Phonology

Article · April 2002

Source: DBLP

Steven Bird
Charles Darwin University

Phonology

Steven Bird
University of Pennsylvania

2002

Abstract

Phonology is the systematic study of the sounds used in language, their internal structure, and their composition into syllables, words and phrases. Computational phonology is the application of formal and computational techniques to the representation and processing of phonological information. This chapter will present the fundamentals of descriptive phonology along with a brief overview of computational phonology.

1 Phonological contrast, the phoneme, and distinctive features

There is no limit to the number of distinct sounds that can be produced by the human vocal apparatus. However, this infinite variety is harnessed by human languages into **sound systems** consisting of a few dozen language-specific categories, or **phonemes**. An example of an English phoneme is *t*. English has a variety of *t*-like sounds, such as the aspirated t^h of *ten* the unreleased *t*˺ of *net*, and the flapped *ɾ* of *water* (in some dialects). In English, these distinctions are not used to differentiate words, and so we do not find pairs of English words which are identical but for their use of t^h versus *t*˺. (By comparison, in some other languages, such as Icelandic and Bengali, aspiration is contrastive.) Nevertheless, since these sounds (or **phones**, or **segments**) are phonetically similar, and since they occur in **complementary distribution** (i.e. disjoint contexts) and cannot differentiate words in English, they are all said to be **allophones** of the English phoneme *t*.

Of course, setting up a few allophonic variants for each of a finite set of phonemes does not account for the infinite variety of sounds mentioned above. If one were to record multiple instances of the same utterance by the single speaker, many small variations could be observed in loudness, pitch, rate, vowel quality, and so on. These variations arise because speech is a motor activity involving coordination of many independent articulators, and perfect repetition of any utterance is simply impossible. Similar variations occur between different speakers, since one person's vocal apparatus is different to the next person's (and this is how we can distinguish people's voices). So 10 people saying *ten* 10 times each will produce 100 distinct acoustic records for the *t* sound. This diversity of tokens associated with a single type is sometimes referred to as **free variation**.

Above, the notion of phonetic similarity was used. The primary way to judge the similarity of phones is in terms of their **place** and **manner** of articulation. The consonant chart of the International Phonetic Alphabet (IPA) tabulates phones in this way, as shown in Figure 1. The IPA provides symbols for all sounds that are contrastive in at least one language.

The major axes of this chart are for place of articulation (horizontal), which is the location in the oral cavity of the primary constriction, and manner of articulation (vertical), the nature and degree of that constriction. Many cells of the chart contain two consonants, one **voiced** and the other **unvoiced**. These complementary properties are usually expressed as opposite values of a **binary feature** [±voiced].

A more elaborate model of the similarity of phones is provided by the theory of **distinctive features**. Two phones are considered more similar to the extent that they agree on the value of their features. A set of distinctive features and their values for five different phones is shown in (1). (Note that many of the features have an extended technical definition, for which it is necessary to consult a textbook.)

THE INTERNATIONAL PHONETIC ALPHABET (revised to 1993)

CONSONANTS (PULMONIC)

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

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

Figure 1: Pulmonic Consonants from the International Phonetic Alphabet

(1)

	t	z	m	l	i
anterior	+	+	+	+	–
coronal	+	+	–	+	–
labial	–	–	+	–	–
distributed	–	–	–	–	–
consonantal	+	+	+	+	–
sonorant	–	–	+	+	+
voiced	–	+	+	+	+
approximant	–	–	–	+	+
continuant	–	+	–	+	+
lateral	–	–	–	+	–
nasal	–	–	+	–	–
strident	–	+	–	–	–

Statements about the distribution of phonological information, usually expressed with rules or constraints, often apply to particular subsets of phones. Instead of listing these sets, it is virtually always simpler to list two or three feature values which pick out the required set. For example [+labial,–continuant] picks out *b*, *p*, and *m*, shown in the top left corner of Figure 1. Sets of phones which can be picked out in this way are called **natural classes**, and phonological analyses can be evaluated in terms of their reliance on natural classes. How can we express these analyses? The rest of this chapter discusses some key approaches to this question.

Unfortunately, as with any introductory chapter like this one, it will not be possible to cover many important topics of interests to phonologists, such as acquisition, diachrony, orthography, universals, sign language phonology, the phonology/syntax interface, systems of intonation and stress, and many others besides. However, numerous bibliographic references are supplied at the end of the chapter, and readers may wish to consult these other works.

2 Early Generative Phonology

Some key concepts of phonology are best introduced by way of simple examples involving real data. We begin with some data from Russian in (2). The example shows some nouns, in nominative and dative cases, transcribed using the International Phonetic Alphabet. Note that *x* is the symbol for a voiceless velar fricative (e.g. the *ch* of Scottish *loch*).

(2)

Nominative	Dative	Gloss
xlep	xlebu	‘bread’
grop	grobu	‘coffin’
sat	sadu	‘garden’
prut	prudu	‘pond’
rok	rogu	‘horn’
ras	razu	‘time’

Observe that the dative form involves suffixation of *-u*, and a change to the final consonant of the nominative form. In (2) we see four changes: *p* becomes *b*, *t* becomes *d*, *k* becomes *g*, and *s* becomes *z*.

Where they differ is in their **voicing**; for example, *b* is a **voiced** version of *p*, since *b* involves periodic vibration of the vocal folds, while *p* does not. The same applies to the other pairs of sounds. Now we see that the changes we observed in (2) are actually quite systematic. Such systematic patterns are called **alternations**, and this particular one is known as a **voicing alternation**. We can formulate this alternation using a **phonological rule** as follows:

(3)

$$\begin{bmatrix} \text{C} \\ -\text{voiced} \end{bmatrix} \rightarrow [+\text{voiced}] \,/\, \underline{\quad}\,\text{V}$$

A consonant becomes voiced in the presence of a following vowel

Rule (3) uses the format of early generative phonology. In this notation, C represents any consonant and V represents any vowel. The rule says that, if a voiceless consonant appears in the **phonological environment** ‘___ V’ (i.e. preceding a vowel), then the consonant becomes voiced. By default, vowels have the feature [+voiced], and so can make the observation that the consonant **assimilates** the voicing feature of the following vowel.

One way to see if our analysis generalises is to check for any nominative forms that end in a voiced consonant. We expect this consonant to stay the same in the dative form. However, it turns out that we do not find any nominative forms ending in a voiced consonant. Rather, we see the pattern in example (4). (Note that č is an alternative symbol for IPA ʧ).

(4)

Nominative	Dative	Gloss
čerep	čerepu	‘skull’
xolop	xolopu	‘bondman’
trup	trupu	‘corpse’
cvet	cvetu	‘colour’
les	lesu	‘forest’
porok	poroku	‘vice’

For these words, the voiceless consonants of the nominative form are unchanged in the dative form, contrary to our rule (3). These cannot be treated as exceptions, since this second pattern is quite pervasive. A solution is to construct an artificial form which is the dative wordform minus the *-u* suffix. We will call this the **underlying form** of the word. Example (5) illustrates this for two cases:

(5)

Underlying	Nominative	Dative	Gloss
prud	prut	prudu	‘pond’
cvet	cvet	cvetu	‘colour’

Now we can account for the dative form simply by suffixing the *-u*. We account for the nominative form with the following **devoicing rule**:

(6)

$$\begin{bmatrix} \text{C} \\ +\text{voiced} \end{bmatrix} \rightarrow [-\text{voiced}] \, / \, \_\_\#$$

A consonant becomes devoiced word-finally

This rule states that a voiced consonant is devoiced (i.e. [+voiced] becomes [–voiced]) if the consonant is followed by a word boundary (symbolised by #). It solves a problem with rule 3 which only accounts for half of the data. Rule 6 is called a **neutralisation** rule, because the **voicing contrast** of the underlying form is removed in the nominative form. Now the analysis accounts for all the nominative and dative forms. Typically, rules like (6) can simultaneously employ several of the distinctive features from (1).

Observe that our analysis involves a certain degree of **abstractness**. We have constructed a new **level of representation** and drawn inferences about the **underlying forms** by inspecting the observed **surface forms**.

To conclude the development so far, we have seen a simple kind of **phonological representation** (namely sequences of alphabetic symbols, where each stands for a bundle of distinctive features), a distinction between levels of representation, and rules which account for the relationship between the representations on various levels. One way or another, most of phonology is concerned about these three things: representations, levels, and rules.

Finally, let us consider the plural forms shown in example (7). The plural morpheme is either *-a* or *-y*.

(7)

Singular	Plural	Gloss
xlep	xleba	‘bread’
grop	groby	‘coffin’
čerep	čerepa	‘skull’
xolop	xolopy	‘bondman’
trup	trupy	‘corpse’
sat	sady	‘garden’
prut	prudy	‘pond’
cvet	cveta	‘colour’
ras	razy	‘time’
les	lesa	‘forest’
rok	roga	‘horn’
porok	poroky	‘vice’

The phonological environment of the suffix provides us with no way of predicting which allomorph is chosen. One solution would be to enrich the underlying form once more (for example, we could include the plural suffix in the underlying form, and then have rules to delete it in all cases but the plural). A better approach in this case is to distinguish two **morphological classes**, one for nouns taking the *-y* plural, and one for nouns taking the *-a* plural. This information would then be an idiosyncratic property of each lexical item, and a morphological rule would be responsible for the choice between the *-y* and *-a* **allomorphs**. A full account of this data, then, must involve phonological, morphological and lexical modules of a grammar.

As another example, let us consider the vowels of Turkish. These vowels are tabulated below, along with a decomposition into distinctive features: [high], [back] and [round]. The features [high] and [back] relate to the position of the tongue body in the oral cavity. The feature [round] relates to the rounding of the lips, as in the English *w* sound.[1]

(8)

	u	o	ü	ö	ı	a	i	e
high	+	–	+	–	+	–	+	–
back	+	+	–	–	+	+	–	–
round	+	+	+	+	–	–	–	–

[1]Note that there is a distinction made in the Turkish alphabet between the dotted *i* and the dotless *ı*. This *ı* is a high, back, unrounded vowel that does not occur in English.

Consider the following Turkish words, paying particular attention to the four versions of the possessive suffix. Note that similar data are discussed in chapter 2.

(9)

ip	‘rope’	ipin	‘rope’s’
kız	‘girl’	kızın	‘girl’s’
yüz	‘face’	yüzün	‘face’s’
pul	‘stamp’	pulun	‘stamp’s’
el	‘hand’	elin	‘hand’s’
çan	‘bell’	çanın	‘bell’s’
köy	‘village’	köyün	‘village’s’
son	‘end’	sonun	‘end’s’

The possessive suffix has the forms *in*, *ın*, *ün* and *un*. In terms of the distinctive feature chart in (8), we can observe that the suffix vowel is always [+high]. The other features of the suffix vowel are copied from the stem vowel. This copying is called **vowel harmony**. Let us see how this behaviour can be expressed using a phonological rule. To do this, we assume that the vowel of the possessive affix is only specified as [+high] and is **underspecified** for its other features. In the following rule, *C* denotes any consonant, and the Greek letter variables range over the + and – values of the feature.

(10)

$$\begin{bmatrix} V \\ +\text{high} \end{bmatrix} \longrightarrow \begin{bmatrix} \alpha\text{back} \\ \beta\text{round} \end{bmatrix} \Big/ \begin{bmatrix} \alpha\text{back} \\ \beta\text{round} \end{bmatrix} C^* \_\_$$

A high vowel assimilates to the backness and rounding of the preceding vowel

So long as the stem vowel is specified for the properties [high] and [back], this rule will make sure that they are copied onto the affix vowel. However, there is nothing in the rule formalism to stop the variables being used in inappropriate ways (e.g. α back $\rightarrow$ α round). So we can see that the rule formalism does not permit us to express the notion that certain features are **shared** by more than one segment. Instead, we would like to be able to represent the sharing explicitly, as follows, where $\pm$H abbreviates [$\pm$high], an underspecified vowel position:

(11)

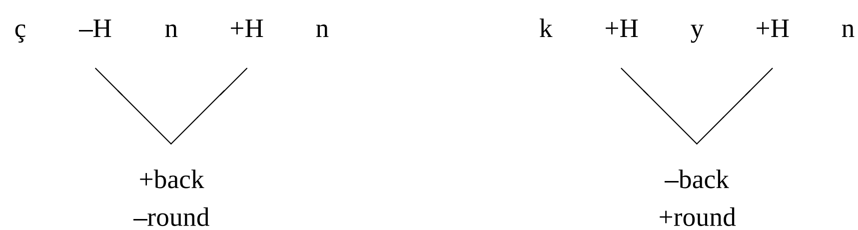

The lines of this diagram indicate that the backness and roundness properties are shared by both vowels in a word. A single vowel property (or type) is manifested on two separate vowels (tokens).

Entities like [+back,–round] that function over extended regions are often referred to as **prosodies**, and this kind of picture is sometimes called a **non-linear** representation. Many phonological models use non-linear representations of one sort or another. Here we shall consider one particular model, namely **autosegmental phonology**, since it is the most widely used non-linear model. The term comes from ‘autonomous + segment’, and refers to the autonomous nature of segments (or certain groups of features) once they have been liberated from one-dimensional strings.

3 Autosegmental Phonology

In autosegmental phonology, diagrams like those we saw above are known as **charts**. A chart consists of two or more **tiers**, along with some **association lines** drawn between the autosegments on those tiers. The

no-crossing constraint is a stipulation that association lines are not allowed to cross, ensuring that association lines can be interpreted as asserting some kind of temporal overlap or inclusion. **Autosegmental rules** are procedures for converting one representation into another, by adding or removing association lines and autosegments. A rule for Turkish vowel harmony is shown below on the left in (12), where *V* denotes any vowel, and the dashed line indicates that a new association is created. This rule applies to the representation in the middle, to yield the one on the right.

(12)

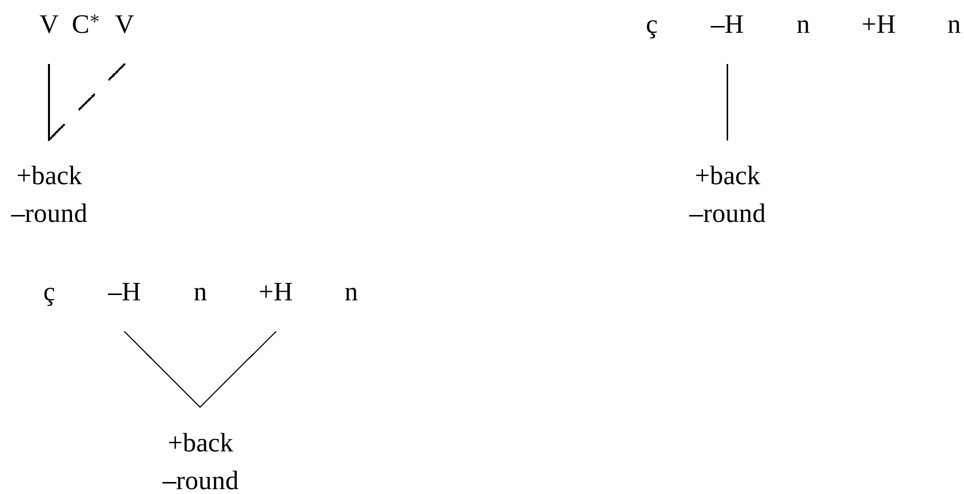

In order to fully appreciate the power of autosegmental phonology, we will use it to analyse some data from an African tone language. Consider the data in Table 1. Twelve nouns are listed down the left side, and the isolation form and five contextual forms are provided across the table. The line segments indicate voice pitch (the fundamental frequency of the voice); dotted lines are for the syllables of the context words, and full lines are for the syllables of the target word, as it is pronounced in this context. At first glace this data seems bewildering in its complexity. However, we will see how autosegmental analysis reveals the simple underlying structure of the data.

Looking across the table, observe that the contextual forms of a given noun are quite variable. For example *bulɑli* appears as –‾–, ‾–_, –‾‾, and ‾––.

We could begin the analysis by identifying all the levels (here there are five), assigning a name or number to each, and looking for patterns. However, this approach does not capture the relative nature of tone, where ––_ is not distinguished from ‾‾–. Instead, our approach just has to be sensitive to *differences* between adjacent tones. So these distinct tone sequences could be represented identically as $+1$, -2, since we go up a small amount from the first to the second tone ($+1$), and then down a larger amount (-2). In autosegmental analysis, we treat **contour tones** as being made up of two or more **level tones** compressed into the space of a single syllable. Therefore, we can treat –⌝ as another instance of $+1$, -2. Given our autosegmental perspective, a sequence of two or more identical tones corresponds to a single spread tone. This means that we can collapse sequences of like tones to a single tone.[2] When we retranscribe our data in this way, some interesting patterns emerge.

First, by observing the raw frequency of these intertone intervals, we see that -2 and $+1$ are by far the most common, occurring 63 and 39 times respectively. A -1 difference occurs 8 times, while a $+2$ difference is very rare (only occurring 3 times, and only in phrase-final contour tones). This patterning is characteristic of a **terrace tone language**. In analysing such a language, phonologists typically propose an inventory of just two tones, H (high) and L (low), where these might be represented featurally as [±hi]. In such a model, the tone sequence HL corresponds to ‾–, a pitch difference of -2.

[2] This assumption cannot be maintained in more sophisticated approaches involving lexical and prosodic domains. However, it is a very useful simplifying assumption for the purposes of this presentation.

Wordform	A. ___ isolation	B. **i ___** 'his ...'	C. **am goro ___** 'your (pl) brother's ...'	D. **___ kũ** 'one ...'	E. **am ___ wo dɔ** 'your (pl) ... is there'	F. **jiine ___ ni** 'that ...'
1. **bɑkɑ** 'tree'	–¯	⋯¯–	⋱⋯⋯–¯	–¯⋯	⋱–¯⋯⋯	⋯⋯¯–⋯
2. **sɑkɑ** 'comb'	–↘	⋯¯–	⋱⋯⋯–↘	–¯⋯	⋱–¯⋯⋯	⋯⋯¯–⋯
3. **buri** 'duck'	¯_	⋯¯_	⋱⋯⋯¯_	¯¯⋯		⋯⋯¯¯⋯
4. **siri** 'goat'	¯↘	⋯¯↘	⋱⋯⋯¯↘	¯¯⋯	⋱¯¯⋯⋯	⋯⋯¯¯⋯
5. **gɑdo** 'bed'	––	⋯––	⋱⋯⋯––	––⋯		⋯⋯––⋯
6. **gɔrɔ** 'brother'	––	⋯–_		––⋯	⋱––⋯⋯	
7. **cɑ** 'dog'	↗	⋯_	⋱⋯⋯↗	–⋯	⋱–⋯⋯	⋯⋯¯⋯
8. **ni** 'mother'	–	⋯–	⋱⋯⋯¯	¯⋯		⋯⋯–⋯
9. **jɔkɔrɔ** 'chain'	––¯	⋯¯_¯	⋱⋯⋯––¯	––¯⋯	⋱––¯⋯⋯	⋯⋯¯_¯⋯
10. **tokoro** 'window'	–––	⋯–––	⋱⋯⋯–––	–––⋯	⋱–––⋯⋯	⋯⋯–––⋯
11. **bulɑli** 'iron'	–¯_	⋯¯–_	⋱⋯⋯–¯_	–¯¯⋯	⋱–¯¯⋯⋯	⋯⋯¯––⋯
12. **misini** 'needle'	¯¯_	⋯¯¯_	⋱⋯⋯¯¯_	¯¯¯⋯	⋱¯¯¯⋯⋯	⋯⋯¯¯¯⋯

Table 1: Tone Data from Chakosi (Ghana)

In terrace tone languages, an H tone does not achieve its former level after an L tone, so HLH is **phonetically realized** as ¯_–, (instead of ¯_¯). This kind of H-lowering is called **automatic downstep**. A pitch difference of $+1$ corresponds to an LH tone sequence. With this model, we already account for the prevalence of the -2 and $+1$ intervals. What about -1 and $+2$?

As we will see later, the -1 difference arises when the middle tone of ¯_– (HLH) is deleted, leaving just ¯–. In this situation we write H!H, where the exclamation mark indicates the lowering of the following H due to a deleted (or **floating** low tone). This kind of H-lowering is called **conditioned downstep**. The rare $+2$ difference only occurs for an LH contour; we can assume that automatic downstep only applies when a LH sequence is linked to two separate syllables (–¯) and not when the sequence is linked to a single syllable (↗).

To summarise these conventions, we associate the pitch differences to tone sequences as shown in (13). Syllable boundaries are marked with a dot.

(13)

Interval	-2	-1	$+1$	$+2$
Pitches	¯_	¯–	–¯	↗
Tones	H.L	H.!H	L.H	LH

Now we are in a position to provide tonal transcriptions for the forms in Table 1. Example (14) gives the transcriptions for the forms involving *bulɑli*. Tones corresponding to the noun are underlined.

(14) **Transcriptions of bulɑli 'iron'**

bulɑli	'iron'	–¯_	<u>L.H.L</u>
i bulɑli	'his iron'	⋯¯–_	H.<u>H.!H.L</u>
am goro bulɑli	'your (pl) brother's iron'	⋱⋯⋯–¯_	HL.L.L.<u>L.H.L</u>
bulɑli kũ	'one iron'	–¯¯⋯	<u>L.H.H</u>.L
am bulɑli wo dɔ	'your (pl) iron is there'	⋱–¯¯⋯⋯	HL.<u>L.H.H</u>.!H.L
jiine bulɑli ni	'that iron'	⋯⋯¯––⋯	L.H.<u>H.!H.H</u>.L

Looking down the right hand column of (14) at the underlined tones, observe again the diversity of **surface forms** corresponding to the single lexical item. An autosegmental analysis is able to account for all this variation with a single spreading rule.

(15) **High Tone Spread**

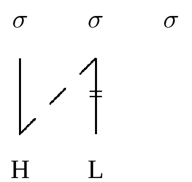

A high tone spreads to the following (non-final) syllable, delinking the low tone

Rule (15) applies to any sequence of three syllables (σ) where the first is linked to an H tone and the second is linked to an L tone. The rule spreads H to the right, delinking the L. Crucially, the L itself is not deleted, but remains as a **floating tone**, and continues to influence surface tone as downstep. Example (16) shows the application of the H spread rule to forms involving *bulɑli*. The first row of autosegmental diagrams shows the underlying forms, where *bulɑli* is assigned an LHL **tone melody**. In the second row, we see the result of applying H spread. Following standard practice, the floating low tones are circled. Where a floating L appears between two H tones, it gives rise to downstep. The final assignment of tones to syllables and the position of the downsteps are shown in the last row of the table.

(16)

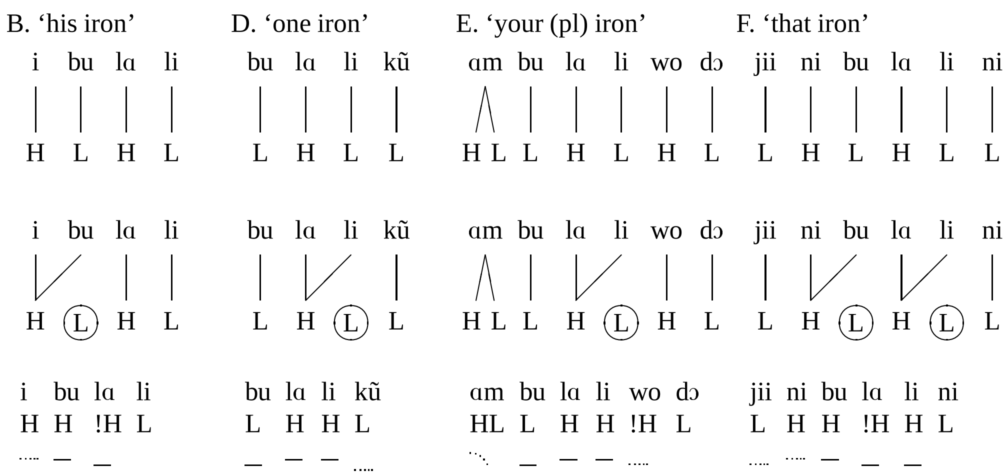

Example (16) shows the power of autosegmental phonology – together with suitable underlying forms and appropriate principles of phonetic interpretation – in analysing complex patterns with simple rules. Space precludes a full analysis of the data; interested readers can try hypothesising underlying forms for the other words, along with new rules, to account for the rest of the data in Table 1.

The preceding discussion of segmental and autosegmental phonology highlights the multi-linear organisation of phonological representations, which derives from the temporal nature of the speech stream. Phonological representations are also organised hierarchically. We already know that phonological information comprises words, and words, phrases. This is one kind of hierarchical organisation of phonological information. But phonological analysis has also demonstrated the need for other kinds of hierarchy, such as the **prosodic hierarchy**, which builds structure involving syllables, feet and intonational phrases above the segment level, and **feature geometry**, which involves hierarchical organisation beneath the level of the segment. Phonological rules and constraints can refer to the prosodic hierarchy in order to account for the observed **distribution** of phonological information across the linear sequence of segments. Feature geometry serves the dual purpose of accounting for the inventory of contrastive sounds available to a language, and for the alternations we can observe. Here we will consider just one level of phonological hierarchy, namely the syllable.

4 Syllable Structure

Syllables are a fundamental organisational unit in phonology. In many languages, phonological alternations are sensitive to syllable structure. For example, *t* has several **allophones** in English, and the choice of allophone depends on phonological context. For example, in many English dialects, *t* is pronounced as the flap [ɾ] between vowels, as in *water*. Two other variants are shown in (17), where the phonetic transcription is given in brackets, and syllable boundaries are marked with a dot.

(17) a. atlas [ætˀ.ləs]
b. cactus [kæk.tʰəs]

Native English syllables cannot begin with *tl*, and so the *t* of *atlas* is syllabified with the preceding vowel. Syllable final *t* is regularly glottalised or unreleased in English, while syllable initial *t* is regularly aspirated. Thus we have a natural explanation for the patterning of these allophones in terms of syllable structure.

Other evidence for the syllable comes from loanwords. When words are borrowed into one language from another, they must be adjusted so as to conform to the legal sound patterns (or **phonotactics**) of the host language. For example, consider the following borrowings from English into Dschang, a language of Cameroon (Bird, 1999).

(18) afruwa *flower*, akalatusi *eucalyptus*, alɛsa *razor*, alɔba *rubber*, aplɛŋgɛ *blanket*, asəkuu *school*, cɛɛn *chain*, dəək *debt*, kapinda *carpenter*, kɛsiŋ *kitchen*, kuum *comb*, laam *lamp*, lɛsi *rice*, luum *room*, mbasəku *bicycle*, mbrusi *brush*, mbərəək *brick*, mɛta *mat*, mɛtərasi *mattress*, ŋglasi *glass*, ɲjakasi *jackass*, mɛtisi *match* nubatisi *rheumatism*, pɔkɛ *pocket* ŋgalɛ *garden*, səsa *scissors*, tɛwɛlɛ *towel*, wasi *watch*, ziiŋ *zinc*,

In Dschang, the **syllable canon** is much more restricted than in English. Consider the patterning of *t*. This segment is illegal in syllable-final position. In technical language, we would say that alveolars are not **licensed** in the syllable coda. In mɛta *mat*, a vowel is inserted, making the *t* into the initial segment of the next syllable. For dəək *debt*, the place of articulation of the *t* is changed to velar, making it a legal syllable-final consonant. For aplɛŋgɛ *blanket*, the final *t* is deleted. Many other adjustments can be seen in (18), and most of them can be explained with reference to syllable structure.

A third source of evidence for syllable structure comes from morphology. In Ulwa, a Nicaraguan language, the position of the possessive **infix** is sensitive to syllable structure. The Ulwa syllable canon is (C)V(V|C)(C), and any **intervocalic** consonant (i.e. consonant between two vowels) is syllabified with the following syllable, a universal principle known as **onset maximisation**. Consider the Ulwa data in (19).

(19)

Word	Possessive	Gloss	Word	Possessive	Gloss
bɑɑ	bɑɑ.**kɑ**	'excrement'	bi.lɑm	bi.lɑm.**kɑ**	'fish'
dii.muih	dii.**kɑ**.muih	'snake'	gɑɑd	gɑɑd.**kɑ**	'god'
ii.bin	ii.**kɑ**.bin	'heaven'	ii.li.lih	ii.**kɑ**.li.lih	'shark'
kɑh.mɑ	kɑh.**kɑ**.mɑ	'iguana'	kɑ.pɑk	kɑ.pɑk.**kɑ**	'manner'
lii.mɑ	lii.**kɑ**.mɑ	'lemon'	mis.tu	mis.**kɑ**.tu	'cat'
on.yɑn	on.**kɑ**.yɑn	'onion'	pɑu.mɑk	pɑu.**kɑ**.mɑk	'tomato'
sik.bilh	sik.**kɑ**.bilh	'horsefly'	tɑim	tɑim.**kɑ**	'time'
tɑi.tɑi	tɑi.**kɑ**.tɑi	'grey squirrel'	uu.mɑk	uu.**kɑ**.mɑk	'window'
wɑi.ku	wɑi.**kɑ**.ku	'moon, month'	wɑ.sɑ.lɑ	wɑ.sɑ.**kɑ**.lɑ	'possum'

Observe that the infix appears at a syllable boundary, and so we can already state that the infix position is sensitive to syllable structure. Any analysis of the infix position must take **syllable weight** into consideration. Syllables having a single short vowel and no following consonants are defined to be **light**. (The presence of onset consonants is irrelevant to syllable weight.) All other syllables, i.e. those which have two vowels, or a single long vowel, or a final consonant, are defined to be **heavy**; e.g. *kah*, *kaa*, *muih*, *bilh*, *ii*, *on*. Two common phonological representations for this syllable structure are the onset-rhyme model, and the moraic model. Representations for the syllables just listed are shown in (20). In these diagrams, σ denotes a syllable, O onset, R rhyme, N nucleus, C coda and μ *mora* (the traditional, minimal unit of syllable weight).

(20) a. **The Onset-Rhyme Model of Syllable Structure**

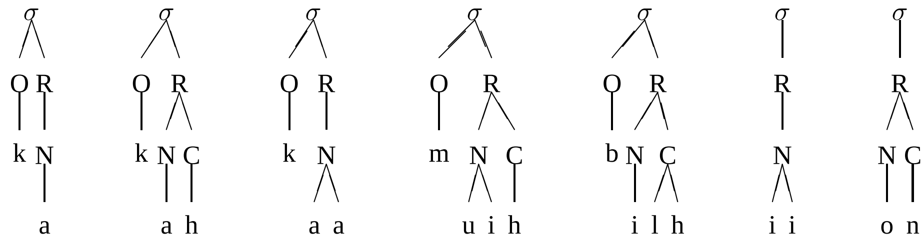

b. **The Moraic Model of Syllable Structure**

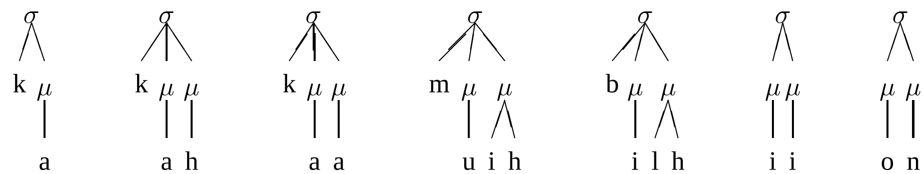

In the onset-rhyme model (20a), consonants coming before the first vowel are linked to the onset node, and the rest of the material comes under the rhyme node.[3] A rhyme contains an obligatory nucleus and an optional coda. In this model, a syllable is said to be heavy if and only if its rhyme or its nucleus are branching.

In the moraic mode (20b), any consonants that appear before the first vowel are linked directly to the syllable node. The first vowel is linked to its own mora node (symbolised by μ), and any remaining material is linked to the second mora node. A syllable is said to be heavy if and only if it has more than one mora.

These are just two of several ways that have been proposed for representing syllable structure. Now the syllables constituting a word can now be linked to higher levels of structure, such as the *foot* and the *prosodic word*. For now, it is sufficient to know that such higher levels exist, and that we have a way to represent the binary distinction of syllable weight.

Now we can return to the Ulwa data, from example (19). A relatively standard way to account for the infix position is to stipulate that the first light syllable, if present, is actually invisible to the rules which assign syllables to higher levels; such syllables are said to be **extra-metrical**. They are a sort of ‘upbeat’ to the word, and are often associated with the preceding word in continuous speech. Given these general principles concerning hierarchical structure, we can simply state that the Ulwa possessive affix is infixed after the first syllable.[4]

In the foregoing discussion, I hope to have revealed many interesting issues which are confronted by phonological analysis, without delving too deeply into the abstract theoretical constructs which phonologists have proposed. Theories differ enormously in their organisation of phonological information and the ways in which they permit this information to be subjected to rules and constraints, and the way the information is used in a **lexicon** and an overarching **grammatical framework**. Some of these theoretical frameworks include: lexical phonology, underspecification phonology, government phonology, declarative phonology, and optimality theory. For more information about these, please see §5.3 for literature references.

5 Computational phonology

When phonological information is treated as a string of atomic symbols, it is immediately amenable to processing using existing models. A particularly successful example is the work on finite state transducers

[3] Two syllables usually have to agree on the material in their rhyme constituents in order for them to be considered rhyming, hence the name.

[4] A better analysis of the Ulwa infixation data involves reference to **metrical feet**, phonological units above the level of the syllable. This is beyond the scope of the current chapter however.

(see chapter 21). However, phonologists abandoned linear representations in the 1970s, and so we will consider some computational models that have been proposed for multi-linear, hierarchical, phonological representations. It turns out that these pose some interesting challenges.

Early models of generative phonology, like that of the Sound Pattern of English (SPE), were sufficiently explicit that they could be implemented directly. A necessary first step in implementing many of the more recent theoretical models is to formalise them, and to discover the intended semantics of some subtle, graphical notations. A practical approach to this problem has been to try to express phonological information using existing, well-understood computational models. The principal models are finite state devices and attribute-value matrices.

5.1 Finite state models of non-linear phonology

Finite state machines cannot process structured data, only strings, so special methods are required for these devices to process complex phonological representations. All approaches involve a many-to-one mapping from the parallel layers of representation to a single machine. There are essentially three places where this many-to-one mapping can be situated. The first approach is to employ multi-tape machines (Kay, 1987). Each tier is represented as a string, and the set of strings is processed simultaneously by a single machine. The second approach is to map the multiple layers into a single string, and to process that with a conventional single-tape machine (Kornai, 1995). The third approach is to encode each layer itself as a finite state machine, and to combine the machines using automaton intersection (Bird and Ellison, 1994).

This work demonstrates how representations can be compiled into a form that can be directly manipulated by finite state machines. Independently of this, we also need to provide a means for phonological generalisations (such as rules and constraints) to be given a finite state interpretation. This problem is well studied for the linear case, and compilers exist that will take a rule formatted somewhat like the SPE style and produce an equivalent finite state transducer. Whole constellations of ordered rules or optimality-theoretic constraints can also be compiled in this way. However, the compilation of rules and constraints involving autosegmental structures is still largely un-addressed.

The finite state approaches emphasise the temporal (or left-to-right) ordering of phonological representations. In contrast, attribute-value models emphasise the hierarchical nature of phonological representations.

5.2 Attribute-value matrices

The success of attribute-value matrices (AVMs) as a convenient formal representation for constraint-based approaches to syntax (see chapter 3), and concerns about the formal properties of non-linear phonological information, led some researchers to apply AVMs to phonology. Hierarchical structures can be represented using AVM nesting, as shown in (21a), and autosegmental diagrams can be encoded using AVM indexes, as shown in (21b).

(21) a.
$$\begin{bmatrix} \text{onset} & \langle \text{k} \rangle \\ \text{rhyme} & \begin{bmatrix} \text{nucleus} & \langle \text{u, i} \rangle \\ \text{coda} & \langle \text{h} \rangle \end{bmatrix} \end{bmatrix}$$

b.
$$\begin{bmatrix} \text{syllable} & \langle \text{i}_{\boxed{1}}, \text{bu}_{\boxed{2}}, \text{l}\text{ɑ}_{\boxed{3}}, \text{li}_{\boxed{4}} \rangle \\ \text{tone} & \langle \text{H}_{\boxed{5}}, \text{L}_{\boxed{6}}, \text{H}_{\boxed{7}}, \text{L}_{\boxed{8}} \rangle \\ \text{associations} & \left\{ \langle \boxed{1}, \boxed{5} \rangle, \langle \boxed{2}, \boxed{5} \rangle, \langle \boxed{3}, \boxed{7} \rangle, \langle \boxed{4}, \boxed{8} \rangle \right\} \end{bmatrix}$$

AVMs permit re-entrancy by virtue of the numbered indexes, and so parts of a hierarchical structure can be shared. For example, (22a) illustrates a consonant shared between two adjacent syllables, for the word *cousin* (this kind of double affiliation is called **ambisyllabicity**). Example (22b) illustrates shared structure within a single syllable *full*, to represent the **coarticulation** of the onset consonant with the vowel.

(22) a.

$$\begin{bmatrix}\text{syllable} & \left\langle \begin{bmatrix}\text{onset} & \langle \text{k} \rangle \\ \text{rhyme} & \begin{bmatrix}\text{nucleus} & \langle \text{ʌ} \rangle \\ \text{coda} & \langle \text{z}\boxed{1} \rangle \end{bmatrix}\end{bmatrix} \begin{bmatrix}\text{onset} & \langle \boxed{1} \rangle \\ \text{rhyme} & \begin{bmatrix}\text{nucleus} & \langle \text{ə} \rangle \\ \text{coda} & \langle \text{n} \rangle \end{bmatrix}\end{bmatrix} \right\rangle \end{bmatrix}$$

b.

$$\begin{bmatrix}\text{onset} & \begin{bmatrix}\text{consonantal} & \begin{bmatrix}\text{grave} & + \\ \text{compact} & - \end{bmatrix} \\ \text{source} & \begin{bmatrix}\text{voice} & - \\ \text{continuant} & + \end{bmatrix} \\ \text{vocalic} & \boxed{1}\begin{bmatrix}\text{grave} & + \\ \text{height} & \text{close} \end{bmatrix}\end{bmatrix} \\ \text{rhyme} & \begin{bmatrix}\text{nucleus} \mid \text{vocalic} & \boxed{1} \\ \text{coda} & \begin{bmatrix}\text{consonantal} & \begin{bmatrix}\text{grave} & - \\ \text{compact} & - \end{bmatrix} \\ \text{vocalic} & \begin{bmatrix}\text{grave} & + \\ \text{compact} & + \end{bmatrix} \\ \text{source} \mid \text{nasal} & 1 \end{bmatrix}\end{bmatrix}\end{bmatrix}$$

Given such flexible and extensible representations, rules and constraints can manipulate and enrich the phonological information. Computational implementations of these AVM models have been used in speech synthesis systems.

5.3 Computational Tools for Phonological Research

Once a phonological model is implemented, it ought to be possible to use the implementation to evaluate theories against data sets. A phonologist's workbench should help people to 'debug' their analyses and spot errors before going to press with an analysis. Developing such tools is much more difficult than it might appear.

First, there is no agreed method for modelling non-linear representations, and each proposal has shortcomings. Second, processing data sets presents its own set of problems, having to do with tokenisation, symbols which are ambiguous as to their featural decomposition, symbols marked as uncertain or optional, and so on. Third, some innocuous looking rules and constraints may be surprisingly difficult to model, and it might only be possible to approximate the desired behaviour. Additionally, certain universal principles and tendencies may be hard to express in a formal manner. A final, pervasive problem is that symbolic transcriptions may fail to adequately reflect linguistically significant acoustic differences in the speech signal.

Nevertheless, whether the phonologist is sorting data, or generating helpful tabulations, or gathering statistics, or searching for a (counter-)example, or verifying the transcriptions used in a manuscript, the principal challenge remains a computational one. Recently, new directed-graph models (e.g. Emu, MATE, Annotation Graphs) appear to provide good solutions to the first two problems, while new advances on finite-state models of phonology are addressing the third problem. Therefore, we have grounds for confidence that there will be significant advances on these problems in the near future.

Further reading and relevant resources

The phonology community is served by an excellent journal *Phonology*, published by Cambridge University Press. Useful textbooks and collections include: (Katamba, 1989; Frost and Katz, 1992; Kenstowicz, 1994; Goldsmith, 1995; Clark and Yallop, 1995; Gussenhoven and Jacobs, 1998; Goldsmith, 1999; Roca et al., 1999; Jurafsky and Martin, 2000; Harrington and Cassidy, 2000). Oxford University Press publishes a series *The Phonology of the World's Languages*, including monographs on Armenian (Vaux, 1998), Dutch (Booij, 1995), English (Hammond, 1999), German (Wiese, 1996), Hungarian (Siptár and Törkenczy, 2000). Kimatuumbi (Odden, 1996), Norwegian (Kristoffersen, 1996), Portuguese (Mateus and d'Andrade, 2000), and Slovak (Rubach, 1993). An important forthcoming survey of phonological variation is the Atlas of North American English (Labov et al., 2001).

Phonology is the oldest discipline in linguistics and has a rich history. Some historically important works include: (Joos, 1957; Pike, 1947; Firth, 1948; Bloch, 1948; Hockett, 1955; Chomsky and Halle, 1968). The most comprehensive history of phonology is (Anderson, 1985).

Useful resources for phonetics include: (Catford, 1988; Laver, 1994; Ladefoged and Maddieson, 1996; Stevens, 1999; International Phonetic Association, 1999; Ladefoged, 2000; Handke, 2001), and the homepage of the International Phonetic Association `http://www.arts.gla.ac.uk/IPA/ipa.html`. The phonology/phonetics interface is an area of vigorous research, and the main focus of the *Laboratory Phonology* series published by Cambridge: (Kingston and Beckman, 1991; Docherty and Ladd, 1992; Keating, 1994; Connell and Arvaniti, 1995; Broe and Pierrehumbert, 2000). Two interesting essays on the relationship between phonetics and phonology are (Pierrehumbert, 1990; Fleming, 2000). Coleman has shown that in Tashlhiyt Berber (Morocco), where many words appear to have no vowels, careful phonetic analysis dramatically simplifies the phonological analysis of syllable structure (Coleman, 2001).

Important works on the syllable, stress, intonation and tone include the following: (Pike and Pike, 1947; Liberman and Prince, 1977; Burzio, 1994; Hayes, 1994; Blevins, 1995; Ladd, 1996; Hirst and Di Cristo, 1998; Hyman and Kisseberth, 1998; van der Hulst and Ritter, 1999). Studies of partial specification and redundancy include: (Archangeli, 1988; Broe, 1993; Archangeli and Pulleyblank, 1994).

Attribute-value and directed graph models for phonological representations and constraints are described in the following papers and monographs: (Bird and Klein, 1994; Bird, 1995; Coleman, 1998; Scobbie, 1998; Bird and Liberman, 2001; Cassidy and Harrington, 2001).

The last decade has seen two major developments in phonology, both falling outside the scope of this limited chapter. On the theoretical side, Alan Prince, Paul Smolensky, John McCarthy and many others have developed a model of constraint interaction called *Optimality Theory* (OT) (Archangeli and Langendoen, 1997; Kager, 1999; Tesar and Smolensky, 2000). The Rutgers Optimality Archive houses an extensive collection of OT papers [`http://ruccs.rutgers.edu/roa.html`]. On the computational side, the Association for Computational Linguistics (ACL) has a special interest group in computational phonology (SIGPHON) with a homepage at `http://www.cogsci.ed.ac.uk/sigphon/`. The organization has held five meetings to date, with proceedings published by the ACL and many papers available online from the SIGPHON site: (Bird, 1994b; Sproat, 1996; Coleman, 1997; Ellison, 1998; Eisner et al., 2000). Another collection of papers was published as a special issue of the journal *Computational Linguistics* in 1994 (Bird, 1994a). Several PhD theses on computational phonology have appeared: (Bird, 1995; Kornai, 1995; Tesar, 1995; Carson-Berndsen, 1997; Walther, 1997; Boersma, 1998; Wareham, 1999; Kiraz, 2000). Key contributions to computational OT include the proceedings of the fourth and fifth SIGPHON meetings, and (Ellison, 1994; Tesar, 1995; Eisner, 1997; Karttunen, 1998).

The sources of data published in this chapter are as follows: Russian (Kenstowicz and Kisseberth, 1979); Chakosi (Ghana: Language Data Series, ms); Ulwa (Sproat, 1992:49).

6 Acknowledgements

I am grateful to D. Robert Ladd and Eugene Buckley for comments on an earlier version of this chapter, and to James Roberts for furnishing me with the Chakosi data.

References

Stephen R. Anderson. *Phonology in the Twentieth Century: Theories of Rules and Theories of Representations*. The University of Chicago Press, 1985.

Diana Archangeli. Aspects of underspecification theory. *Phonology*, 5:183–207, 1988.

Diana Archangeli and D. Terence Langendoen, editors. *Optimality Theory: An Overview*. Oxford: Blackwell, 1997.

Diana Archangeli and Douglas Pulleyblank. *Grounded Phonology*. MIT Press, 1994.

Steven Bird, editor. *Computational Linguistics: Special Issue on Computational Phonology*, volume 20(3). MIT Press, 1994a.

Steven Bird, editor. *Proceedings of the First Meeting of the ACL Special Interest Group in Computational Phonology*. Somerset, NJ: Association for Computational Linguistics, 1994b.

Steven Bird. *Computational Phonology: A Constraint-Based Approach*. Studies in Natural Language Processing. Cambridge University Press, 1995.

Steven Bird. Dschang syllable structure. In Harry van der Hulst and Nancy Ritter, editors, *The Syllable: Views and Facts*, Studies in Generative Grammar, pages 447–476. Berlin: Mouton de Gruyter, 1999.

Steven Bird and T. Mark Ellison. One level phonology: autosegmental representations and rules as finite automata. *Computational Linguistics*, 20:55–90, 1994.

Steven Bird and Ewan Klein. Phonological analysis in typed feature systems. *Computational Linguistics*, 20:455–91, 1994.

Steven Bird and Mark Liberman. A formal framework for linguistic annotation. *Speech Communication*, 33:23–60, 2001.

Juliette Blevins. The syllable in phonological theory. In John A. Goldsmith, editor, *The Handbook of Phonological Theory*, pages 206–44. Cambridge, MA: Blackwell, 1995.

Bernard Bloch. A set of postulates for phonemic analysis. *Language*, 24:3–46, 1948.

Paul Boersma. *Functional Phonology: Formalizing the Interactions Between Articulatory and Perceptual Drives*. PhD thesis, University of Amsterdam, 1998.

Geert Booij. *The Phonology of Dutch*. The Phonology of the World's Languages. Oxford: Clarendon Press, 1995.

Michael Broe. *Specification Theory: the Treatment of Redundancy in Generative Phonology*. PhD thesis, University of Edinburgh, 1993.

Michael Broe and Janet Pierrehumbert, editors. *Papers in Laboratory Phonology V: Language Acquisition and the Lexicon*. Cambridge University Press, 2000.

Luigi Burzio. *Principles of English Stress*. Cambridge University Press, 1994.

Julie Carson-Berndsen. *Time Map Phonology: Finite State Models and Event Logics in Speech Recognition*, volume 5 of *Text, Speech and Language Technology*. Kluwer, 1997.

Steve Cassidy and Jonathan Harrington. Multi-level annotation of speech: An overview of the emu speech database management system. *Speech Communication*, 33:61–77, 2001.

J. C. Catford. *Practical Introduction to Phonetics*. Oxford: Clarendon Press, 1988.

N. Chomsky and M. Halle. *The Sound Pattern of English*. New York: Harper and Row, 1968.

John Clark and Colin Yallop. *An Introduction to Phonetics and Phonology*. Oxford: Blackwell, 1995.

John Coleman, editor. *Proceedings of the Third Meeting of the ACL Special Interest Group in Computational Phonology*. Somerset, NJ: Association for Computational Linguistics, 1997.

John Coleman. The phonetics and phonology of tashlhiyt berber syllabic consonants. *Transactions of the Philological Society*, 99:29–64, 2001. `http://www.phon.ox.ac.uk/~jcoleman/TPS.html`.

John S. Coleman. *Phonological Representations — their Names, Forms and Powers*. Cambridge Studies in Linguistics. Cambridge University Press, 1998.

Bruce Connell and Amalia Arvaniti. *Papers in Laboratory Phonology IV: Phonology and Phonetic Evidence*. Cambridge University Press, 1995.

Gerard J. Docherty and D. Robert Ladd, editors. *Papers in Laboratory Phonology II: Gesture, Segment, Prosody*. Cambridge University Press, 1992.

Jason Eisner. Efficient generation in primitive optimality theory. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, pages 313–320. Somerset, NJ: Association for Computational Linguistics, 1997.

Jason Eisner, Lauri Karttunen, and Alain Thériault, editors. *Proceedings of the Fifth Meeting of the ACL Special Interest Group in Computational Phonology*. Somerset, NJ: Association for Computational Linguistics, 2000.

Mark Ellison, editor. *Proceedings of the Fourth Meeting of the ACL Special Interest Group in Computational Phonology*. Somerset, NJ: Association for Computational Linguistics, 1998.

T. Mark Ellison. Phonological derivation in optimality theory. In *Proceedings of the Fifteenth International Conference on Computational Linguistics*, pages 1007–13. International Committee on Computational Linguistics, 1994.

John Rupert Firth. Sounds and prosodies. In *Papers in Linguistics 1934–1951*, pages 121–138. London: Oxford: Clarendon Press (1957), 1948. Originally published in Transactions of the Philological Society, 1948:127–52.

Edward Fleming. Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology*, 17, 2000.

Ram Frost and Leonard Katz, editors. *Orthography, Phonology, Morphology and Meaning*, volume 94 of *Advances in Psychology*. Amsterdam: North-Holland, 1992.

John A. Goldsmith, editor. *The Handbook of Phonological Theory*. Cambridge, MA: Blackwell, 1995.

John A. Goldsmith, editor. *Phonological Theory: The Essential Readings*. Cambridge, MA: Blackwell, 1999.

Carlos Gussenhoven and Haike Jacobs. *Understanding Phonology*. Edward Arnold, 1998.

Michael Hammond. *The Phonology of English: A Prosodic Optimality-Theoretic Approach*. Oxford: Clarendon Press, 1999.

Jürgen Handke. *The Mouton Interactive Introduction to Phonetics and Phonology*. Berlin: Mouton de Gruyter, 2001.

Jonathan Harrington and Steve Cassidy. *Techniques in Speech Acoustics*. Kluwer, 2000.

Bruce Hayes. *Metrical Stress Theory: Principles and Case Studies*. University of Chicago Press, 1994.

Daniel Hirst and Albert Di Cristo, editors. *Intonation Systems: A Survey of Twenty Languages*. Cambridge University Press, 1998.

Charles F. Hockett. *A Manual of Phonology*. Baltimore: Waverly Press, 1955.

Larry M. Hyman and Charles Kisseberth, editors. *Theoretical Aspects of Bantu Tone*. CSLI Publications / Cambridge University Press, 1998.

International Phonetic Association. *Handbook of the International Phonetic Association: A Guide to the Use of the International Phonetic Alphabet*. Cambridge University Press, 1999.

Martin Joos, editor. *Readings in Linguistics I: The Development of Descriptive Linguistics in America, 1925–56*. The University of Chicago Press, 1957.

Daniel Jurafsky and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall, 2000.

René Kager. *Optimality Theory*. Cambridge University Press, 1999.

Lauri Karttunen. The proper treatment of optimality in computational phonology. xxx.lanl.gov/abs/cmp-lg/9804002, 1998.

Francis Katamba. *An Introduction to Phonology*. Addison Wesley, 1989.

Martin Kay. Nonconcatenative finite-state morphology. In *Proceedings of the Third Meeting of the European Chapter of the Association for Computational Linguistics*, pages 2–10. Somerset, NJ: Association for Computational Linguistics, 1987.

Patricia A. Keating. *Papers in Laboratory Phonology III: Phonological Structure and Phonetic Form*. Cambridge University Press, 1994.

Michael Kenstowicz. *Phonology in Generative Grammar*. Blackwell, 1994.

Michael Kenstowicz and Charles Kisseberth. *Generative Phonology: Description and Theory*. Academic Press, 1979.

John Kingston and Mary E. Beckman, editors. *Papers in Laboratory Phonology I: Between the Grammar and the Physics of Speech*. Cambridge University Press, 1991.

George Kiraz. *Computational Approach to Non-linear Morphology*. Studies in Natural Language Processing. Cambridge University Press, 2000.

András Kornai. *Formal Phonology*. New York: Garland Publishing, 1995.

Gjert Kristoffersen. *The Phonology of Norwegian*. The Phonology of the World's Languages. Oxford: Clarendon Press, 1996.

William Labov, Sharon Ash, and Charles Boberg. *Atlas of North American English*. Berlin: Mouton de Gruyter, 2001.

D. Robert Ladd. *Intonational Phonology*. Cambridge University Press, 1996.

Peter Ladefoged. *Vowels and Consonants: An Introduction to the Sounds of Languages*. Cambridge, MA: Blackwell, 2000.

Peter Ladefoged and Ian Maddieson. *The Sounds of the World's Languages*. Cambridge, MA: Blackwell, 1996.

John Laver. *Principles of Phonetics*. Cambridge Textbooks in Linguistics. Cambridge University Press, 1994.

Mark Y. Liberman and Alan S. Prince. On stress and linguistic rhythm. *Linguistic Inquiry*, 8:249–336, 1977.

Helena Mateus and Ernesto d'Andrade. *The Phonology of Portuguese*. The Phonology of the World's Languages. Oxford: Clarendon Press, 2000.

David Odden. *The Phonology and Morphology of Kimatuumbi*. The Phonology of the World's Languages. Oxford: Clarendon Press, 1996.

Janet Pierrehumbert. Phonological and phonetic representation. *Journal of Phonetics*, 18:375–94, 1990.

Kenneth L. Pike. *Phonemics: A Technique for Reducing Language to Writing*. Ann Arbor: University of Michigan Press, 1947.

Kenneth L. Pike and E. V. Pike. Immediate constituents of Mazateco syllables. *International Journal of American Linguistics*, 13:78–91, 1947.

Iggy Roca, Wyn Johnson, and Aggy Roca. *A Course in Phonology*. Cambridge, MA: Blackwell, 1999.

Jerzy Rubach. *The Lexical Phonology of Slovak*. Oxford: Clarendon Press, 1993.

James Scobbie. *Attribute-Value Phonology*. New York: Garland Publishing, 1998.

Péter Siptár and Miklós Törkenczy. *The Phonology of Hungarian*. The Phonology of the World's Languages. Oxford: Clarendon Press, 2000.

Richard Sproat. *Morphology and Computation*. Natural Language Processing. Cambridge, MA: MIT Press, 1992.

Richard Sproat, editor. *Computational Phonology in Speech Technology: Proceedings of the Second Meeting of the ACL Special Interest Group in Computational Phonology*. Somerset, NJ: Association for Computational Linguistics, 1996.

Kenneth N. Stevens. *Acoustic Phonetics*. MIT Press, 1999.

Bruce Tesar. *Computational Optimality Theory*. PhD thesis, Rutgers University, 1995.

Bruce Tesar and Paul Smolensky. *Learnability in Optimality Theory*. MIT Press, 2000.

Harry van der Hulst and Nancy Ritter, editors. *The Syllable: Views and Facts*. Studies in Generative Grammar. Berlin: Mouton de Gruyter, 1999.

Bert Vaux. *The Phonology of Armenian*. The Phonology of the World's Languages. Oxford: Clarendon Press, 1998.

Markus Walther. *Declarative prosodic morphology - constraint-based analyses and computational models of Finnish and Tigrinya*. PhD thesis, Heinrich-Heine-Universität, Düsseldorf, 1997. thesis in German.

Todd Wareham. *Systematic Parameterized Complexity Analysis in Computational Phonology*. PhD thesis, University of Victoria, 1999.

Richard Wiese. *The Phonology of German*. The Phonology of the World's Languages. Oxford: Clarendon Press, 1996.