex consonants to the set in (35), and consider palatals to be [+coronal]):

(36)	dental	alveolar	postalveolar	retroflex	palatal
	t̪ d̪	t d		ʈ ɖ	c ɟ
	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ
[cor]	+	+	+	+	+
[ant]	+	+	–	–	–

Thus the definition of [anterior] is retained, but now applies only to consonants produced with the tongue-blade as active articulator.

Two other features, originally proposed by Chomsky and Halle (1968), provide further distinctions amongst the various coronal segments. These are [strident] and [distributed].

The feature [strident] distinguishes various members of the class of fricatives, on the basis of the relative amount of 'high-frequency noise' involved. Thus sounds such as [s z ʃ ʒ] have a relatively large amount of high-frequency noise, and are [+strident] as opposed to their counterparts [θ ð ç ʝ]. Although this feature is not defined in terms of place, we include it here, as it serves to subcategorise consonants in the coronal area. Its phonological relevance can be seen in processes such as that illustrated in (37):

[21] Hall (1997: §1.2) suggests that the uncertainty surrounding the status of palatals with respect to the feature [coronal] is due to the fact that palatal fricatives such as /ç ʝ/ behave differently from non-continuants such as /c ɟ ɲ/. The fricatives behave as dorsals rather than as coronals. Hall cites a process from a dialect of German spoken near Düsseldorf, in which uvular /ʀ/ is devoiced to [χ] before coronal /t s ʃ/, but not before /p f k ç/. Thus /ç/ here patterns with the non-coronals. On the other hand, non-continuants (/c ɟ ɲ/) typically pattern with coronals, rather than with dorsals. Hall concludes that the continuants are [–coronal] and the non-continuants [+coronal], and attributes this to an articulatory difference; the 'palatal' stops are in fact alveolo-palatal, he claims.

[22] In §2.2 we consider the question of how the notion of a feature not being relevant to a segment can be formalised.

(37) a.	masses	MASS+PL	/mæsəz/	b.	moths	MOTH+PL	/mɒθs/
	buzzes	BUZZ+PL	/bʌzəz/		lathes	LATHE+PL	/leɪðz/
	coshes	COSH+PL	/kɒʃəz/				
	edges	EDGE+PL	/ɛdʒəz/				

In the formation of English plurals we find a different form of the suffix after the [+strident] fricatives in (37a) (/-əz/, or /-ɪz/, depending on dialect) than after the [–strident] fricatives in (b).

Stridency does not seem to be directly relevant to plosives;[23] in (38) we show how it further subcategorises the set of coronals:

(38)	dental	alveolar	postalveolar	retroflex	palatal
	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ
[cor]	+	+	+	+	+
[ant]	+	+	–	–	–
[strid]	–	+	+	+	–

[+distributed] sounds are those in which the consonantal stricture is 'relatively long'. This feature is often used to distinguish tongue-blade (laminal) sounds, which are [+distributed], from tongue-tip (apical) sounds [–distributed], and non-retroflex [+distributed] from retroflex [–distributed].

The feature [distributed] is often seen as somewhat unsatisfactory (see e.g. Keating 1991). On the one hand, it is not clear that the phonetic definition is entirely appropriate, and on the other, evidence that [+distributed] consonants can function as a class as opposed to [–distributed] consonants is limited. However, Pulleyblank (1989: 384–5) cites processes in Australian languages which have the four coronal stops given in (39):

(39)	laminal (inter)dental	apical alveolar	apical postalveolar	laminal palatal alveolar
	d̻	d	d̺	ɟ

Pulleyblank, following Dixon (1980: §6.4), notes that the palatal and (inter)dental stops, which are laminal, pattern together in phonological processes in these languages, as opposed to the alveolar and postalveolar stops, which are apical. A grouping of this sort can be satisfactorily characterised by the feature [distributed]. In view of facts like these, we include [distributed] here, although we suspect that it need only be used for stops – for fricatives, [strident] is sufficient to characterise the oppositions found in languages. (40) shows how [distributed] subcategorises the coronal series of stops:[24]

[23] However, notice that Jakobson *et al.* (1951: 24) use the feature [strident] to distinguish affricates from plosives, irrespective of place of articulation.

[24] We distinguish apicals from laminals by means of the diacritics ̺ and ̻, respectively.

(40)

	apical dental	laminal dental	apical alveolar	laminal alveolar	retroflex	palatal
	t̪ d̪	t̪̻ d̪̻	t̺ d̺	t̻ d̻	ʈ ɖ	c ɟ
[cor]	+	+	+	+	+	+
[ant]	+	+	+	+	–	–
[distr]	–	+	–	+	–	+

Notice that the distinction between [+distributed] and [–distributed] segments is independent from that between dentals and alveolars, which can both be either laminal or apical. However, it is generally claimed that languages do not *contrast*, say, laminal dentals with laminal alveolars, so that in a language which has both /t̪/ and /t/, the two stops will be distinguished not only by place, but also by laminality vs apicality; [distributed] can then be used to characterise the difference between them.

This concludes our discussion of individual features. It is of course clear that additional features will be required in phonological descriptions (see again the references in §1.3), in particular features characterising the various airstream mechanisms (i.e. pulmonic, glottalic and velaric) utilised in the languages of the world, and those characterising different phonation types (voice, voicelessness, breathy and creaky voice and aspiration) and tonal contrasts. We introduce these in the course of our discussion, as and when they become necessary.

1.3.5 The characterisation of grouping

We have now identified a set of features which can be used in phonological description, and have also uncovered some evidence which suggests that these features might be organised into groups. Evidence for the grouping of features can be found in hierarchy-related processes such as weakening, for example, but also in cases where two or more features together define a class of segments which functions together in some phonological process, and thus forms a natural class. Thus we saw in (15) that [sonorant] and [consonantal] behave in this way; the conjunction of the values [+sonorant] and [+consonantal] in the last segment of a noun in Dutch, among other things, determines that a schwa will be inserted before the diminutive suffix.

Similar evidence can be found in the process of nasal place assimilation, which we discussed at some length in §1.2, but whose formulation we have not yet considered in terms of the features developed above.[25] A direct 'translation' of (7), our last attempt at a formulation of nasal assimilation, would yield (41):

[25] For the purposes of this discussion, we utilise the feature system proposed in *SPE* for the characterisation of place of articulation, rather than that based on articulator theory.

(41)

$$[+\text{nas}] \rightarrow \left\{ \begin{array}{l} \begin{bmatrix}+\text{ant}\\-\text{cor}\end{bmatrix} / \_\_ \begin{bmatrix}+\text{ant}\\-\text{cor}\end{bmatrix} \\ \begin{bmatrix}+\text{ant}\\+\text{cor}\end{bmatrix} / \_\_ \begin{bmatrix}+\text{ant}\\+\text{cor}\end{bmatrix} \\ \begin{bmatrix}-\text{ant}\\-\text{cor}\end{bmatrix} / \_\_ \begin{bmatrix}-\text{ant}\\-\text{cor}\end{bmatrix} \end{array} \right\}$$

In assimilations of place like these, then, [anterior] and [coronal] form a group, as we would expect.[26]

Examples like these could be multiplied, and we can identify other kinds of phonological behaviour which provide further evidence for feature grouping. For example, it has been argued that /h/ in many languages is a 'defective' segment, in that it lacks a particular group of features, those characterising place of articulation. In other words, /h/ has no independent place of articulation, but acquires its specification for these features from a following vowel: in English *heat* /hiːt/, for example, the articulators adopt the position of the following high vowel during the production of /h/, whereas in *harp* /hɑːp/, the articulation of /h/ is very different.

If feature grouping plays such an important role in phonology, it is clear that this must be formalised in some way. In other words, the simple list of features found in the *SPE* approach must be structured in such a way that it reflects our claims about the grouping of features.

Perhaps the most straightforward solution, originally proposed in Lass and Anderson (1975) and Lass (1976), is simply to divide the feature-matrix into submatrices, or **gestures**. Thus we might distinguish a **categorial gesture**, containing the features [sonorant], [consonantal], [continuant] and [voice], i.e. the group of features which seems to be involved in the expression of the relative sonority of segments, and a **place gesture**, containing the features characterising the vowel space and defining consonantal place of articulation. However, the categorial gesture might have to show a further subdivision, given our claim that [sonorant] and [consonantal] are more closely related to each other than to [continuant] and [voice], and vice versa. Let us call these the **major class** and **manner** gestures. The feature-matrix for English /θ/ might have the representation in (42):[27]

[26] The vowel features [high], [low], [back] and [round] would also be required in a full statement of the process; for simplicity, we omit these features here.

[27] Note that, for ease of exposition, we restrict ourselves to a subset of the features which we have introduced. In addition, we do not consider the incorporation into (42) and (43) of features characterising tone and phonation.

(42)

−sonorant, +consonantal	major class	categorial gesture
+continuant, −voice	manner	
+anterior, +coronal	consonantal	place gesture
−high, −low, −back, −round	vowel	

However, in recent work in phonology the notion of grouping has been presented in terms of a rather different set of notational conventions from that given in (42). Furthermore, the terminology used is different. Within the model of **autosegmental phonology**, which we introduce in §1.4, the segment has in recent years been represented in terms of a **feature geometry**, i.e. a tree-like structure similar to that which we introduced in (11):

(43)

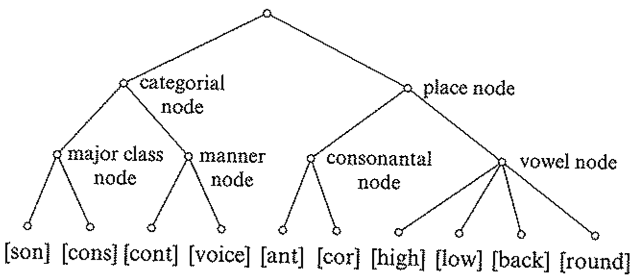

Notice that although (42) and (43) look very different, they are entirely equivalent in the claims made about the structure of the segment; a labelled bracketing like (42) and a tree structure (or feature geometry) like (43) are notational variants.

Explicit proposals for a geometrical structure of the type in (43) were first made by Clements (1985), and have been developed in a number of publications since (e.g. Sagey 1986; McCarthy 1988; Clements and Hume 1995; Pulleyblank 1995). The particular structure in (43) is merely intended to show how the gestural representation in (42) can be 'translated' into a geometrical representation. Proposals within geometrical phonology differ in various respects from (43), and we return in §1.4 to a discussion of some of the substantive issues involved. Here we consider only the general claims made about the nature of segmental structure within a geometrical approach, rather than the specific form of the geometry in (43).

In a geometry such as (43), features are seen as labels for NODES (represented as small circles), and, as in a model incorporating gestures, are grouped together, under higher nodes. These higher nodes, which are referred to as **class nodes**, have essentially the same status as gestures and sub-gestures, in that they can act as units in phonological rules. In other words, the set of features grouped under a particular class node can be appealed to in a phonological rule, just as a set of features forming a gesture or sub-gesture can be appealed to. Similarly, a segment may lack one or more of the nodes, and hence all the features dominated by the nodes in question. Compare for example the representation of English /θ/ in (44a) with that of /h/ in (44b):

(44) a.

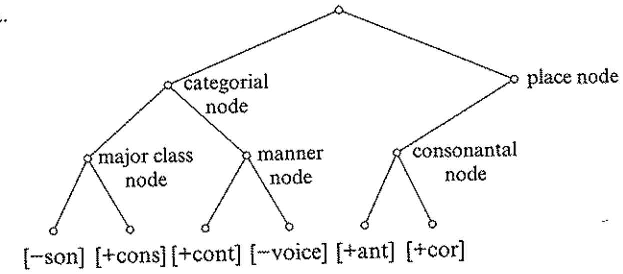

b.

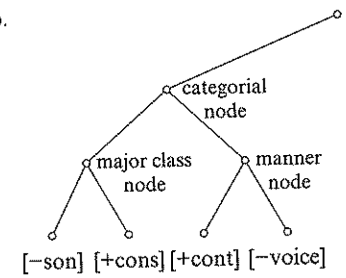

The representation for /h/ lacks a place node, and hence displays none of the features dominated by that node, while /θ/ and /h/ have identical specifications for those features dominated by the categorial node.[28]

It is clear that structures such as (42) and (43) allow us to make reference in phonological rules to groups of features. Indeed, if we find rules which refer to the particular groupings suggested here, we are thereby providing support for the subdivisions proposed. However, we have not yet addressed the question of how these structures play a role in the operation of phonological rules, a problem to which we now turn.

[28] We are further assuming here that /θ/ and /h/ lack a vowel node. See e.g. Kenstowicz (1994: §9) and Clements and Hume (1995) for discussion of vowel and consonant place features.

1.4 Autosegmental phonology

Consider again the formulation of nasal place assimilation in (41). It is clear that this formulation is inadequate in at least three respects. In the first place, it does not show that the two features involved form a group. Secondly, it fails to show that the value for each of the features which is changed by the rule must be the same as the value for the corresponding feature in the segment to which the nasal assimilates. Thirdly, just as in (9) above, there is nothing to prevent us having a rule like (45):

(45)
$$[+\text{nas}] \rightarrow \left\{ \begin{array}{l} \begin{bmatrix} +\text{ant} \\ -\text{cor} \end{bmatrix} / \_\_ \begin{bmatrix} +\text{son} \\ -\text{cons} \end{bmatrix} \\ \begin{bmatrix} +\text{ant} \\ +\text{cor} \end{bmatrix} / \_\_ \begin{bmatrix} +\text{high} \\ +\text{son} \end{bmatrix} \\ \begin{bmatrix} -\text{ant} \\ -\text{cor} \end{bmatrix} / \_\_ \begin{bmatrix} -\text{round} \\ -\text{voice} \end{bmatrix} \end{array} \right\}$$

(45) is formally just as easy to express as (41).

With respect to the first two objections to (41), the generalisation we are trying to express is that the features of the place gesture must be identical for the two segments. However, the formulation of this generalisation in terms of the kind of notation we have been using up to now is, perhaps not surprisingly, not straightforward: there is no obvious way of expressing the notion of identity, although attempts involving the use of 'Greek letter variables' have been offered, as in (46), where each feature is bound by a variable (see e.g. *SPE*: 352):

(46)
$$[\text{nas}] \rightarrow \begin{bmatrix} \alpha\,\text{ant} \\ \beta\,\text{cor} \end{bmatrix} / \_\_ \begin{bmatrix} \alpha\,\text{ant} \\ \beta\,\text{cor} \end{bmatrix}$$

The use of a particular Greek letter variable denotes identity for the feature in question, and the incorporation of this convention takes care of the second objection to (41), by extending the domain of the Greek letter variable convention to all the features dominated by a particular class node, as in (47):

(47)
$$[\text{nas}] \rightarrow \alpha[\text{PLACE}] / \_\_ \alpha[\text{PLACE}]$$

However, such linear formulations, even patched up with devices like these, are hardly adequate to characterise assimilation processes, in that they are unable to deal with the third objection raised above. There is again nothing to prevent us replacing the [PLACE] in the environment of (47) by [CATEGORIAL], for example.

What we are failing to characterise in (47) is the fact that the nasal preceding a consonant in assimilation cases acquires its specification for place of articulation from the consonant by a process referred to as SPREADING, i.e. it does not have an independent set of features characterising place, but 'shares' its place features with the following consonant. (48) is a formulation of this, using the conventions of non-linear autosegmental phonology (see e.g. Goldsmith 1976, 1990; Clements 1977):

(48)

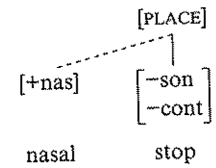

In this formulation, the fact that the two consonants in this environment share a single specification for place of articulation is made explicit. The 'direction' of the assimilation is indicated by the dashed line in (48): the features of the place gesture which are associated with the stop spread from the stop to the nasal.

Cases like these provide the basis for the theory of autosegmental phonology. What we have here is an example of one set of features (the place features) operating *independently* (hence the name of the theory) from the other features. With respect to the place features in (48), there is only one specification – a single **autosegment** – while for all other features there are two specifications, and hence two segments. Among other things, then, autosegmental phonology is concerned with the characterisation of cases where two segments necessarily share the same specification for a feature or group of features.

In this connection, it is perhaps useful to compare words like *camber* and *linger*, which display nasal place assimilation, with English words containing a lateral consonant followed by an obstruent. At first sight, the forms in (49) suggest that such forms do not behave in the same way as the nasals discussed above:

(49) gulp [gʌlp]
kilt [kɪlt]
milk [mɪlk]

In all cases in (49) the lateral, which is phonologically alveolar, is phonetically realised as an alveolar, so that English does not seem to have a 'lateral place assimilation' rule of the same type as the nasal place assimilation rule discussed earlier.[29] This raises the question of how the sequence of consonants

[29] Notice that the /l/ is velarised in many dialects of English in such contexts; this, however, does not affect the issue at hand.

in *kilt* should be represented. It is of course true that the /l/ and the /t/ have the same place of articulation, and thus the same features in the place gesture, but the /l/ is not alveolar *because* the /t/ is alveolar. Rather, the /l/ is alveolar 'in its own right', as is shown by (49). Thus, in a representation of the same type as (48), it might appear that we should assign a separate (although identical) specification for the place gesture to /l/:

(50)

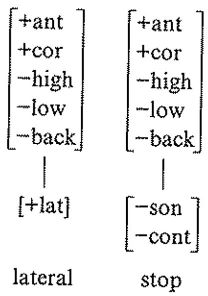

(Notice that we here introduce a feature [lateral], whose function will be clear.)

However, there are at least three problems associated with a representation such as (50). In the first place, there has been a great deal of discussion as to whether (50) is a well-formed phonological representation, or whether successive identical specifications for particular gestures must eventually be collapsed by what is known as the 'Obligatory Contour Principle', to give a structure more like (48) (see e.g. Odden 1988; Yip 1988). Secondly, the question arises as to whether laterals in English have to be specified underlyingly as having place features at all. English has only one lateral, /l/, and so place of articulation is not distinctive for a segment which is [+lateral] – its place of articulation is always predictable. Thus it may not be appropriate to include a place specification for the lateral in the representation in (50). We return to this aspect when we introduce the notion of underspecification in §2.2.

We devote some space to the third problem at this point. It does not seem to be strictly true that English has no assimilation process affecting laterals. Consider the forms in (51):

(51) health [hɛl̠θ]
kilt [kɪlt]
Welsh [wɛl̠ʃ]

(We use [l̠] to represent a retracted alveolar lateral.) The forms in (51) appear to be fairly common in English (although notice that many speakers have no central closure at all, especially in the case of a following postalveolar fricative /ʃ/). Thus it appears that there *is* a process of lateral assimilation, which

is restricted to cases where the following consonant is a coronal obstruent. This phenomenon lends further support to two aspects which we have been considering. In the first place, we have another case of spreading, and secondly, we have uncovered further evidence for treating [coronal] as one of the major articulation types, as in (35): the features which distinguish the various [+coronal] segments ([anterior], [strident] and [distributed]) form a group, as all and only these features appear to be involved in the spreading process in (51). Thus we can argue that the [coronal] node is what spreads in the English lateral place assimilation process, giving the formulation in (52a), whereas dorsal segments do not spread to laterals, as in (52b):

(52) a.

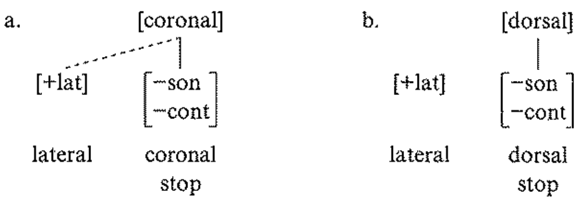

We emphasise at this point that this formulation is a very tentative one, which will be subject to revision in the light of the theoretical developments which we will consider in the remainder of this chapter and in Chapter 2. However, we include it at this point because it throws light on the kinds of considerations that have contributed to these theoretical developments.

We turn now to a more detailed account of autosegmental phonology, starting with a consideration of the ways in which it differs from *SPE*.

In *SPE*, phonological representations are conceived of as unilinear strings of segments, where segments are **unstructured**, **unordered** and **non-overlapping** sets of binary features:

(53) $\begin{bmatrix} +F \\ -G \\ +H \end{bmatrix} \begin{bmatrix} -F \\ -G \\ +H \end{bmatrix} \begin{bmatrix} -F \\ +G \\ +H \end{bmatrix} \begin{bmatrix} +F \\ +G \\ +H \end{bmatrix}$

As we have seen, the introduction of the concept of feature grouping brings some degree of structure to the segment.

The claim that the features in a matrix are sequentially unordered entails that a single segment cannot have two values for a particular feature, as in (54), as this would lead to an anomalous specification:

(54) $\begin{bmatrix} +\text{son} \\ +\text{cons} \\ +\text{nas} \\ -\text{nas} \end{bmatrix}$

At first sight, it is not obvious why a feature specification with opposing values for a single feature would be desirable in any case. However, there are a number of phenomena involving **contour segments**, such as prenasalised consonants, affricates (cf. the discussion above) and contour-toned vowels, which appear to point to the desirability of a single segment being able to have two – perhaps sequentially ordered – specifications for a single feature. The term contour segment is generally used to describe two distinct 'events' which appear to function as a single segment in the phonology of a language. Thus, as we have seen, affricates display both a closure phase and a release phase, but behave phonologically as single segments. Similarly, prenasalised consonants involve both a nasal and an oral phase, but, for phonological reasons, are not interpreted as a sequence of two consonants. 'Contour tone' is the term used to describe the interpretation of either a falling or a rising tonal pattern as a sequence of high tone + low tone (H+L) or low tone + high tone (L+H), respectively. As such tonal patterns can be associated with single (short) vowels, we again have two apparently sequentially distinct feature specifications associated with a single segment.

As an example of the relevance of this type of phenomenon for phonology, consider the forms from Apinayé in (55), presented by Anderson (1976):

	a.	b.
(55)	[V b d V]	[V b V]
	[Ṽ m d V]	[Ṽ m͡b V]
	[V b n Ṽ]	[V b͡m Ṽ]
	[Ṽ m n Ṽ]	[Ṽ m Ṽ]

These forms illustrate that the nasality of Apinayé consonants depends on the nasality of the contiguous vowel.[30] Those in (55a) show that the first consonant in a sequence of two takes its specification for nasality from a preceding vowel, and the second from a following vowel. In these cases, then, each consonant has simply one specification for nasality. In (55b), however, we see what happens when there is only a *single* intervocalic consonant. If both vowels have the same value for nasality, there is no problem – the consonant has the same value. However, in cases where the vowels have opposite values, it seems that *both* vowels spread their values, to give either a prenasalised or a postnasalised stop. Intuitively, the generalisation here is quite obvious: the first part of the consonant gets its nasality specification from the preceding vowel; the second part from the following vowel. However, such a state of affairs is difficult to express in the formalism of *SPE*,

[30] Strictly speaking, the data in (55) only shows that the nasality of the vowels and consonants are not independent, rather than showing the directionality of the relationship. For a discussion of the details of this phenomenon, see Anderson (1976).

which would require the introduction of features such as [prenasalised] and [postnasalised], together with two distinct rules, to express the sequential change within the segment, as in (56) and (57):

(56) a. $\begin{bmatrix} \text{+cons} \\ \vdots \\ \text{+nas} \\ \text{+prenas} \end{bmatrix}$ b. $\begin{bmatrix} \text{+cons} \\ \vdots \\ \text{+nas} \\ \text{+postnas} \end{bmatrix}$

(57) a. $C \rightarrow [\text{+prenas}] \,/\, \begin{bmatrix} V \\ \text{+nas} \end{bmatrix} \_\_$

b. $C \rightarrow [\text{+postnas}] \,/\, \_\_ \begin{bmatrix} V \\ \text{+nas} \end{bmatrix}$

However, (58), in which there are *two* occurrences of [nasal], seems more adequately to express the phonetic state of affairs in Apinayé, which is, quite simply, that at the left prenasalised consonants show the same behaviour as nasals do, whereas at the right they pattern with non-nasals (cf. Anderson 1976):

(58) $\begin{bmatrix} \text{+cons} \\ \vdots \\ \text{+nas, −nas} \end{bmatrix}$

Such phenomena are often referred to as **edge effects**, and many similar cases can be found. A particularly strong case can be made with reference to contour-toned vowels. Indeed, the analysis of tone systems in languages provided much of the initial impetus for the development of autosegmental phonology (see Goldsmith 1976, for example), and is still commonly used to introduce the theory (e.g. Goldsmith 1990). We look briefly here at the issues involved.

In many languages of the world lexical items are specified as bearing a particular tone. Let us consider a simple system with only two tones, High (H) and Low (L), such as that of Mende, a language spoken in Sierra Leone (cf. Halle and Clements 1983). (59) gives some Mende lexical items, in which the tones are marked on the vowels of each item (we follow the normal convention of marking high tone with an acute accent (´) and low tone with a grave accent (`)):

(59)	kɔ́	'war'
	pɛ́lɛ́	'house'
	bɛ̀lɛ̀	'trousers'
	nàvó	'money'

In (59) we see that each vowel has a single tonal specification, so that at first sight we might assume that tone can be dealt with linearly, just like features such as [nasal] in the *SPE* account. Things become more complicated, however, when we consider forms such as those in (60):

(60) mbû 'owl'
mbǎ 'rice'
njàhâ 'woman'

These forms display what are referred to as 'contour tones', i.e. falling tones (ˆ) and rising tones (ˇ). Does this mean that we also have to recognise contour tones as basic tonal types in Mende?

This question can be answered by considering what happens to the tonal pattern of the various forms in (59) and (60) when they combine with the suffix *-ma*, meaning 'on':

(61) a. kɔ́-má b. mbú-mà
pɛ́lɛ́-má mbà-má
bɛ̀lɛ̀-mà njàhá-mà
nàvó-má

The forms in (61a) demonstrate that the suffix *-ma* has no independent tonal specification. Rather it acquires its tone by spreading of the tone which is associated with the last vowel of the stem, as in (62):

(62) a. b.

Thus a single tonal specification may be associated with more than one segment. But we can go further than this. Compare (61) with (59) and (60). We see that the tone of the suffix can differ from the final tone of the stem, but only if the final tone of the stem when it occurs without a suffix is a contour tone. Thus, if the final tone of the stem in isolation is a falling tone, the suffix acquires a low tone, while the final tone of the stem changes into a high tone, as is evidenced by [mbû] vs [mbúmà]. If the final tone is a rising tone, exactly the reverse holds ([mbǎ] vs [mbàmá]).

What does this phenomenon tell us about tonal representations? Among other things, it suggests that contour tones are not independent basic entities, but, rather, are realisations of a *sequence* of two tones (H+L or L+H), associated with a single vowel. Thus [mbû] and [mbǎ] might be represented as:

(63) a. b.

We have now uncovered a case in which a single segment is associated with two tones, in addition to the cases given in (62), where a single tone is associated to two segments. Thus the tonal specification appears to be independent of the segmental representation, as is further evidenced by the behaviour of the forms in (61b). These forms not only confirm that contour tones are sequences of simple tones, but also show that underlyingly the tones are not associated with individual vowels, but are **floating**: they form part of the lexical representation of the morpheme in question, but do not link up with the vowels. Rather, the association of tones to vowels takes place in Mende only after suffixation.

Before giving an account of how we derive the forms in (61b), let us consider how the process of association operates. The general principle of association is clearly that in (64):

(64) Associate tones to vowels

However, (64) can be seen as the result of a combination of three subprinciples which ensure that the process of association produces well-formed surface representations, i.e. representations in which every tone is associated with at least one vowel, and every vowel with at least one tone. These subprinciples are given in (65) (cf. van der Hulst 1984):

(65) a. *Mapping*
Associate each tone with a vowel, working from left to right.
b. *Spreading*
If there are fewer tones than vowels, associate the final tone with all remaining vowels.
c. *Dumping*
If there are fewer vowels than tones, associate all remaining tones with the final vowel.

In the light of these principles (which we have formulated here to account for the Mende data only), we can derive various of the forms given above as follows. (66a) shows the derivations of stems alone; (66b) of stems when combined with a suffix. We indicate 'new' associations at each stage with a dashed line; existing associations at any stage are indicated by a solid line:

(66)

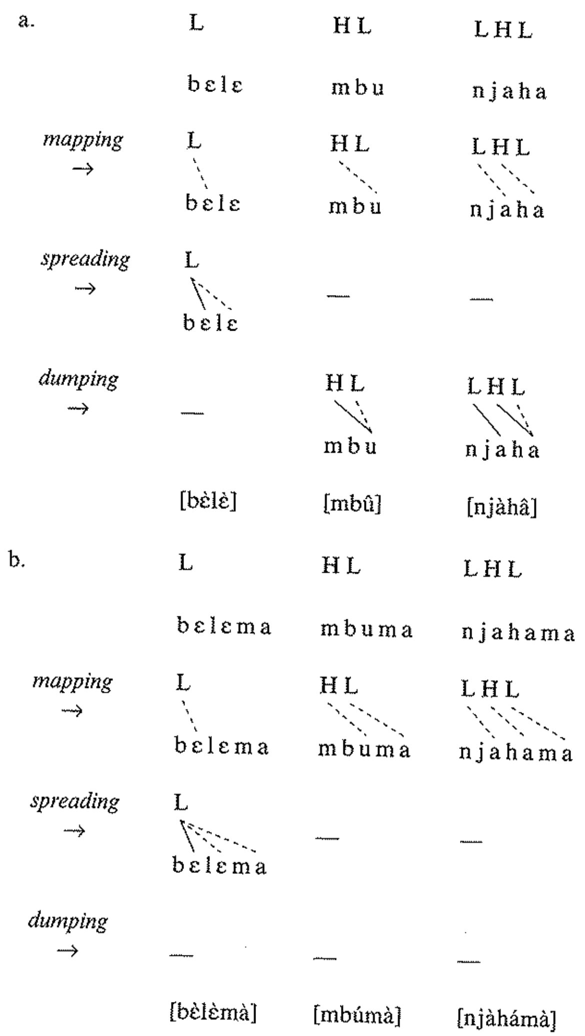

What (66) shows us is how Mende responds to a mismatch between the number of tones and the number of vowels, or more correctly the number of **tone-bearing units** (TBUs) in any word. The phonology of the language demands that all tones are realised, and so, where there are more tones than available TBUs, two tones must share a TBU. Equally, though, each TBU must be realised with a tone; spreading ensures that any toneless TBU can share a tone with some other TBU. This means that a stem-final contour tone is split into its two component tones when the toneless suffix *-ma* is attached to the stem.

However, not all tone languages allow contour tones. Goldsmith (1990: 20ff.) considers a dialect of Mixtecan, a tone language of Mexico, which 'has the property of requiring each vowel to have – maximally and minimally – exactly one tone'. What happens in this kind of language when there are more tones than vowels?

Goldsmith illustrates this with the following three words (we use M as a shorthand for Mid tone):

(67)

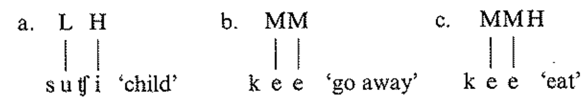

This dialect of Mixtecan has High, Mid and Low tones, which are associated on a one-to-one basis with the vowels, as shown in (67a, b). As we would expect, when these two words are combined as in (68), the associations between tones and vowels remains unchanged:

(68)

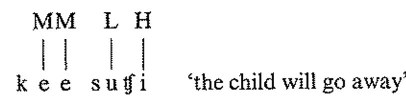

However, certain words, such as *kee* 'eat' in (67c), 'idiosyncratically have a High tone suffixed to them which is not realized when the word is pronounced in isolation, but which is realized when there is a following word for it to associate to'. When this word combines with /suʧi/, there are more tones than vowels, i.e. TBUs:

(69)

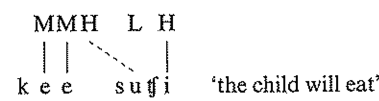

The final floating H tone of *kee* 'eat' displaces the L tone of /suʧi/, which is then left unassociated. Because, unlike Mende, the language does not permit contour tones, this displaced L tone is simply not realised. Thus different languages may respond to similar situations in rather different ways, depending on the phonological possibilities available to them.

Representations such as the ones we have been developing for tone capture the insight that the phenomena considered are manifestations of a more general property of speech, namely that segmental representations should allow for **overlapping**. In other words, these phenomena suggest that we should abandon what we might refer to as the **strict segment hypothesis**, embodied in representations such as those in (53), and allow single segments to be linked to more than one value of the same feature, and also a single feature to be associated with two segments.

This, then, is the central property of the theory of autosegmental phonology. However, it is of course not the case that feature spreading is an unrestricted process; rather, there are severe restrictions on the way in which features can spread from one segment to another. To illustrate this point in a preliminary way, we turn to the phenomenon of nasal spreading in English. In a word such as *pan*, the vowel is generally phonetically nasalised, to give [pæ̃n], as represented in (70):

(70)

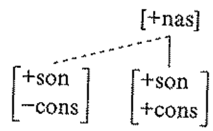

Here a single feature, [nasal], spreads to the preceding vowel. However, spreading is restricted, as is apparent if we consider words such as *kiln*, in which no spreading can take place. A pronunciation such as [kɪ̃ln], in which nasality spreads across an intervening liquid, is not possible. If we assume for the moment that the liquid is specified as [−nasal], this possibility is ruled out by a basic principle of autosegmental phonology, which is usually referred to as the **no-crossing condition**, formulated in (71) (see e.g. Goldsmith 1976 and, for discussion, Hammond 1988; Sagey 1988; Coleman 1998):

(71) *The no-crossing condition*
Association lines may not cross

The condition prevents segments sharing a feature specification if an intervening segment has the opposite value for that feature. Thus in *kiln*, for example, spreading is prevented because the two segments are not **adjacent**, as shown in (72):

(72)

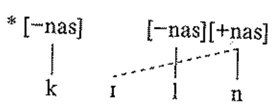

The non-adjacency of the two segments means that the [+nasal] specification for /n/ would have to spread *across* the association line linking the nasality node to the rest of the feature specification for /l/, thus violating (71).[31]

It appears then that the fact that nasality does not spread across the lateral in *kiln* is because spreading must be **local**, i.e. a feature can only spread to an adjacent segment. But it is clearly not the case that spreading processes are necessarily local in this strict sense. For example, in our discussion of Mende tone spreading above, we saw that tones could spread from one vowel to another, ignoring intervening consonants. A very similar situation is found in vowel harmony processes, such as the ATR harmony process of Akan illustrated in (29). Again, consonants are 'invisible' to the spreading harmony feature, so that the appropriate value of the feature [ATR] can spread from one vowel to another. Similar phenomena, involving various features, can readily be found.

[31] We should note that this treatment of the no-crossing condition is grossly simplified, in that we have not yet introduced the notion of underspecification (see §2.2).

Nevertheless, the claim that spreading is local is a sound one. How can we explain such phenomena, without abandoning this hypothesis? Consider a possible representation of the result of the ATR harmony process of Akan:

(73)

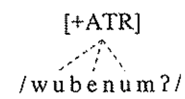

We can maintain the hypothesis by assuming that the segments which intervene between the various vowels in (73) have no representation for [ATR]. That is /b/ and /n/ in (73) are 'invisible' to the spreading feature – the vowels, then, are adjacent with respect to [ATR].

We do not at this stage investigate the issue of the circumstances under which segments can lack specifications for particular features. However, notice that this approach suggests that individual features are independent of each other in the sense that the node dominating a feature (or group of features) occupies its own **tier**, and leads to a representation in which adjacency on one tier need not correspond to adjacency on another tier.

On the assumption that each node occupies its own tier, we provide in (74) a representation of the English word *plank*, realised as [pl̥æ̃ŋk], and thus displaying spreading processes involving the tiers occupied by the features [voice], [nasal] and [back] (we ignore all other features here, and represent them simply by orthographic symbols):

(74)

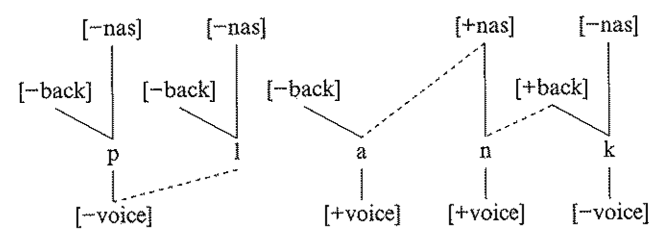

This conception of phonological structure is notoriously difficult to represent on paper, as it is essentially three-dimensional. (75) is an attempt at a more or less complete phonological representation of the English word *cab* /kæb/, in which no feature sharing is involved (as we have not yet developed a theory of the conditions under which segments can lack a specification for particular features, we here provide each segment with a specification for every feature):

(75)

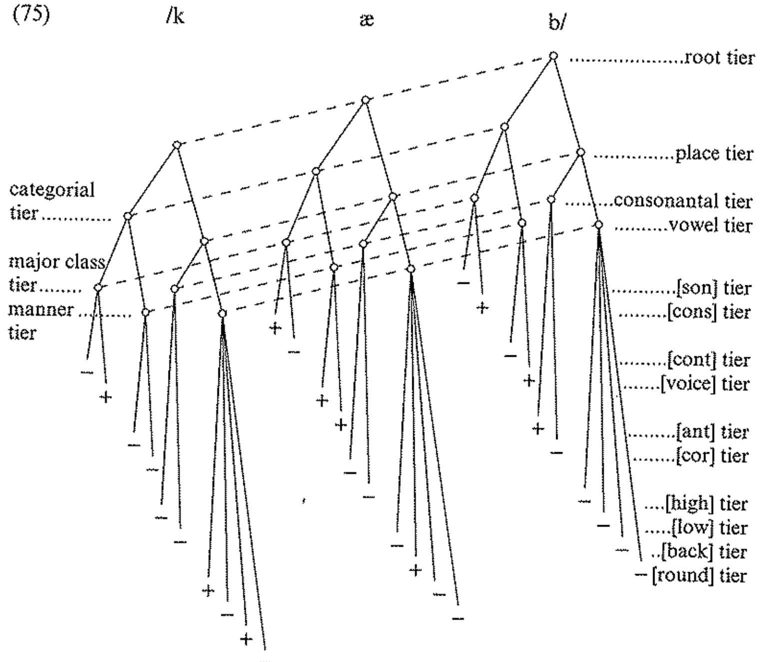

In (75) we see that each tier consists of a set of nodes. Thus, on the major class tier, we find three class nodes, one for each segment, showing that each segment has some specification for the features dominated by the major class node. On the [sonorant] tier, we again find three specifications. These, however, bear a label, i.e. either '+' or '–', given the fact that this is a content tier, rather than a class tier. Thus, for /k/ and /b/ we find the value '–', and for /æ/ the value '+'.

(75), then, shows both the independence of individual features and their grouping. It does not, however, show any examples of the phenomenon we have been discussing, the overlapping of segments. It is, however, clear how this can be achieved, although, again, the representation of what is quite

a simple notion is difficult in terms of the kind of formalism proposed in (75).

Let us consider again the phenomenon of assimilation of place of articulation in forms such as *camber*, where the place of articulation of the nasal is determined by that of the stop. As we demonstrated in (48), this state of affairs can be analysed as involving the spreading of the place features of the stop to the nasal, which has no independent place specification.[32] In terms of the formalism in (75), this can be characterised as in (76), which shows the underlying representation for the sequence /Nb/ (i.e. nasal unspecified for place of articulation followed by /b/), and in (77), the surface representation for [mb] after spreading:[33]

(76)

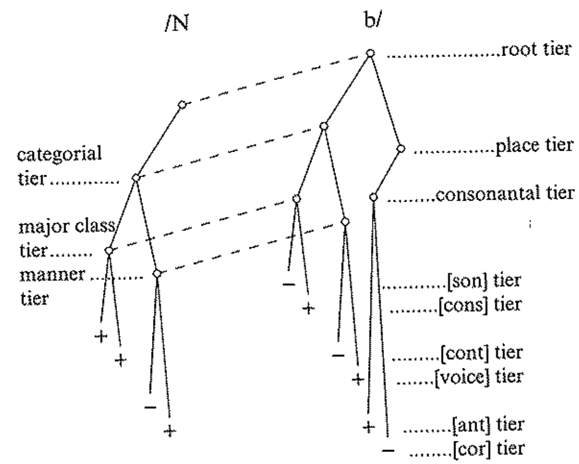

In (76) we find no place node for the nasal consonant. As a result of spreading, however, we have (77):

[32] In terms of a Prague School approach to phonology (see Trubetzkoy 1939), what we have here is a case of neutralisation: an opposition which is found elsewhere in the language (e.g. /ræm/ vs /ræn/ vs /ræŋ/) is not found in a particular environment, as evidenced by (1). In this approach, appeal was also made to a kind of underspecification: the segment occurring in a neutralisation environment was referred to as an 'archiphoneme', i.e. a phonological unit lacking those phonological properties neutralised in the environment in question, in this case the properties characterising place of articulation. Thus the set of words in (1) would be represented in Prague School phonology as /kæNbə/, /kæNtə/ and /kæNkə/, where /N/ is the archiphoneme characterising a nasal consonant, without a specification for place of articulation, which is not distinctive *in this environment*. (For a discussion, see Lass 1984a: ch. 3.) As we shall see in §2.2, the essence of this approach is contained in modern theories of underspecification, although rather different terminology is used; notice that the term 'archisegment' rather than archiphoneme is commonly used.

[33] We assume that non-dorsal consonants such as those in (76) and (77) lack features dominated by the vowel node in (75).

(77)

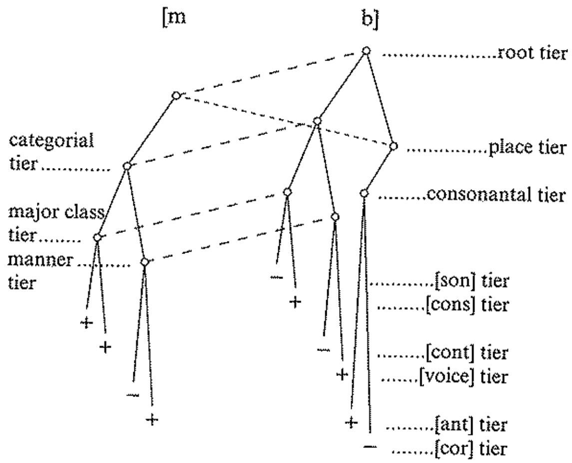

As in (48), the spreading of the place node of /b/ to the nasal 'segment' is denoted by the dashed line linking the root node of /m/ to the place node of /b/. Notice that, by convention, if a node spreads in this way, all the nodes which it dominates spread as well. In this case all the nodes characterising place of articulation spread as a result of the spreading of the place node, thus assigning labiality to the first of the two segments.

We see in (77) the consequences of abandoning the strict segment hypothesis. Although in the cases of the other tiers in (77) we find two nodes, there is only a single node on the place tier, and on all tiers dominated by it.

It is clear that the formalism used here can be used to account for the various other types of phenomena we have been considering. However, in view of the excessive complexity of diagrams like (75)–(77), we shall henceforth simplify our representations, using instead something rather more like (74) to characterise the notion of segments being made up of independent tiers, with the possibility of 'node sharing'. The reader should be aware, however, that these are merely shorthand representations for the fuller diagrams we have just introduced. Unless the hierarchical structure of the tiers is at issue, then, we will make use of the simpler form of representation.

In this book, we will assume that the kind of feature geometry outlined above is indeed appropriate for the representation of segmental structure, but we will not consider in any detail the question of the exact nature of this structure. In other words, we shall not consider at this point the question of whether the particular structure in (42) or (43), which we have adopted in the immediately preceding discussion, is appropriate, or whether some other organisation of the features into groups is to be preferred. Rather, we now

apply the model developed above to two processes which seem particularly susceptible to an autosegmental treatment.

1.4.1 Old English i-*umlaut*

Old English displays a number of morphophonemic alternations, resulting from the operation of a sound-change known as Old English *i*-umlaut (OEIU), which involves various vowel changes triggered by the presence of /i/ or /j/ in a following syllable. These alternations are illustrated in (78), taken from Lass and Anderson (1975: 117):

(78)

Alternation					
[uː]	[yː]	cūþ	'known'	cȳþan	'make known'
[u]	[y]	burg	'city'	byrig	'city (DAT SG)'
[oː]	[øː] ([eː])	dōm	'judgement'	dēman	'judge'
[o]	[ø] ([e])	ofost	'haste'	efstan	'hasten'
[ɑː]	[æː]	hāl	'whole'	hǣlan	'heal'
[ɑ]	[e]	mann	'man'	menn	'men'
[ɑ]	[æ]	faran	'go'	færst	'go (2SG)'

As will be seen from (78), by the Old English period most of the umlauted forms (those in the second column) had lost the triggering environment /i/ or /j/ (but note *byrig*, which retains /i/). However, whatever the philological details of OEIU, and whether or not it forms part of the *synchronic* phonology of Old English, it is clear that we are dealing here with a (historical) phonological process whereby, among other things, back vowels become front under the influence of a high front vowel or approximant.[34] Thus (79) shows the way in which some of the relevant forms underwent OEIU:[35]

(79)

*cūþ+i+an know-CAUSE-INF	>	*cȳþ+i+an	>	cȳþan
*hāl+i+an whole-CAUSE-INF	>	*hǣl+i+an	>	hǣlan
*far+ist go-2SG	>	*fær+ist	>	færst

In terms of an autosegmental approach, it is clear that what we have here is the spreading of an autosegment on the backness tier from the vowel of the suffix. A suitable representation might be that in (80):

[34] OEIU also affects low front vowels, as well as various diphthongs, but we ignore this here.
[35] Note the use of * here to denote that we are dealing with a reconstructed, rather than an attested, historical form.

(80)

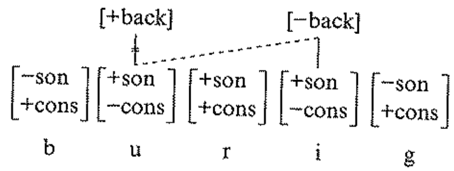

There are various aspects of OEIU which differ from certain of the autosegmental processes which we have previously considered. In the first place, notice that the spreading here is to a segment which apparently already bears a value on the backness tier. The spreading value thus *replaces* the original value, which must then be **delinked** (represented by the double bar through the association line). The process thus appears to be **feature-changing**. Notice, too, that although the consonants bear no specifications on the backness tier, so that they do not block spreading by violation of the no-crossing condition, it might well be possible to argue that they are also affected by spreading, as in cases where a phoneme may have either a [−back] or a [+back] allophone (e.g. [ç] vs [x] in the case of final *g* in Old English), it is the [−back] allophone which is found in the umlauted cases.

1.4.2 Vowel harmony in Turkish

We have just considered a phonological process of English which involves two vowels showing agreement in the values for a particular feature, a situation which is naturally represented within a model incorporating the notion of feature spreading. As we have seen with respect to the feature [ATR] (§1.3.3), however, there are even more spectacular examples of agreement of this sort, in which, within a particular phonological 'domain' in a language – say the word – all the vowels must have the same value for a particular feature (or indeed for a number of features). These vowel harmony phenomena have provided a fertile source of exemplification for proponents of autosegmental phonology. For this reason, we consider here in a little detail one of the most familiar processes of vowel harmony, that found in Turkish.

Turkish has a system of eight distinctive vowels, which are classified on phonological grounds into two height classes:[36]

(81)

	[−back]		[+back]	
	[−round]	[+round]	[−round]	[+round]
[+high]	i	y	ɨ	u
[−high]	e	ø	ɑ	o

[36] We ignore here the question of the precise phonetic realisation of the non-high series – we make no claim as to the phonetic accuracy of the transcriptions used here.

As shown in (82), all eight vowels can appear in monosyllabic stems, thus implying that the eight-way contrast has phonemic status:

(82)

çift	[çift]	'couple'
üç	[yç]	'three'
ek	[ek]	'affix'
köy	[køj]	'village'
kız	[kɨz]	'girl'
kurt	[kurt]	'worm'
at	[ɑt]	'horse'
son	[son]	'end'

In polysyllabic words, however, a number of restrictions hold. We give a preliminary formulation of these restrictions as (83):

(83) a. All vowels must have the same value for the feature [back].
b. A high vowel must have the same value for [round] as the directly preceding vowel (if any).

These requirements hold both in underived stems, i.e. words consisting of single morphemes, and in derived words, i.e. words consisting of a stem and a suffix (although there are many exceptional forms in underived words). The following examples (from van der Hulst and van de Weijer 1991) illustrate some typical derived forms:

(84)

ABS SG	POSS SG		ABS SG	POSS SG	
son	sonu	'end'	dere	deresi	'river'
boru	borusu	'pipe'	ɑt	ɑtɨ	'horse'
køj	køjy	'village'	tɑt	tɑdɨ	'taste'
kurt	kurdu	'worm'	kɨz	kɨzɨ	'girl'
tilki	tilkisi	'fox'	kɑp	kɑbɨ	'container'
inek	inei	'cow'	yty	ytysy	'iron'

We should notice that Turkish is an agglutinative language, which typically forms words by adding one or more suffixes to the stem. Thus the words in (84) consist of a stem morpheme, which can either occur on its own as the absolutive singular form (e.g. *boru*), or be followed by a possessive suffix (which, because of the harmony rules, may have any one of the realisations [i y ɨ u]).[37]

The requirement in (83a) represents the typical situation for vowel harmony: all vowels within some domain (the word, in this case) agree with respect to some feature ([back], in this case). It is clear that this kind of situation can be

[37] Notice that other segments in the stem morpheme may be affected by the suffixation process – this need not concern us here.

characterised by the autosegmental mechanism we have been developing, although we should notice that (83a) represents a different state of affairs from OE *i*-umlaut, in that it does not seem to involve spreading from one vowel to another; rather, *all* the vowels within the domain of the harmony process seem to have the same value for the feature in question.

(83b) is a more complicated restriction, which differs from (83a) in two ways. Firstly, the restriction concerns only a subclass of the vowels, the high vowels, and secondly, it has a directional aspect, which is not found in (83a).

The additional set of paradigms in (85) show the effect of (83b): non-high vowels do not display rounding harmony:

(85)

ABS SG	ABL SG	ABS PL	POSS PL	
son	sondan	sonlar	sonlarɨ	'end'
boru	borudan	borular	borularɨ	'pipe'
køj	køjden	køjler	køjleri	'village'
kurt	kurttan	kurtlar	kurtlarɨ	'worm'
tilki	tilkiden	tilkiler	tilkileri	'fox'
inek	inekten	inekler	inekleri	'cow'
dere	dereden	dereler	dereleri	'river'
at	attan	atlar	atlarɨ	'horse'
tat	tattan	tatlar	tatlarɨ	'taste'
kɨz	kɨzdan	kɨzlar	kɨzlarɨ	'girl'
kap	kaptan	kaplar	kaplarɨ	'container'
yty	ytyden	ytyler	ytyleri	'iron'

(where [den, dan] – or [ten, tan] – are the realisations of the ablative morpheme, [ler, lar] those of the plural morpheme, and, as above, [i, ɨ] two of the realisations of the possessive morpheme).

Low vowels, then, need not be in rounding harmony with other vowels, as shown by the ablative singular and absolutive plural forms of *son* [son] 'end', where a back rounded vowel is followed by a low unrounded vowel, and by the same forms for *köy* [køj] 'village', where front vowels are involved.

The restrictions on rounding harmony can be derived from a more general constraint:

(86) The vowels /o/ and /ø/ only occur in initial syllables (i.e. low vowels are never [+round] except in initial syllables).

Let us now turn to the directional aspect of rounding harmony. Consideration of the possessive plural forms in (85) shows that a vowel takes its value for [round] from the immediately preceding vowel, rather than from the first one. Consider *ütüler* [ytyler], the plural form of *ütü* 'iron'. The plural morpheme contains a low vowel, which is therefore unrounded, even though the vowels of *ütü* [yty] are rounded (cf. (86)). The singular possessive form is

ütüsü [ytysy], with the possessive morpheme being realised as [y] – because it is a high vowel, it undergoes rounding harmony. However, the plural possessive form is *ütüleri* [ytyleri]. Here the possessive morpheme is realised as [i], i.e. the vowel is unrounded. This is clearly because the immediately preceding vowel – the vowel of the plural morpheme – is unrounded, and so the vowel of the possessive morpheme harmonises in the expected way.

Thus the rounding of a high vowel cannot spread across a following low (and therefore unrounded) vowel. As we have already seen, then, harmony statements apply to adjacent vowels. More precisely, in this case rounding harmony only applies to segments which are adjacent *on the roundness tier*.

The harmony processes of Turkish mean that the full set of vowel contrasts need only be specified on initial vowels. In non-initial syllables we only find a contrast between [+high] and [–high], with the values of the features [round] and [back] being determined by spreading from the initial syllable to non-initial syllables.

Many autosegmental treatments of Turkish vowel harmony propose that, lexically, the features [back] and [round] are floating, i.e. they are not lexically associated to a particular vowel. In order to derive the correct surface forms, we need two association rules and one condition, given in (87):

(87) a. Associate [back] and [round] to the first vowel (initial association).
b. Associate [back] and [round] to the remaining vowels (spreading).
c. [+round] may not associate to non-initial non-high vowels (condition on target vowel).

Consider now the following derivation of the form *boruları* [borularɨ]:

(88)

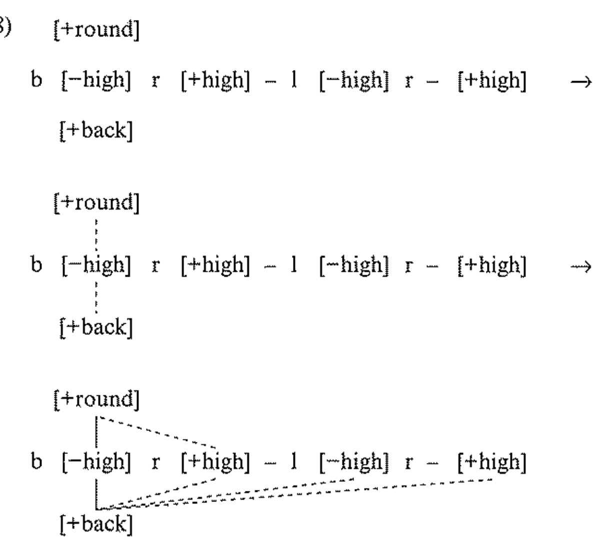

Observe that the impossibility of [+round] spreading to the low vowel of the plural suffix also blocks its spreading to the possessive suffix. Furthermore, in order to derive the surface representations, we must assume that all vowels not specified for [round] are automatically assigned the value [–round].

1.5 Summary

This chapter has been concerned with phonological segments, in particular their internal structure. We have shown that segments are not the basic atoms of phonological structure; much of §1.2 dealt with the evidence for postulating units smaller than the segment. These units, commonly referred to as features, are the true atoms of segmental structure, and are primarily motivated by the way in which particular groups of segments recurrently pattern together as natural classes in phonological generalisations (i.e. constraints and processes), generally as a result of some shared phonetic property, which in turn provides the basis for the definition of the relevant feature. For example, the vowels /u o ɔ/ in some language may function as a natural class on the basis of the shared phonetic property of lip-rounding, thus leading us to postulate a feature [round] in our model of representation. The simple enumeration of sets of segments would leave unexplained why these sets of segments form natural classes and why certain changes occur in certain environments.

In §1.3 we discussed a set of features that can be motivated on the basis of various well-known constraints and processes, starting with major class and manner features and then proceeding with vocalic features and consonantal features, respectively. It has not been our intention to suggest that this is the only possible set of features; indeed, we have considered a number of alternatives for features characterising particular phonetic and phonological dimensions. In the domain of vowel features, for example, we proposed that different features may be necessary for several closely related dimensions (height, tongue root position and tenseness). Our discussion of consonantal features included surveys of systems based on place of articulation as well as those based on the active articulator. Nonetheless, the features which we have discussed (or close equivalents) are encountered in most models of segmental representation, and as such provide a general grounding in feature theory. At the end of §1.3, we introduced the idea that the groups that we identified in our survey of features may themselves form part of the segmental structure, which thereby becomes a hierarchical structure, often represented as a tree. This allows rules and constraints to refer not only to the features themselves but also to organisational nodes dominating one or more features.

In §1.4 we showed that in some circumstances an individual feature may simultaneously be assigned to more than one segment. In other words, features are to some extent independent of the notion of segment itself; a single feature may have as its domain a consonant cluster, a syllable or even a complete word. The fact that the consonants in an English nasal + stop cluster such as [mp] or [ŋk] *necessarily* have the same place of articulation, for example, suggests that the two segments involved have only a single feature characterising the place of articulation. We introduced this autosegmental approach to the characterisation of assimilation and harmony processes by looking at tonal phenomena, and illustrated it further with several examples involving umlaut and vowel harmony, in the course of which we also examined some of the principles that govern the association between features and segments. In the following chapter, we will show that our concept of the phonological feature will have to be substantially refined to take account of the full range of phonological processes which we encounter in language; in addition, we will consider the formal nature of features in rather more detail than here.

1.6 Further reading

This chapter has been concerned with the phonological segment and its internal organisation in terms of features. The transcription system of the IPA (see Appendix) reflects a view of phonology in which segment-sized units are the central units in phonological analysis. In this view phonological properties are primarily properties of segments, not of larger units such as syllables.

The view that segments consist of smaller phonological units (§1.2) originates with the work of Trubetzkoy (1939), who proposes to group segments according to the distinctive oppositions in which they participate. Jakobson *et al.* (1951) and Jakobson and Halle (1956) formalise the notion in terms of binary distinctive features (§1.3). See also Ladefoged (1980), Halle (1983). In *SPE* we find a modified system and an emphasis on articulatory definitions. Hyman (1975) and Baltaxe (1978) contain extensive discussion of the binary tradition, and Keating (1988a) provides an overview. See also Jakobson and Waugh (1979) for a study of features in a somewhat wider context. Kaye (1989: ch. 2) discusses the distributional motivation for grouping segments into natural classes. Within the theory of **articulatory phonology**, an approach which we have not considered here, Browman and Goldstein (1986, 1989, 1992) consider the basic units of phonology to be 'articulatory gestures', rather than features. See also Clements (1992).

Major class features (§1.3.1) are discussed in Selkirk (1984b), Anderson and Ewen (1987), McCarthy (1988), Clements (1990), van der Hulst and Ewen

(1991), Kaisse (1992) and Hume and Odden (1995). Manner features include features for voicing (cf. also below) and continuancy (Davis 1989), for the distinctive properties of liquids (Spencer 1985; Lindau 1985; Walsh Dickey 1997) and for nasality (see Anderson 1976 and the papers in Ferguson *et al.* 1975 and Huffman and Krakow 1993). The latter property has received a great deal of attention, particularly in the context of harmony processes; cf. Herbert (1986) and Cohn (1990, 1993), Piggott (1988), Piggott and van der Hulst (1997). Vowel features (§1.3.2) are discussed in Lindau (1978), Wood (1982), Fischer-Jørgensen (1985), van der Hulst (1988), Clements (1989) and Odden (1991). Clements (1991) addresses the issue of vowel height and related dimensions, which he proposes to capture in terms of a single aperture dimension. See also Goad (1993). The discussion on this dimension (and on vowel features in general) often centres around the proper treatment of a large variety of vowel harmony systems; cf. the papers in Vago (1980), as well as van der Hulst (1988) and van der Hulst and van de Weijer (1995). There is extensive literature on ATR harmony systems (§1.3.3); see, e.g., Archangeli and Pulleyblank (1994). The four articulator-based features (§1.3.4) proposed by McCarthy (1988) characterising the major places of articulation ([labial], [coronal], [dorsal], [radical]) have also been adopted in the description of vowels (Sagey 1986; see also Hume 1992, Clements and Hume 1995). The status of coronals has been an issue of some debate; see for example the papers in Paradis and Prunet (1991), as well as Lahiri and Blumstein (1984), McCarthy and Taub (1992) and Hall (1997). For discussion of other consonantal features, see Hayward and Hayward (1989), Ladefoged and Maddieson (1989), Trigo (1991), McCarthy (1994), Ní Chiosáin (1994) and Rice (1995). On proposals for a feature [grave], see Hyman (1973), Vago (1976) and Odden (1978). There is a large amount of literature on the concept of grouping features into some kind of hierarchical structure (§1.3.5). Den Dikken and van der Hulst (1988) and McCarthy (1988) discuss a wide variety of proposals in this area (see also Odden 1991; Halle 1995; Pulleyblank 1995). Developments of the concept of grouping proposed in dependency phonology are discussed in van der Hulst (1995).

All discussion of features depends either on the analysis of phonological processes or on the reliability of information on segment inventories. With respect to the latter, Maddieson (1984) and Ladefoged and Maddieson (1996) are extremely valuable. The former study contains inventories of over 400 languages, and separate chapters dealing with the various classes of segments (e.g. vowels, liquids, etc.). The latter contains a wealth of information on potential contrasting properties of speech sounds. See also Crothers (1978).

In §1.4 we introduced the model of autosegmental phonology. This approach can be traced back to a theory which we have not considered in this chapter, that of the Firthian school of **prosodic analysis** (see e.g. Firth 1948 and the other papers in Palmer 1970; also Langendoen 1968), which is essentially non-segmental in approach. See the series of papers by Goldsmith (1992), Ogden and Local (1994) and Goldsmith (1994) for a discussion of the relationship between autosegmental phonology and prosodic analysis. Firthian phonology has also been combined with work in **declarative phonology** (Ogden 1999; see also Coleman 1998). An alternative model of 'non-segmental' phonology is offered by Griffen (1976). Analyses of the formal properties of autosegmental phonology are offered by Bird (1995; see also Goldsmith 1997), Kornai (1995) and Scobbie (1997).

For work on tone, see the papers in Fromkin (1978), as well as Maddieson (1978). The study of tonal phenomena has been of crucial importance for the development of autosegmental phonology; see the studies in Clements and Goldsmith (1984) and van der Hulst and Snider (1993). Tone features, because of their relationship to laryngeal features, are often discussed together with features for phonation (voicing types), leading to proposals for a unified set of laryngeal features for tonal contrasts and phonation contrasts. See Bao (1990), Duanmu (1990), Odden (1995) and Yip (1995) for further studies on tone.

On airstream mechanisms, see Catford (1977) for a description, and Ladefoged (1971), Lass (1984a), Ladefoged and Traill (1994) on possible feature systems.

Old English *i*-umlaut (§1.4.1) has been the subject of numerous diachronic and synchronic studies; see e.g. Lass and Anderson (1975), Hogg (1992a, b). For Turkish vowel harmony (§1.4.2), see Clements and Sezer (1982) and van der Hulst and van de Weijer (1991).

General overviews of non-linear phonology are given by van der Hulst and Smith (1982a, 1985), Goldsmith (1990) and many of the articles in Goldsmith (1995). A historical overview of some of the developments in non-linear phonology from the mid-eighties up to the mid-nineties can be found in van der Hulst and van de Weijer (1995).

For studies dealing with the structure of the segment in relation to phonological acquisition, see Vihman (1978), Levelt (1994), Stoel-Gammon and Stemberger (1994) and Rice and Avery (1997).

The structure of the segment in the phonology of sign languages is considered by Liddell and Johnson (1989), Sandler (1989), van der Hulst (1993) and Brentari (1999).