Features and parameters for different purposes

Peter Ladefoged

Linguistics Department, UCLA, Los Angeles, CA 90095-1543

Abstract

The phonetic description of a language must be related to the phonology. A computerized description of a language can have a very faithful phonetic component, but its phonetic structures are not appropriate for a phonological description. In current systems of linguistic analysis there are three aspects of phonology: (1) the representation of the lexical contrasts in a language; (2) the specification of the constraints on the sounds in lexical items; and (3) the description of phonological patterns of sounds as evident in the relations between the underlying lexical items and the observable phonetic output. There is a conflict between the phonetic component required for the first of these goals and that required for the other two. Characterizing the sounds of languages can be done most efficiently by using a large number of features, all defined in articulatory terms. This will result in having more features than are necessary to characterize phonological patterns efficiently. In addition, some phonological patterns depend on auditory characteristics which will require auditorily defined features. Yet other patterns are observable in a language considered as a social institution rather than a mental concept.

There are many ways in which one can make a description of the sounds of a language, and linguists often forget about the most obvious one. Figure 1 is an example of part of a description of the sounds of English, namely their waveforms.

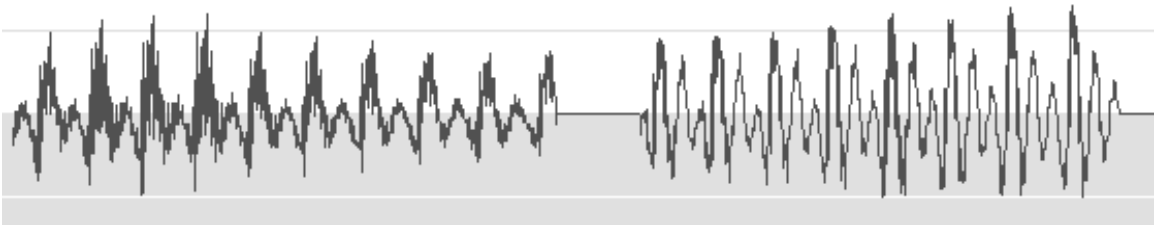

Figure 1. The waveforms of the English vowels /i/ and /u/.

This is a very complete and accurate description, albeit a rather lengthy one. It would take several pages to show the waveforms necessary to describe all the sounds of English. But providing a set of sound waves is certainly one way of formalizing the phonetic component of a language.

Most people also overlook the next most obvious form of description, an x-ray or MRI movie, comprised of frames as in Figure 2. This is not quite as good a description as a waveform— one cannot reconstruct the complete sounds just from these MRI data. But it is not hard to envisage ways in which this description could be elaborated by the addition of physiological data so that it would form a reasonably complete description of the sounds of English.

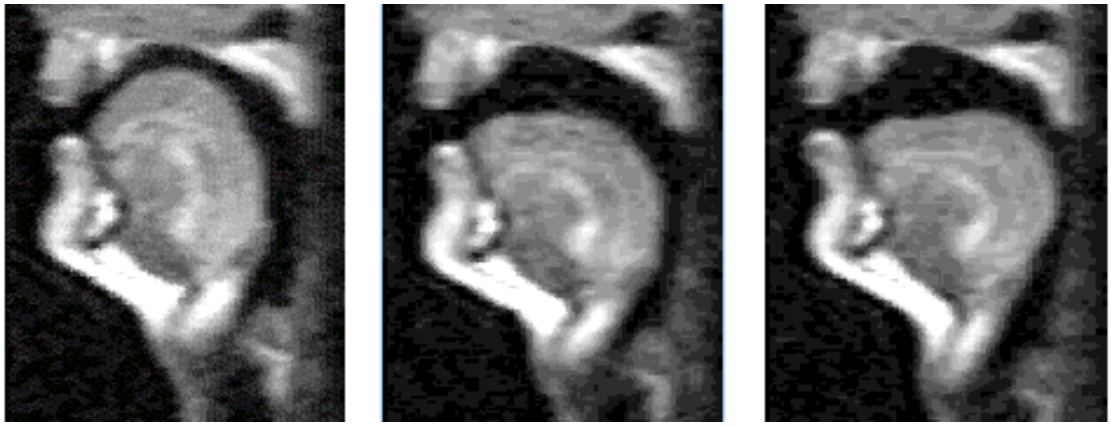

Figure 2. Three frames from an MRI movie of the vowels of French.

These rather unconventional forms of description might well form the phonetic component of a certain kind of grammar, one describing how children acquire language. A set of waveforms is one of the principal sources that a child has when learning to talk. The other principal component available to a child is the second set of data mentioned above, knowledge of the physiological mechanisms involved. Furthermore, if we want to write a grammar that provides the basis for teaching people the language and how to pronounce it, these forms of description might well be part of it. Opinions differ on how best to teach a language, and obviously it depends on who is doing the teaching, who is doing the learning, and why they want to learn it. But many authorities would say that often the best way is to listen to lots of the language, hear carefully constructed phrases and learn when and where to use them. The phonetic component of a grammar of this kind could well be a set of waveforms as illustrated above.

A conventional grammar of a language must account for everything between the thoughts behind an utterance — the semantics — and the corresponding sounds — the phonetics. In this paper, when considering such a grammar we will assume that enough is known about the semantic component for it to start generating a sentence from possible thoughts, and that the syntactic and phonological components are well formalized, and can turn the semantic component into a phonological output. Our interest is in the appropriate phonetic component.

At the moment the ways in which we can model a grammar on a computer involve only limited thoughts and sentences. Those working in semantics and syntax still have a lot of work to do before all the sentences of a language could be generated by a computer. But given a prompt, such as the question ‘How do I measure the lengths of the vowels in the word *aware*?’ a computer can be programmed to take this question, work out what it means, and generate an appropriate response such as ‘Please ask your TA to help you’. The phonetic component of a language description involves turning this answer into sounds, which it can do quite easily, as

anyone can test by typing these or other phrases into Rhetorical Systems demonstration at: http://www.rhetorical.com/cgi-bin/demo.cgi.

When people hear computer generated speech, they generally think that the phonetic component is fairly good. The intonation could often be improved, but the general speech quality is fine. Of course, the semantic component of a computer model of a grammar could be better programmed so that it was a little more helpful, but that is not our concern here. Speech synthesis shows that we know how to make phonetic descriptions so that we can go from the output generated by the semantic and syntactic components of a language generator to reasonably high quality speech. This is a grammar in accordance with the definition given at the beginning of this paper, an account of a language that goes from the thoughts behind an utterance to its realization as sounds that we can hear.

What is the phonetic component in this grammar doing in most of the current computer speech synthesis systems? The answer is that it does not use phonemes or features or any of the usual linguistic units to generate sound waves. The phonetic component of a computer system may use an orthographic text, but it does not really require one. Chinese characters would do just as well, as long as there was a character for each possible English syllable. The way most speech synthesis systems work is to search through a vast store of several hours of recorded speech and find the longest possible sequences that match the desired output. In practice this is done with reference to something like a segmental transcription of words, but there is nothing in the theory of concatenative speech synthesis, as this process is called, that requires traditional phonetic segments. In fact it virtually never uses them. The smallest lengths of sounds that are ever pulled out of the store are diphones — the last half of one segment and the first half of another — and even these are used only very occasionally, when there is no matching word or syllable in store. Concatenative synthesis nearly always manages to join whole syllables or even whole long words together. Listeners to high quality synthetic speech agree that the phonetic component of a language generator works quite well by simply finding appropriate bits of stored sounds and then joining them together. Given simple questions like “Where and when is the Linguistics 523 class?” it is quite within the capabilities of the semantic and syntactic components to give a correct analysis that the phonetic component can use to produce the answer: “The seminar is on Wednesdays, from 12 to 3. It is in Guildford, room 112.” A good synthesizer even gets the proper name correctly, and all without using conventional linguistic units.

Automatic speech recognition systems also work reasonably well without using conventional linguistic units. None of them is as advanced as the speech synthesis we have been considering. Speech systems can generate, with minor faults, almost any sentence from a slice of Shakespeare to a teenage conversation. Speech recognition systems are limited in the topics they can interpret at any given moment, but they can handle quite long sentences if the topic is known. They do this with very little reference to linguistic units such as phonemes. Instead they search large stores corresponding to sound waves, seeking the most probable matches that lead to possible words.

When workable speech synthesis systems and speech recognition systems use grammars that pay so little attention to the customary units found in linguistic analyses perhaps it is time to ask if there are fundamental problems with our linguistic analyses.

Linguistic phonetics

In current systems of linguistic analysis the phonetic component of a language is considered to be part of its phonology. A traditional view is that there are three aspects of phonology (Goldsmith 1998): (1) the representation of the lexical contrasts in a language; (2) the specification of the constraints on the sounds in lexical items in a language (or, as some would put it, the definition of a well formed syllable); and (3) the formal description of the relations between the underlying lexical items and the observable phonetic output (or, putting it another way, the explication of the sound patterns of a language). All these aspects of phonology require a well defined phonetic component. In this part of this paper we will see that there is a conflict between the phonetic component required for the first of these goals, representing lexical contrasts, and that required for the other two, constraining or defining words and syllables, and explicating the sound patterns. Any theory of grammar that neglects this conflict is flawed.

The usual way for phonologists to describe the phonetic component is in terms of phonological features such as those originally proposed by Chomsky and Halle (1968) and now somewhat modified. There is no fully agreed set, but the core features probably include: LABIAL, CORONAL, DORSAL, RADICAL; Voice, Spread Glottis, Constricted Glottis; Syllabic, Sonorant, Consonantal; Distributed, Anterior, ATR, Strident, Lateral, Nasal, Continuant; High, Low, Back, Round, Tense. The first four of these features are in capital letters, as they can have only one possible value. Most of the remaining features are binary — they must have one or other of two possible values. It is sometimes held, however, that the features can be arranged in a hierarchy (Clements and Hume, 1995). From this point of view, there are some dependent features that should also be considered as potentially having no value. For example, Distributed cannot have a value for sounds that are not Coronal.

In theory, each of these features is definable in terms of certain phonetic properties. These properties can be considered to be either part of a Universal Grammar that forms part of a child's innate endowment, or, in an alternative view that will be promoted here, they can be considered to be simply phonetic properties of the language being described. If features are part of a Universal Grammar they may be said to be universally defined by particular properties, but if features are considered one language at a time, they can be defined locally with reference to the properties of that language. A phonology in which the features are locally defined just for a particular language has both strengths and weaknesses. When the phonetic details are glossed over, it is possible to group together somewhat disparate sounds that act together in a phonological process. For example, in a recent phonological description of this sort Local and Lodge (2004) give a good account of Kalenjin harmony patterns in both consonants and vowels, if it is assumed that the ATR feature has variable consonantal exponents. The weakness of a feature system that does not require fixed universal definitions of each feature is that the features do not form part of the innate endowment of all children. In addition, the feature system may not fulfill the first aim of phonology listed above, the ability to distinguish lexical contrasts in different languages. It allows for good statements of phonological patterns, fulfilling the second and third aims, at the expense of not involving a Universal Grammar.

Some of the data that have to be considered

The first problem that arises in setting up an extrinsically defined feature system that could be part of a Universal Grammar is that we have to decide when a sound in one language is different from a similar sound in another language. Obviously, two sounds in different languages

should be considered as distinct if we can point to a third language in which the same two sounds distinguish words. But if there is no such language we can only estimate whether the difference between the two sounds in different languages is great enough so that there might be (or might have been or will some time be) a language in which they are used to distinguish words.

Making an estimate of this kind is not at all easy. Consider some of the known sounds that are used in languages. Who would have thought that people could distinguish bilabial [ɸ] and labiodental [f] as they do in Ewe /éɸá/ 'he polished' vs. /éfá/ 'he was cold'? Or that there could be a difference between a voiceless nasalized alveolar click [n̊ǃ] and voiceless aspirated alveolar click [kʰǃ] as there is in Nama /n̊ǃóas/'narrating' vs. /kʰǃóas/ 'belt'?

The makers of the early IPA charts were faced with this problem. They noted the difference between [ɸ] and [f] (using the symbols **ꜰ** and **f**) although the phoneticians at the time did not know of any language that used either that distinction or the voiced counterparts, which they symbolized as **ʋ** vs. **v**. The notes accompanying the first IPA chart describe **ꜰ** (along with **ʋ**) as 'une fricative bilabiale simple'. It does not illustrate **ꜰ** at all, but says **ʋ** occurs "dans le Flamand *wrocht*, dans l'Allemand du Sud *wesen* and elsewhere" (Association Phonétique Internationale 1900:8). The 1904 English version describes **ꜰ** as "the sound we produce when blowing out a candle" and **ʋ** using similar dialect examples as in the French version (International Phonetic Association 1904:8). It seems that right from the start the IPA tradition was one of providing symbols and categories not only for sounds that had actually been observed to contrast within a language but also for sounds that might be different.

As an example of the problems involved in setting up a classificatory system using features that have absolute phonetic definitions we will consider in some detail the varying sibilants that occur in different languages. Lexical contrasts involving sibilants occur in Polish, Chinese, Ubykh (North Caucasian), Toda (Dravidian) and English. All of these languages except English have three or four contrasting sibilant fricatives, sounds that would all be called [Coronal] and [+ strident] in most feature systems. The articulatory and acoustic characteristics of these fricatives are discussed at length by Ladefoged and Maddieson (1996), who provide palatographic, x-ray and acoustic data to support their conclusions. Their findings indicate that there are 8 possible coronal strident gestures as shown in Table 1, with two of these gestures (those in rows 2 and 6) allowing alternate possibilities (and thus giving us 10 rows). We will consider the alternate possibilities first, so that we can concentrate on the eight clearly contrasting gestures.

Table 1. Types of sibilants (based on Ladefoged and Maddieson, 164:1996).

		ARTICULATORY GESTURES	EXEMPLIFYING LANGUAGES
1	s̪	dental	Polish
2	s	apical or laminal alveolar	Chinese, English, Ubykh
3	s̻	laminal alveolar	Toda
4	ṣ	laminal flat post-alveolar	Chinese, Polish, Ubykh
5	s̺	apical post-alveolar	Toda
6	ʃ	apical or laminal domed post-alveolar	English
7	ʃ̻	laminal domed post-alveolar (palatoalveolar)	Toda
8	ɕ	laminal palatalized post-alveolar (alveolopalatal)	Chinese, Polish, Ubykh
9	ŝ	laminal closed post-alveolar ("hissing-hushing")	Ubykh
10	ʂ	sub-apical post-alveolar (sub-apical retroflex)	Toda

Rows 2 and 6 in Table 1 list languages in which there are articulations that some speakers make using the top of the tongue, making them apical, and other speakers use a laminal articulation. In these languages the apical/laminal distinction, usually represented by the feature Distributed, is neutralized. Thus, in English, speakers may use either an apical or a laminal articulation when producing [s] and the same is true for their production of [ʃ]. In other languages, such as Toda, the apical laminal distinction plays a key role in distinguishing sibilants. Eliminating the rows in which the apical/laminal distinction is neutralized, we still have 8 articulatory gestures that we have to distinguish. Words illustrating these contrasts are given in Table 2.

Table 2. Words illustrating Polish, Chinese, Ubykh and Toda sibilant sounds. The columns are numbered as in Table 1, with columns 2 and 6 being omitted.

	1	3	4	5	7	8	9	10
Polish	**s̪ali**		**ṣali**			**ɕali**		
	sali		*szali*			*siali*		
	'room' (gen.)		'scale' (gen.)			'sown'		
Chinese		**saˀ**	**ṣaˀ**			**ɕaˀ**		
		'three (items)'	'sand'			'blind'		
Ubykh		**saːba**	**ṣaˀ**			**ɕaɕa**	**ŝa**	
		'why'	'head'			'mother-in-law'	'three'	
Toda	**koːs̻**			**poːs̺**	**poːṣ**			**poː ʃ̻**
	'money'			'milk'	(clan name)			'language'

There is no known language that contrasts the 8 distinct gestures involved in producing the sibilants shown in table 2. If we are considering phonological distinctions, do we need a classificatory system that formalizes the possibility of all these articulations? The claim here is that we do need such a system. Table 3 in part validates this claim by showing some of the contrasts that have been observed using the data in Tables 1 and 2, with the addition of data from English

It is difficult to know precisely how many of the gaps in Table 3 are due to our not having found a language using a given contrast as opposed to being due to over-differentiation of

categories. Consider, for example, dental vs. alveolar sibilants. There is no doubt that several languages have dental sibilants. In addition to Toda, exemplified above, many languages of California, such as Karok (Bright 1978), have dental [s̪]. But in these languages the nearest other sibilant is post-alveolar [s̺] rather than alveolar [s]. It may well be that dental and alveolar sibilants are not sufficiently different to provide reliable linguistic contrasts. But from a classificatory point of view we are not extending the system by allowing for the possibility of contrasting dental and alveolar sibilants, as we already need the dental–alveolar contrast for other sounds, such as nasals in Malayalam (Ladefoged and Maddieson 1996).

Table 3. A matrix showing observed contrasts among sibilant fricatives in different languages.

	Alveolar	Laminal flat Post-alveolar	Apical post - Alveolar	Laminal domed Post-alveolar	laminal palatalized post-alveolar	Closed post alveolar	Sub-apical retroflex
	s	s̺	s̠	ʃ	ɕ	ŝ	ʂ
Dental s̪	xxx	Polish	Toda	Toda	Polish		Toda
Alveolar s		Chinese, Ubykh		English	Chinese, Ubykh	Ubykh	
Laminal Flat Post-Alveolar s̺			xxx		Chinese, Polish, Ubykh	Ubykh	
Apical Post -Alveolar s̠				Toda			Toda
Laminal Domed Post-Alveolar ʃ					xxx		Toda
Alveolo palatal ɕ						Ubykh	
Closed post alveolar ŝ							xxx

In the second row in Table 3, the contrast between alveolar [s] and post-alveolar [s̠] cannot be demonstrated from the sibilant data at hand. But these appear to be distinct gestures and could well contrast. The same is even more true for the other gap in that row, between alveolar [s] and sub-apical alveolar [ʂ]. In fact we can safely assert that all the gaps in the sub-apical column are due to the rarity of these sounds, rather than to their being a non-contrastive variant of some other gesture. The same applies to the gaps in the closed post-alveolar column. These are unusual sounds, but clearly sounds with distinct gestures.

This leaves us with two suspicious gaps to consider. There is no contrast that we know of between laminal flat post-alveolar [s̺] and apical post-alveolar [s̠], nor is there one between laminal domed post-alveolar (palatoalveolar) [ʃ] and alveolo-palatal sibilants [ɕ]. The first of these gaps does not cause any extension of the classificatory system as we have to have a feature characterizing the apical — laminal contrast for other sounds. The feature Distributed is usually assigned to this task. However, within a hierarchical feature system, we should note that it seems likely that Distributed should have no value in this context.

The lack of a contrast between laminal domed post-alveolar (palatoalveolar) [ʃ] and alveolo-palatal [ɕ]may be caused by the auditory similarity of these sounds. They may be not sufficiently distinct to be able to sustain a reliable linguistic contrast. Nevertheless speakers of languages such as Polish and Chinese are keenly aware of the difference between their [ɕ] and the similar but distinguishable [ʃ] sound that occurs in English. They protest loudly when phoneticians confuse the two. We will take it that this is a case where we have not yet found a language that uses this contrast rather than ruling out the possibility of a contrast on the grounds that the two sounds are too similar. However, this may be a wrong judgment. It must be admitted that these sounds are considered to be distinct because, simply as a matter of opinion, it seems likely that they *could* be used as distinctive categories, and with no real scientific basis for this conclusion.

How can the 8 sibilant gestures that we have established as being distinct be described in terms of phonological features? They are all CORONAL, and [+ strident]. They are also all [– voice], [+ spread glottis], [– constricted glottis], [– syllabic], [– sonorant, [+ consonantal], [– ATR], [– lateral], [– nasal], [+ continuant], and [– round]. Considering the features listed above, this leaves us with High, Low, Back, Distributed, Anterior, and Tense to account for these 8 distinctions. Low and Back as usually defined for vowels and seem non-definable for any of the sibilant gestures, as does Tense. This leaves us with Anterior, Distributed and High to distinguish these 8 sounds.

Table 4. The classification of sibilants in terms of the features Anterior, Distributed and High.

	Dental	Alveolar	Laminal flat post-alveolar	Apical post - alveolar	Laminal domed post-alveolar	Laminal palatalized post-alveolar	Closed post alveolar	Sub-apical retroflex
	s̪	s	s̡	s̲	ʃ	ɕ	ŝ	ʂ
Anterior	+	+	–	–	–	–	–	–
Distributed	(+)	(–)	+	–	+	+	+	–
High	–	–	–	–	–	+	–	–
Retroflex	–	–	+	–	–	–	–	+
Closed	–	–	–	–	–	–	+	–

Table 4 shows how these three features can be used, together with two more that prove necessay. The usual definition of Anterior is that [+ anterior] sounds are articulated forward of the alveolar ridge, i.e. equivalent to what we have been calling dental or alveolar, [– anterior] sounds are post-alveolar. We have already noted that we do not know of a language that distinguishes dental and alveolar sibilants, but if there were one it would seem likely that one would be [+ distributed] and the other [– distributed]. The remaining 6 sounds cannot be distinguished by just the two features Distributed and High. If Distributed is defined such that [+ distributed] is equivalent to what we have been calling laminal, and [– distributed] equivalent to apical we can fill in the third row in Table 4 as shown. This leaves us with having to distinguish [s̡, ʃ, ɕ, ŝ], all of which are [+ distributed], and [s̲] and [ʂ], which are both [– distributed]. The feature High can be used to distinguish [ɕ] from [s̡, ʃ, ŝ] by calling [ɕ] [+ high], but this feature does not really distinguish any of the others. If we are to distinguish all these sounds in terms of a Universal Grammar we would have to add two new features as shown in the last two rows of Table 4. Retroflex will separate [s̡] from [ʃ, ŝ], and Closed will separate [ŝ] from [ʃ].

Similar arguments can easily be made showing that more features are needed for classifying vowels. Thinking just of how many vowel heights there are, Danish has four front vowels, each of which can be long or short. These vowels cannot be said to differ in terms of the features Tense or ATR as usually defined. As well as the features High and Low, we will have to add some other feature, such as Mid, which we can define as having first formant frequencies in the middle of the range. We can then say that the Danish vowels are high, mid high, mid low, low, much as the IPA does. This would also enable us to characterize Germanic dialects, such as the Dutch dialect of Weert (see Table 5) that have 5 vowel heights among front vowels. We can describe these vowels in terms of a single parameter of vowel height as indicated by the terms in the first column of the table, and in terms of binary features as shown in the last three columns.

Table 5. Words illustrating the front unrounded vowels of the Dutch dialect of Weert (data from Heijmans and Gussenhoven,1998).

	Long	Short	High	Mid	Low
High	**β̟iːt** 'far'	**ʁit** 'Mary'	+	–	–
Mid-high	**ʁeːt** 'reed'	**hɪtst** 'heat'	+	+	–
Mid	**'blɛːcə** 'leaf' (dim.)	**'zɛgə** 'to say'	–	+	–
Mid-low	**tæːnt** 'tent'	**slæt** 'dishcloth'	–	+	+
Low	**naːt** 'wet'		–	–	+

It is easy to carry this line of argument further and show that we need still more features for a truly Universal grammar. For example, we need features for the lexical contrasts formed by the 83 clicks that occur in !Xóõ, and for the phonation types in the Austro-Asiatic languages of South-East Asia, among many other contrasting sounds.

We should also note that any system based on our current linguistic knowledge must be incomplete for two reasons. Firstly, as we have admitted, some distinctions permitted within the system are simply estimates of what distinctions are possible within a language. Some future language may arise that proves these estimates wrong. We may find, for example, a language that distinguished not only the tense and modal phonations of Bruu, but also murmured and creaky phonations. Secondly, we cannot account for what is not systematic. Despite de Saussure's claim that 'la langue est une systéme ou toute sa tient" (Sausure 1968), it is not true that each language forms a system in which everything holds together. There are always odds and ends that may, or may not, be considered part of the language. Ladefoged and Everett (1996) point to a number of cases in which a language has an usual sound that occurs in a dozen or fewer regular lexical items, such as the alveolar released bilabial trill that occurs in Oro Win in words such as [t͡ʁum]'a small boy'.

Other properties of feature systems

So far the burden of the argument has been that the set of features necessary for describing lexical contrasts in a Universal Grammar is large, cumbersome and different from the set of features needed for describing phonological patterns. Describing lexical contrasts necessarily requires different features than those required for describing phonological patterns. The reverse is also true. Describing phonological patterns requires different features from those needed for describing the lexicon. We can make complete descriptions of the lexical contrasts that occur in the world's languages by referring to properties of the ways in which the sounds are made. We do not have to refer to ways in which they are heard. Insofar as features are simply within the minds of speakers they can be said to have both articulatory and auditory correlates. But when we are accounting for patterns that occur in a language we need some features that have articulatory correlates and others that have auditory correlates.

Many years ago Martinet (1955) described the two principal causes of sound changes and the resulting patterns of sounds: articulatory ease (which produces, for example, assimilation in plural forms such as *cats* and *ducks* with [s], as opposed to *dogs* and *lions* with [z]), and auditory distinctiveness (which produces plural forms such as *horses* and *fishes*, in which the two sibilants are kept separate by an epenthetic vowel). Even in this single, well known, linguistic phenomenon, we need to refer to the assimilation of an articulatory feature, Voice, and the separation of sounds with the feature, Strident, an auditory feature that is plainly distinguished by its acoustic characteristics rather than by its manner of articulation.

Ladefoged (1971, 1992) has described several other features that are characterized by auditory properties rather than by articulations. These include Sonorant, Rhotic, and the features that specify vowels, such as High, Low and Back, which, despite popular belief, do not have well defined articulatory correlates (which is why the correlates of Mid were given in terms of formant frequencies in the preceding section)

The implications of there being more than one origin of phonological patterns have not been apparent in classificatory systems. Neither Jakobsonian distinctive features nor Chomsky and Halle (1968) phonological features took articulatory ease and auditory distinctiveness into account. They had different aims. Jakobson, Fant and Halle (1951) were interested in developing a minimal classificatory system rather than one that helped explain the observed patterns. Chomsky and Halle were interested in explaining observed sound patterns, but they considered their feature set to have both articulatory and acoustic properties that speakers know about. They did not consider sound patterns in terms of two distinct sets of features.

To appreciate the difference between the view presented here and that of Chomsky and Halle we have to consider the nature of language. What is a grammar trying to specify? From the viewpont being presented here it is not just something in a speaker's mind. There are around half million words in the Oxford English Dictionary, and no individual knows all of them. Yet they are all good lexical items that a grammar must contain. This cannot be considered as simply a matter of the competence of an ideal speaker as opposed to the performance of an individual. We cannot claim that an ideal speaker could know all these words. It would not be true. The brain is not structured that way. Nevertheless there are around half a million words in present day English. They are all part of the language considered as a social institution.

There are other properties of the language that are not part of a speaker's competence, but are part of the language as a social institution. There are many observable patterns of English sounds

that a grammar should describe as part of the language. For example there is a constraint in English against having two non-coronal stops, oral or nasal, at the end of a word, as in *bomb* vs, *bombardier*, *iamb* vs. *iambic, paradigm* vs. *paradigmatic* etc. The same provision applies at the beginning of a word, as in *mnemonic* vs. *amnesia, Gnostic* vs. *agnostic*, and even *pterygoid* vs. *helicopter*. The pattern is there, and if we were ever invaded by pterygoid aliens with six wings we would no doubt coin a new word and regard them as hexapters. The existence of a pattern, even one that is productive for literate speakers, is not proof that this is part of a speaker's knowledge of the language, just as the existence of half a million words is not proof that they could all be part of an ideal speaker's competence. But if a pattern is there it should be described in a grammar that regards a language as a social institution

Our current linguistic analyses may be far from reflecting mental processes for many reasons. In writing grammars we have been strongly influenced not so much by observations of what goes on in a person's mind as by observations of the development of the language. Viewed as a social institution, a language changes because of military or economic conquests, and because social groups are always striving for a new way of speaking to mark their cohesiveness. The current form of a language is a reflection of its past history, and any grammar describing the sound patterns that occur must take this into account.

A language is an intricate social institution, like the national economy. Languages and economies are best discussed as self organizing systems subject to the pressures of particular societies. No one would describe the economy in terms of the competence of an ideal user of money. Similarly we have to think of languages as institutions that are molded by many factors. We linguists are a little better off than economists. We have identified some of the more important factors — the influences of other languages, the pressures exerted by laziness or efficiency resulting in articulatory changes, and the desire to maintain clarity and produce a distinct auditory message. To these well known influences on languages we must add a third that we will call organizational economy. This is the counterpart within a self-organizing system of what Hockett (1955) calls pattern congruity, the tendency of a language to fill in the gaps so that if it has a set of voice stops /b, d, g/ it is also likely to have a set of voiceless stops /p, t, k/. Maddieson (1996) refers to similar tendencies as gestural economy, and Clements (2004) provides good arguments for considering it as simply feature economy.

Failure to recognize the nature of language may explain why the features of a Universal Grammar as defined by various authorities are not as useful as might be expected in describing phonological patterens. In a survey of 561 languages Mielke (2004) found that the best known feature sets frequently fail to define appropriate natural classes that operate in phonological rules. It seems that on many occasions languages act in individual ways. The different forces acting on a language are too complex to be encapsulated by the feature sets of a Universal Grammar.

The three constraints on the evolution of the self-organizing sound systems of languages must be reflected in any feature system that is trying to account for phonological patterns. Some of the features must be defined with reference to articulatory properties, and some with reference to auditory properties such as those described by Ladefoged (1971, 1992). A hierarchical arrangement of the features will reflect constraints imposed to achieve organizational economy. But this feature system is redundant with regard to specifying the lexicon. As we have noted, the lexicon can be specified without using auditory features.

Learning and speaking a language

Arguing within the assumptions of the current paradigm, we have shown that language patterns arise through a variety of causes that cannot be explicated by reference to a single set of features. We have also shown that the features necessary to specify all the distinct sounds in the languages of the world would be a large set that is not likely to be a part of a Universal Grammar, and would be difficult for a child to manage. But at this point we should step outside the current paradigm and ask why should we imagine that a child learns a language by reference to an innate set of features? Learning a language involves tuning the phonetic parameters to the right values, which can be done without any thought of features. Indeed, agents governed by a computer program can learn the properties of a given vowel system (de Boer 2000), and systems for learning consonants in terms of articulatory phonology have been described by Goldstein and Fowler (2003). Children almost certainly learn whole syllables or larger utterances without reference to features. Toddlers leaving a room know that something has to be said. Some say ['baɪ'baɪ] *bye bye*, others know that the correct thing to say is the single word ['bɪraɪ'bæ] *be right back*. This is unlikely to be represented as shown in part in Table 6. We certainly cannot demonstrate that a feature specification could be used to control the vocal organs. No one has yet produced a computer model that will generate speech sounds directly from a feature matrix.

Table 6. A partial feature representation of ['bɪraɪ'bæ].

Syllabic	–	+	–	+	–	–	+
Sonorant	–	+	+	+	+	–	+
Consonantal	+	–	–	–	–	+	–
Distributed	–	–	–	–	–	–	–
Anterior	–	–	–	–	–	–	–
ATR	–	–	–	–	–	–	–
Strident	–	–	–	–	–	–	–
Lateral	–	–	–	–	–	–	–
Nasal	–	–	–	–	–	–	–
Continuant	–	+	+	+	+	–	+
High	–	+	–	–	+	–	–
Low	–	–	–	+	–	–	+
Back	–	–	–	–	–	–	–
Round	–	–	–	–	–	–	–
Tense	–	–	–	–	–	–	–
Voice	+	+	+	+	+	+	+
Spread glottis	–	–	–	–	–	–	–
Constricted glottis	–	–	–	–	–	–	–
LABIAL	1	0	0	0	0	1	0
CORONAL	0	0	1	0	0	0	0
DORSAL	0	0	0	0	0	0	0
RADICAL	0	0	0	0	0	0	0

Phonological features are best regarded as artifacts that linguists have devised in order to describe linguistic systems. What linguists are describing is a social institution, which, like the

economy, or the way we dress, is subject to the effects of wars, fashion, laziness and communicative efficiency. Phonological features are great for describing the patterns that occur in a language, but learning a language, and the acts of speaking and listening all involve adjusting articulatory parameters not phonological features. What speakers and listeners do may be better described in terms of articulatory phonology and direct perception as suggested by Goldstein and Fowler (2003), rather than by the features that are needed to describe linguistic patterns. In summary phonologists can make good descriptions of languages using features. Phoneticians can make good descriptions using parameters such as those of articulatory phonology or acoustic phonetics. Both phonologists and phoneticians should also be describing language as a social institution. Language is not only in the mind. There's much more of it outside.

References

Association Phonétique Internationale. 1900. *Exposé des Principes*. Bourg-la-Reine, France.

Bright, W. 1978. Sibilants and naturalness in aboriginal California. *Journal of California Anthropology. Papers in Linguistics*, 1, 39-63.

Chomsky, N. 1964. *Current Issues in Linguistic Theory*. The Hague: Mouton and Co., 1964.

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

Clements, G. N. 2004. Feature economy in sound systems, *Phonology.* 20.3. 287-333.

Clements, G. and E. Hume. 1995. The Internal Organization of Speech Sounds. In: J. Goldsmith (ed.), *The Handbook of Phonological Theory*. Cambridge, Mass.: Blackwell. 245-306.

deBoer, B. 2000. Self-organization in vowel systems. *Journal of Phonetics* 28: 441-465.

Goldsmith, J.A. 1996. *Handbook of Phonological Theory*. Blackwells.

Goldstein, L. & Fowler, C. 2003. Articulatory phonology: a phonology for public language use. In *Phonetics and Phonology in Language Comprehension and Production: Differences and Similarities*, edited by Antje S. Meyer and Niels O. Schiller. Mouton de Gruyter.

Hockett, C. F. 1955. *A manual of phonology*. Baltimore: Waverly Press.

International Phonetic Association. 1904. *Aim and Principles.* Bourg-la-Reine France.

International Phonetic Association 1999. *Handbook of the International Phonetic Association*, Cambridge: Cambridge University Press

Jakobson, R., Fant, G., & Halle, M. 1951. *Preliminaries to speech analysis: The distinctive features and their correlates.* Cambridge, MA: MIT.

Ladefoged, P, 1971. Preliminaries to linguistic phonetics, Chicago: Chicago University Press.

Ladefoged, P. 1992. The many interfaces between phonetics and phonology. In W. U. Dressler,H. C. Luschitzky,O. E. Pfeiffer, & J. R. Rennison (Eds.), *Phonologica 1988* , pp. 165-179. Cambridge: Cambridge University Press.

Ladefoged, P., & Everett, D. 1996. The status of phonetic rarities. *Language*, *72* (4), 794-800.

Ladefoged, P., & Maddieson, I. 1996. *Sounds of the World's Languages*. Oxford: Blackwells.

Local, J. & Lodge 2004. Some auditory and acoustic observations on the phonetics of [ATR] harmony in a speaker of a dialect of Kalenjin, *Jornal of the International Phonetic Association*, 34.1, 1-16.

Maddieson, I. 1995 Gestural economy. In *Proceedings of the 13th International Congress of Phonetic Sciences* Stockholm (Edited K. Elenius and P. Branderud), Volume 4: 574-577.

Martinet, A. 1955. *Economie des changements phonétiques*. Berne: Francke.

Mielke, J. 2004. The emergence of distinctive features. Ph.D. Dissertation, Ohio State University.

Saussure, Ferdinand de 1968. *Cours de linguistique générale*. Paris: Payot.