Phonology 5 (1988) 275–292
Printed in Great Britain

*Underspecification in phonetics**

Patricia A. Keating
University of California, Los Angeles

It is often assumed in work on the phonological underspecification of segments that while representations may at first be underspecified, they end up fully specified. Various kinds of rules are posited to ensure that in output forms all features have values for all segments. In this paper I will consider an alternative view: that underspecification may persist into phonetic representations. I begin by reviewing some of the relevant phonological phenomena and mechanisms, and some of the history of underspecification in phonetic studies. I then show how phonetic data may be taken to reflect the presence or absence of feature values in surface forms. Finally, the extent to which surface specification depends on segmental contrasts is examined. The aim of this paper is to present some phonetic phenomena that are potentially relevant to any theory of underspecification.

1 Underspecification in phonological representations

1.1 Phonological behaviour of underspecified segments

Two types of arguments are often offered in support of phonological underspecification: what I will call arguments from variability and arguments from transparency. It is worthwhile to characterise these types of arguments as they are used in phonology, so that their use with respect to phonetic data can be considered later. An example of a phonological argument from variability is that a vowel undergoes vowel harmony, where a vowel's value for a certain feature depends on values of other vowels in the word. Because the vowel that varies seems to have no particular original value of its own, it may be analysed as being unspecified for that feature, and subject to a rule of harmony. That is, the segment is analysed as being underspecified because it takes on more than one value for a feature, as a result of some kind of feature assigning rule. Steriade (1987) dismisses this kind of argument because a segment could be originally specified, but undergo a delinking rule before subsequent feature assignment. However, Clements (1987) and others are more willing to argue from such cases. By contrast, Steriade stresses the value of arguments from transparency. In these cases, a segment is analysed as being unspecified for a feature because it seems to have no value at all. An example of an argument from transparency is that a consonant does not trigger, undergo, or block vowel harmony – that it is completely invisible to the harmony process.

It should be noted that these two situations are not incompatible: a given segment could behave transparently with respect to one rule, then acquire a feature value by a later rule. Transparency indicates that at a given point, a segment lacks a value; variability indicates that by another point, a segment has acquired a value. We can then ask by what kinds of rules underspecified segments acquire values.

1.2 Phonological rules to supply unspecified feature values

At least three broad types of rules are able to provide missing values of phonological features. One kind applies context-freely; two others apply context-sensitively. All of these kinds of rules are exemplified (though not distinguished in this way) in Stevens *et al.*'s (1986) discussion of redundant feature values, and some of their examples are repeated here.

Fill-in rules. This category comprises rules that fill in feature values without reference to context beyond the segment in question, and includes rules that either refer to other feature values of the segment, or refer to no other information. In the case of contrasts on some feature, it may be that only one value of the feature is underlying, so that some segments must receive their values by rule. For example, a fill-in rule may introduce the feature value [+voice] for sonorants, or the feature value [−low] for [+high] segments; or, if only [+voice] is underlying, a fill-in rule may introduce [−voice] for any segment lacking a value for this feature. Thus this broad category includes Archangeli & Pulleyblank's (forthcoming) default and complement rules, and Steriade's (1987) redundancy rules.

Sometimes, despite lack of contrast, a segment consistently has a particular phonetic quality, so that it behaves as if it has a value for the particular feature, albeit predictably. An example would be [s], which always appears to be phonetically [+high], though this feature value is not contrastive. The tongue body simply needs to be high to position the blade appropriately. (At the same time, [s] lacks the extreme fronting or backing that would make it palatalised or velarised.) Another example would be the relation between backness and redundant rounding for vowels in many languages: back vowels tend to be rounded, and front vowels unrounded. This kind of relation is called 'enhancement' by Stevens *et al.* (1986), meaning that the redundant value of one feature is chosen to enhance the distinction conveyed by another feature. Thus the value [+high] enhances the value [+strident] for the alveolar [s],[1] while [+round] enhances [+back] for vowels. Most of Stevens *et al.*'s examples are of this type. The crucial point is that these enhancement rules, like fill-in rules generally, supply feature values by table look-up, in these cases on the basis of the individual segment or segment classes.

Position rules. At least two kinds of context-sensitive rules are possible. In both cases a segment clearly seems to have a non-contrastive feature value, but that value depends on the context. I call 'position rules' those rules that fill in feature values on the basis of a segment's position in the string of segments, in syllable structure, etc., but without reference to melodic feature values of neighbouring segments. Stevens *et al.* give fewer examples of this kind of rule, but mention [+spread glottis] as enhancing [−voice] in syllable or word-initial position, and [+nasal] enhancing [+voice] in intervocalic position. The feature values assigned by position rules would then be available for later rules in just the way that values given by fill-in rules would be.

Context rules. I call 'context rules' those rules that fill in feature values on the basis of neighbouring segments' feature values. These are rules of assimilation, including harmony, and dissimilation. Stevens *et al.* extend their notion of enhancement to include the possibility that a feature value for one segment is chosen to enhance the value of a different feature in a neighbouring segment. They mention vowel lengthening as an enhancement of consonant voicing, as well as assimilation of certain coronal consonants to vowels with respect to backness. Also in this category would be rules that assign values of [back] to velars on the basis of contextual vowels; a case of this sort will be discussed below.

As a result of these three kinds of phonological feature fill-in rules, the output of the phonology may be considerably more specificed than are lexical representations. In principle, segments could be completely specified by these rules, and it is often assumed that indeed they are. Stevens *et al.* focus on the role of universal principles in motivating particular redundant feature values, especially through the fill-in kind of rule. They distinguish cases where redundant feature values are determined by inherent properties from cases determined by enhancement relations, but they do not state whether they believe that all redundant feature values are filled in by rules.

2 Underspecification in phonetic representations

Consider now another possibility for underspecified segments: that they remain underspecified even in surface representations. This situation could occur if for some reason no segmental feature rules operated to provide a given segment with some missing phonological feature value. Various specific interpretations of this general idea are possible, as will be discussed later, but for the moment let us say simply that some segments might exit the segmental, categorial, phonology with incomplete feature specifications.

The idea that even phonetic representations may contain underspecified segments is not new. Within speech science, during many years of work on coarticulation, some kind of phonological underspecification has been assumed, and in some proposals the underspecification persists into phonetic representations. I begin discussion of surface underspecification by reviewing some of this work. I do not thereby mean to claim that such proposals provide viable theories of underspecification, but only that the concerns which this work addressed may also be relevant to phonological theory.

Phoneticians working on coarticulation, or the articulatory overlap of segments, have assumed since at least the 1930s that coarticulation occurs in part because segments may lack inherent specification for particular articulations. That is, a given segment is characterised by some articulations, but is neutral with respect to other articulations. For example, an English alveolar consonant is characterised by a tongue blade articulation, but is neutral regarding lip rounding. Three influential proposals about coarticulation from the 1960s relied on such a distinction between specified and unspecified articulations. One proposal was from Kozhevnikov & Chistovich (1965): they proposed that the articulation of lip rounding in Russian was begun at the onset of the syllable, where 'syllable' simply meant a vowel plus any number of preceding consonants. Presumably such coarticulation over a whole syllable would not be possible if segments had various contradictory articulatory specifications, and so this proposal depended on underspecification of consonants for lip rounding.

A second influential contribution was Henke's (1966) computer model of the articulation of English stop + vowel sequences. The model took as input a sequence of segmental phonemes, associated each phoneme with an articulatory goal, and gave as output a sequence of vocal tract configurations (like those shown in X-ray tracings) at 1 msec intervals. Henke's main concern appears to have been the design and implementation of a computer program that could access and integrate various kinds of information needed to move articulators from one goal to the next. However, the work is best known for its 'look-ahead' mechanism for anticipatory coarticulation. Henke rejected the view that segments always have complete targets; a stop consonant will always have a goal for the stop occlusion, but not necessarily other articulatory goals. With 'look-ahead', Henke proposed that as soon as the stop contact is made (but no earlier), the stop looks ahead to the vowel's targets for other articulators. Henke found that anticipation of this very limited kind sufficed for the X-ray data on CV syllables he was modelling; Gay's (1977) X-ray data on tongue movements are similar to Henke's in this respect. Henke's proposal is thus similar to Kozhevnikov & Chistovich's; differences between them arise because they examined different segment sequences.

The third proposal, by Öhman (1966, 1967), was addressed to coarticulation between vowels. Öhman's model, like the others, assumed that segments are not specified for all articulators. However, Öhman's model incorporated a very explicit and restrictive version of this idea, namely that vowels and consonants are articulated by independent articulators, called 'channels of articulation'. Though Öhman's model was in terms of articulators, not phonological features, this is much like saying that vowels and consonants are specified for different sets of features. There is a set of features for which vowels are specified but consonants unspecified, and another set of features for which the reverse holds. I call this proposal 'complementary underspecification', since the sets of features are complementary across segments. With complementary underspecification, a given feature usually does not have values on adjacent segments, and this, in Öhman's view, is the source of vowel interactions: when articulatory values are connected up, segments without values will be invisible. With respect to vowel articulations, the influence of one vowel will extend through an adjacent consonant and right into the edges of the next vowel.

In all of these formulations, the basic descriptive dimension is the articulator and its gestures or goals, rather than phonological features *per se*. Thus all three incorporate underspecification relative to fairly shallow representations. Interestingly, this is not true of subsequent, more overtly phonological, work on coarticulation. In later work (e.g. Moll & Daniloff 1971; Benguerel & Cowan 1974; Kent *et al.* 1974; also Sharf & Ohde 1981), vocal tract dimensions and articulations were cast as binary features, on the model of binary features posited by linguists for phonological rules. The features could then be manipulated by rules of coarticulation. In particular, Henke's look-ahead mechanism for anticipatory coarticulation was extended to provide a 'feature spreading' mechanism.[2] Under feature spreading, binary feature values were copied from upcoming specified segments to unspecified segments, for an unbounded span of segments. The result was full specification of all segments for all features (except presumably word-final segments). To see how the feature copying worked, consider three segments S1, S2, S3, and three features, F, G, H, as in (1):

(1)
$$\begin{array}{llll} & \text{S1} & \text{S2} & \text{S3} \\ \text{F} & & & + \\ \text{G} & + & & - \\ \text{H} & & - & + \end{array} \rightarrow \begin{array}{lll} \text{S1} & \text{S2} & \text{S3} \\ + & + & + \\ + & - & - \\ - & - & + \end{array}$$

Thus segments are underspecified in their input forms, but not in their output forms. The model and its look-ahead mechanism say nothing about the precise realisation of these feature values; an additional component, such as Henke's own model, is still needed to describe when a given property would be observed in the physical signal. As a result of this lacuna, which was not acknowledged, the binary feature spreading model was found wanting in its ability to appropriately predict details of timing of coarticulation (Kent *et al.* 1974; Kent & Minifie 1977). Other difficulties are discussed by Gay (1978), Sharf & Ohde (1981) and Kent (1983). It should be noted also that the role of prosodic structure in determining coarticulatory effects was a matter of some interest that has never been resolved.

In sum, all of these accounts of coarticulation relied on underspecification at the level of representing segment contrasts; the underspecification assumed in this work was absence of a feature value for a class of segments. For example, in English the feature [nasal] would have a plus value for nasal stops, a minus value for oral stops, and potentially no value for any other segment. However, such representations were often conceived in terms of articulations rather than phonological features. To the extent that articulatory plans are phonetic representations, then some of these accounts also incorporated phonetic underspecification.

Other, more recent, work in phonetics has also assumed some degree of underspecification, though again not necessarily of features. For example, Fowler (1980) and others at Haskins Laboratories have adopted from Öhman the idea of separate vowel and consonant articulatory cycles, by which information about consonants is transparent to vowels and *vice versa*. Other information may also be unspecified in these representations, as in the work of Browman & Goldstein (1986). Their representations, which do not use segments or features, contain information about certain articulatory gestures while leaving other information unspecified. For example, a positive glottal spreading gesture will be specified in a representation, but in the absence of such a specification, a position suitable for voicing is to be assumed. Part of the job of their dynamic articulatory model is to provide such redundant information. Also proposing rather sparse representations is Pierrehumbert (1980 and later work). Her underlying representations for intonation patterns use tonal elements, which map into a fundamental frequency contour. A major result of this work has been that the number of tones needed to specify the fundamental frequency pattern is often far less than the number of syllables, or potential tone bearing units, in the text. See Pierrehumbert & Beckman (forthcoming) for further discussion of underspecification in their analysis of Japanese.

In sum, the idea of surface underspecification is not new in phonetics, and has been explicitly discussed in some current work. Let us consider now what kinds of arguments can be made about phonetic underspecification of segmental features.

In doing so, it is important to make clear which, if any, of the assumptions from the earlier work carry over into this discussion. As will be seen below, the assumption of complete surface specification that characterised the work on 'feature spreading' coarticulation will be abandoned. Furthermore, given the theoretical innovations of CV phonology, any assumption that surface representations consist of discrete segments in separate slices of time is unnecessary. Even in phonological representations, segments and feature values are not constrained to be in a one-to-one relation, and in phonetic representation internal timing structure is accessible. Feature values need not cooccur in time, and so can be variously aligned. Some of those feature values are underlying, while others are redundant and provided by rule.

I assume also that features have both articulatory and acoustic definitions, with one of the roles of features being to mediate between representations in these two physical domains. Each feature value is related somehow to physical targets. Targets can be context-sensitive, so it is not the case that a given feature value corresponds to one and only one numerical value along a physical scale. I also assume that the relation between features and dimensions may vary to some extent on a language-specific basis, with dimensions 'marshalled' to implement contrasts, somewhat along the lines of Stevens *et al.* (1986). For example, the feature [voice] will in one language marshall only the vocal cords, while in another language it will also marshall other articulators so as to sustain robust voicing across contexts. In the same way, a given acoustic dimension could reflect more than one feature at once. There is no direct relation between the individual articulatory and acoustic dimensions that the feature values map to.

From a physical record, such as an acoustic signal, we can infer something about the presence of the physical targets that correspond to the feature values. At best, then, we see targets, and not feature values directly. That is, strictly speaking, phonetic data bears only on specification for targets, not features. I assume, probably somewhat controversially, that identifiable targets in a physical signal reflect feature values, in that each feature value, and only feature values, get targets of some kind. Recall, however, that assimilations can arise through feature spreading. Thus a velar consonant can be unspecified for backness, but acquire a backness value from a contextual vowel, and therefore also acquire a target along the relevant physical dimensions. Where I take this to be controversial is with allophones that are not arguably the result of spreading. I interpret identifiable phonetic targets in such cases to reflect redundant feature values, which must have been supplied somewhere in the derivation.

3 Phonetic effects of phonetic underspecification: phonetic transparency

Given a framework allowing surface underspecification, suppose a segment exits from the phonology with no value for a feature, having failed to receive one either lexically or by a later rule. That segment will then be transparent in the phonetics to any rule sensitive to that feature. In particular, it cannot be assigned a quantitative target, and so when phonetic rules build trajectories between segments, an unspecified segment will contribute nothing of its own to the trajectory. In this section I describe cases that illustrate this point. There are two ways in which the phenomenon can be examined, both of which are analogues of arguments from phonological transparency. First, the underspecified segment itself can be studied: is it, on the relevant dimension(s), entirely dynamic or transitional in nature, lacking any steady state portion or non-transitional movement, as would be expected if it has no feature value of its own to contribute? Second, the effects of the underspecified segment on its neighbours can be studied: it should be effectively transparent and should permit its neighbours to interact with each other by, for example, showing mutually dependent formant transitions. The next sections consider examples of such effects.

3.1 Segment-internal effects

The transparency of an underspecified segment is perhaps clearest in the case of /h/. As a glottal approximant, /h/ has no intrinsic oral feature values, and is usually described by phonologists as a voiceless version of the following vowel, the result of feature spreading assimilation. But in connected speech, /h/ is not assimilated to a following vowel. Rather, in most instances /h/ is transparent; the acoustics of intervocalic /h/ show that rather than acquiring feature values, /h/ is simply interpolated right through. Figs. 1 and 2 show wideband spectrograms from three languages in which certain formant frequencies of /h/'s clearly interpolate between the surrounding vowels. The second formant is outlined in the figures to aid the reader; note that the surrounding vowels provide turning points in the formant trajectories, while the /h/'s do not. The apparent assimilatory effect on the /h/ is a dynamic, transitional one, not a static one. /h/ starts out unspecified for oral features, and it remains so. That is, sometimes 'assimilation' is a lower-level phonetic effect, rather than spreading of phonological feature values. With /h/, all of the oral features are involved, so the effect on the formant frequencies is quite striking. Other continuant segments are specified for at least some place of articulation features, so show less obvious, though often still visible, effects of this kind.

3.2 Phonetic effects on context

Just as in the phonology, a missing feature value should be apparent from a segment's failure to block a rule from applying across it. A useful test for phonetic underspecification of consonants (though one that sometimes requires judicious interpretation) can be found in the presence or absence of vowel interactions across consonants. Vowel-to-vowel coarticulation, or the effect of one vowel on another even across an intervening consonant, was described by Öhman (1966). He found that in English and Swedish there was a small but consistent effect of stressed vowels on the transitions of other stressed vowels; for example, in a VCV, the frequency value of the onset of the second formant for the second vowel depends partly on the offset value of that formant during the first vowel. Others have since found stronger effects on the centres of unstressed vowels in English (Fowler 1981; Bell-Berti & Harris 1976; Magen 1984; Huffman 1986), and on vowels in other languages (Magen 1984; Manuel & Krakow 1984; McAdams 1987). Though sometimes subtle, vowel-to-vowel effects can be quantified by measuring variation in vowel formant frequencies across contexts. Fig. 3 illustrates a vowel-to-vowel effect across /b/ in English. In these spectrograms the second formant values closest to the consonant closure are marked, with the value for an /a/ compared before /a/ *vs.* /i/. The second formant of /a/ is lower before /a/ than before /i/. Many other examples of these effects are available in the references cited.

Öhman explained vowel-to-vowel coarticulation through consonants as a consequence of the underspecification of consonants for vowel

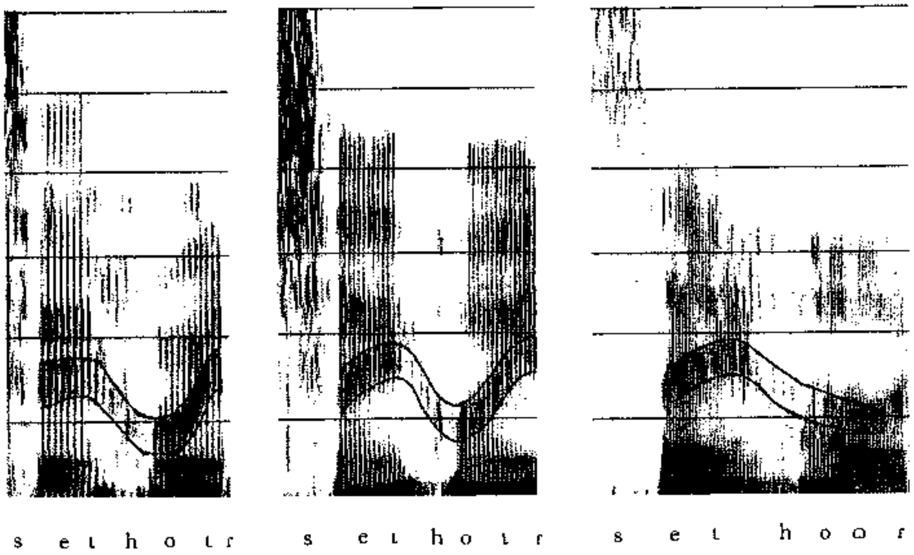

Figure 1

Examples of English intervocalic /h/, in the words *Hoyd* and *hoed*, taken from sentences 'Say h— again'. The second formant trajectory is outlined. Calibration lines are at 1000 Hz intervals.

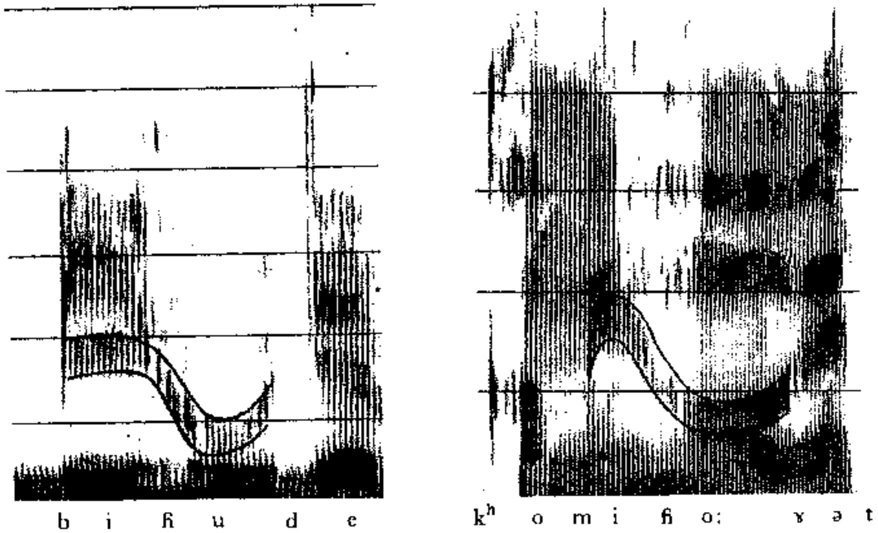

Figure 2

Examples of intervocalic /ɦ/ in Farsi *bihude* 'useless' and Swedish *komihåg att* 'remember to (do something)'. Both are strongly voiced. The second formant trajectory is outlined. Though the vertical scale differs in these two spectrograms, in both cases calibration lines are at 1000 Hz intervals.

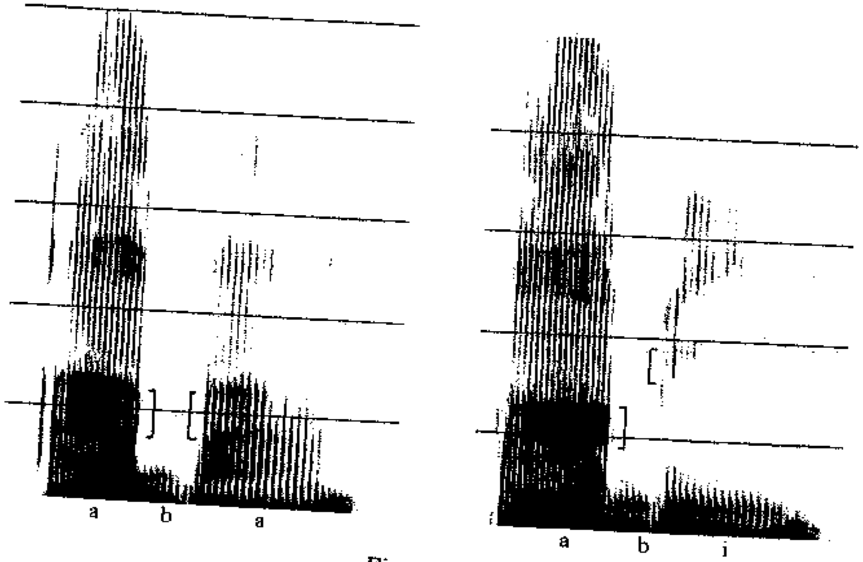

Figure 3
Example of vowel-to-vowel coarticulation across /b/ in English nonsense items. The offset value for the second formant of the first vowel differs in the two items because the second vowels differ. Calibration lines are at 1000 Hz intervals.

articulations. In his data (contrary to Henke's) the tongue body made a single smooth gesture from vowel to vowel, as if for a diphthong, overlapping with any consonant gestures. On his account, the influence of the first vowel extends right through the consonant (which the vowel gesture does not see) and into the next vowel, and *vice versa* in the other direction.

4 Comparison of output types

Particular acoustic patterns can be observed during segments that plausibly lack certain feature values. Thus these acoustic patterns suggest tests that can be applied to determine whether a segment has a particular feature value. However, it is important to compare the results expected as outputs of the various rule types, so as to make clear how they are to be distinguished.

4.1 Context rules *vs.* phonetic rules

It is not trivial to distinguish the output of context rules, which provide segmental feature values on the basis of contextual values, from the output of phonetic rules, which carry out all subsequent quantitative steps, ending with building continuous contours along phonetic dimensions. In the case at hand, that of underspecified segments, the question is how to distinguish acquisition of a feature value, which allows the segment to

acquire a phonetic target along the relevant dimension(s), from lack of a feature value (and hence lack of a target), but interpolation through the segment. The difference seems clear enough in principle: if a segment acquires a feature value from an adjacent segment, it will share a phonetic property with that segment across most or all of its duration; if a contour is built through a segment it will have a more or less continuously changing, transitional, quality from beginning to end that will depend on context on either side. The examples given above of phonetic rules display this quality. In practice, the intended distinction may not always be obvious; however, an example of a context rule whose output is clearly different from the output of phonetic rules is provided by fronting of the velar fricative /x/ in Russian. Although most Russian consonants contrast as either palatalised or not, /x/ does not, but instead depends on context for its value for [back].

If a segment acquires a feature value from a neighbour, then it should display the relevant property much as the neighbour does. Fricatives allow us to examine the time course of any assimilatory effect. To this end, Fig. 4 compares Russian /x/ in mirror-image contexts before and after /i/, as well as between two /a/ vowels. The second formant, which reflects backness, is indicated on the spectrograms in this figure. The two tokens of /x/ adjacent to /i/ are very different. In /ixa/, the fricative shows a continuous change from more front to more back, like other examples of phonetic rules considered above. By contrast, in /axi/ the fricative is extremely fronted throughout its duration. This steady quality of the /x/ is very different from the preceding /a/, and necessitates an extreme transition, especially of the second formant, between the two segments. This phenomenon provides support for the idea that the fricative is [−back] over its entire duration, rather than having only a transitional front quality. The /x/ before /i/ acquires a value for [back] by spreading from /i/, whereas other /x/'s remain unspecified for [back], and therefore show interpolatory effects. (When the two flanking vowels are identical, then interpolation will result in a steady transition, and thus reveal little about the nature of the transitions.) The same is observed for other speakers of Russian, and in fact in other Slavic languages, and will be discussed in detail in a later paper.

4.2 Fill-in rules

Fill-in values, once assigned, should presumably have the same status as underlying contrastive feature values. One test of this would be whether a fill-in value on a consonant blocks vowel-to-vowel coarticulation across that consonant. As discussed earlier, the consonant [s] in English is usually made with a very high tongue body position. Tongue height is reflected acoustically mainly in the frequency of the first formant, but since during voiceless fricatives this formant is not visible, only effects on neighbouring segments can be examined. The lowest two formant frequencies were examined for effects of vowel-to-vowel coarticulation across [s] in VCV

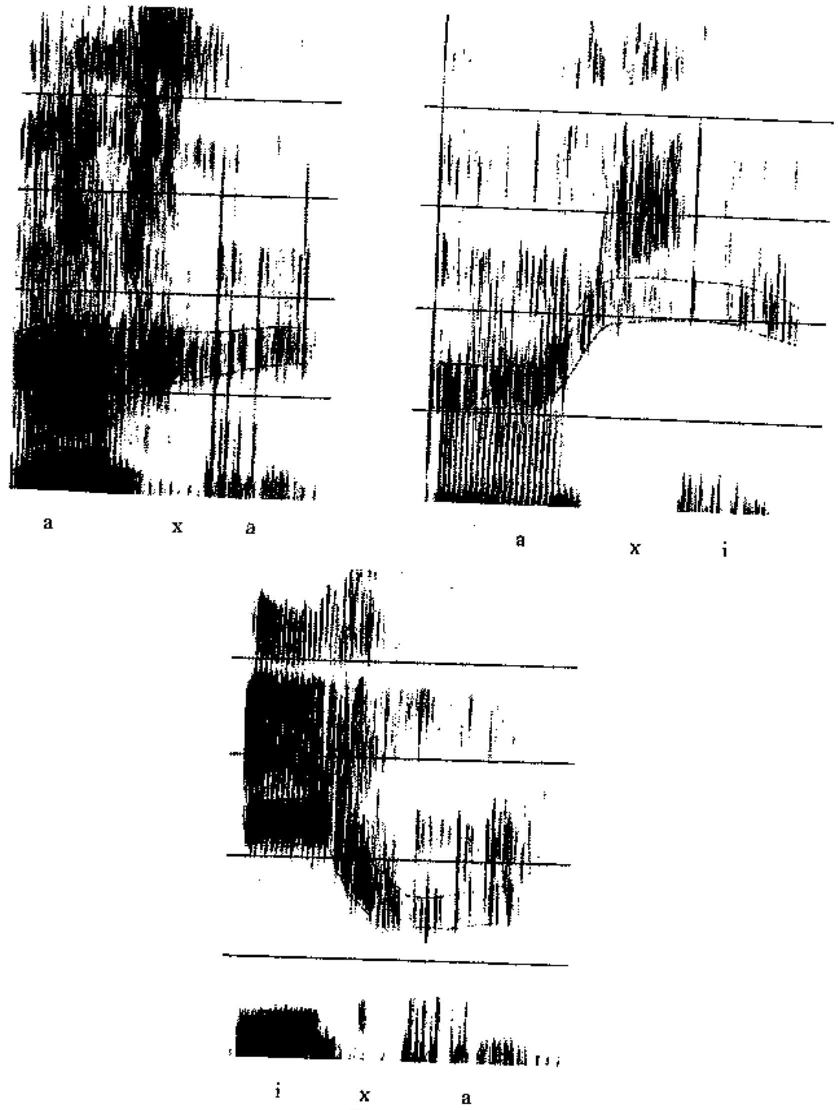

Figure 4

Example of fronting of Russian /x/ before /i/. Top left, steady low frequency component in fricative. Top right, before /i/, steady higher frequency component in fricative. Bottom, after /i/, transitional component in fricative between the two qualities of top row examples.

nonsense items. The vowels were /i/, /a/ and /u/, with the first vowel receiving main stress. Two repetitions of each were measured for two speakers of American English, using wideband spectrograms to measure the formant frequencies at the end of V1 and the beginning of V2. If vowel-to-vowel coarticulation occurs, then the formant frequency for one vowel will depend in part on the identity of the other vowel. For F2, the amount of coarticulation found was similar to what is reported in the literature across other consonants, with carry-over effects stronger than anticipatory. In tokens where F2 can be seen faintly during the [s], it is clearly moving from the first to the second vowel's value. On the other hand, F1 is consistently low in the vowels around the [s] and shows little if any coarticulation. That is, F1, unlike F2, seems to have a fairly fixed value for [s]. (It must be noted that stop consonants have been found to allow vowel-to-vowel coarticulation in F1, so the effect with [s] is not a general property of F1.) Interpreting these data in terms of feature values, we can say that [s] behaves as if it has a value for [high], namely [+high], this value presumably the result of a fill-in rule; and that [s] behaves as if it had no value for [back].

4.3 Three-way comparison

For a given segment, three different output representations may emerge from the phonology, with different acoustic reflexes. These possibilities are illustrated by the schematic representations shown in (2), using VCV sequences:

(2)
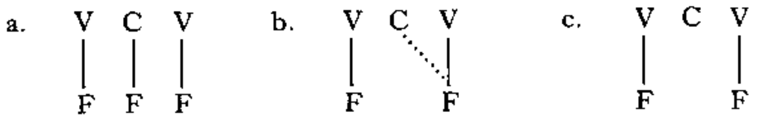

In (a), each segment has its own value for the feature F, and so each will have its own phonetic quality. There should be no vowel-to-vowel effects in either direction, because the feature value on the consonant will block them; there may be extreme transitions between segments that differ in values for F. In (b), the consonant has acquired a value for F from V2, by a spreading rule. There will be no effect of V1 on V2, because the consonant's value for F again serves to block interactions. However, there will appear to be an effect of V2 on V1, because the consonant's feature value will affect V1. The transitions between V1 and C will look just like those of case (a), but since in (b) C's value depends on V2, V1's variation will depend on V2 as well as C. Thus there will be apparent asymmetrical vowel-to-vowel effects, but these are instead purely local. Finally, in (c), there will be vowel-to-vowel effects in both directions, and the consonant will lack its own phonetic quality, showing instead only a gradual transitional quality.

Blocking vowel-to-vowel coarticulation, as in (a), simply shows that a segment has a value for a feature in the output from the phonology. It does not indicate whether that value was present underlyingly, or whether it is a later fill-in value. The phonology of the language (the system of contrasts, behaviour of the segment in other rules) still has to be consulted to discover the source of the feature value.

5 Role of contrast in phonetic behaviour

In the several versions of underspecification theory elaborated recently, systems of contrast do not by themselves determine what features will be underspecified for what values, but they enter in as one factor. For segment classes in which a given feature is distinctive, either one or both values may be given in underlying representations. The kinds of phonetic data described above can bear on the nature of underspecification, providing evidence about which segments are underspecified for which features. Evidence from vowel-to-vowel interactions shows that some, but not all, segments behave as if they have feature values on the surface. In this section I present an example where segment contrasts affect coarticulation, and then discuss the possibility that even certain contrasting segments may lack feature values.

5.1 Underspecification depending on contrast within a segment class

In general, when a segment bears a feature value by virtue of participating in a contrast, we expect that feature value to block phonetic vowel-to-vowel interactions; and when a segment does not contrast and has no feature value, we expect it to allow such interactions. This should be so even within a segment class and within a language. More attention to contrast seems to resolve an apparent contradiction in the literature about vowel-to-vowel coarticulation in Russian. Öhman's model with channels of articulation predicts no coarticulation across Russian stops that contrast in palatalisation, and both Öhman and Purcell (1979) found none. However, the same prediction is made for contrasting fricatives, yet Derkach *et al.* (1970) report anticipatory vowel-to-vowel coarticulation across Russian voiceless fricatives. This result appears to be a counterexample to Öhman's model (though Derkach *et al.* did not make this interpretation), and to the claim that phonetic data can bear on questions of feature specification. However, Derkach *et al.* do not report their results in great detail, and they appear to have treated all the fricatives as a single group, not differentiating the various places of articulation, even though some places have pairs contrasting in palatalisation while others do not. The surface contrasts in voiceless fricatives are shown in (3). The palato-alveolar [ʃ] is usually described as redundantly [+back] (it takes back vowel allophones), and for many speakers it contrasts on the surface with a palatalised segment derived through cluster simplification, shown here in parentheses. The velar /x/ is described as allophonically variable. Even speakers who allow a contrast in the velar stops have none for the fricative:

(3) f s ʃ x
f’ s’ (ʃ’)

Examination of my own data on individual Russian fricatives reveals that the various fricatives do not behave identically. In particular, vowels before the velar fricative show strong effects of following vowels. There may also be lesser effects across the palato-alveolar, but my data are unclear on this point. The clear effects observed with /x/ are due to the assimilation of the fricative to a following /i/, as discussed above. Because the /x/ is affected by the /i/, and because the effect is so strong and prolonged, the /x/ itself influences the preceding vowel, as seen in Fig. 4 above. An extreme transition is required from the vowel to the /x/, and this gives rise to strong measured vowel interactions. That is, what looks like a vowel-to-vowel (through the consonant) effect in this case is actually a vowel-to-consonant-to-vowel effect. It is not really an appropriate situation to examine for vowel-to-vowel effects, since phonetically the consonants are not the same at all. Conceivably this effect alone accounts for Derkach *et al.*'s finding; grouping all fricatives together obscures this principled difference. The status of the palato-alveolar with respect to these phonetic effects remains to be determined.

5.2 Underspecified contrasts

When two segments contrast within a class, must both the plus and minus values of the relevant feature be used? Phonologists distinguish cases in which both feature values characterise a contrast from those in which only one feature value is present. Thus Steriade (1987) includes [round] as necessarily single-valued, and [voice] and [lateral] as underlyingly single-valued in some languages. According to Steriade, it is rare for a feature to be both distinctive and clearly single-valued within a language, but a few fairly clear cases are nonetheless known.

The same hypothesis can be entertained for phonetic representations, and in fact has been the subject of discussion in the traditional coarticulation literature. In some cases, both members of a contrast behave as if they possess a feature value. In articulatory terms, this is equivalent to saying that one member has one articulation, while the other member has a contrary articulation, such as lip rounding *vs.* lip spreading. But in other cases, only one member of the contrast behaves as if it possesses a feature value. This is equivalent to saying that only one member has an articulation, such as lip rounding *vs.* no rounding.

Again, the occurrence of vowel-to-vowel coarticulatory effects can be taken as an indication that a segment lacks its own feature value. For example, consonants with contrastive secondary articulations, differentiated by their values for the feature [back], can be examined to see whether each consonant blocks vowel interactions, as would be expected if each has its own feature value. If only one member of the contrast is specified for [back], then that segment should block vowel interactions, while the other member should be transparent to them. Published data on Russian stop consonants, and other data on fricatives, were discussed above. These data suggest that in Russian both members of the contrasting pairs of consonants are specified for [back], since all contrasting consonants seem to block vowel-to-vowel coarticulation.

However, other data on secondary articulations also suggest the opposite: that in some languages, only half of the contrasting consonants are specified for [back]. One case is that of pharyngealisation in Arabic. Dialects differ, but the basic idea is that certain coronal obstruents (plus sometimes other consonants) are paired as having or not having the tongue body in a low back position. The question here is whether the consonants that are not pharyngealised have some other property of their own by way of contrast. My own preliminary data suggest that they do not, that is, that they are simply non-pharyngealised and otherwise variable, and therefore allow vowel interactions. The same finding for Arabic is presented by Hussein (1987). That is, a feature such as [back] may in some cases (such as Russian) occur in a given segment class with both of its values, and in other cases (such as Arabic) occur with only one value. This in effect restates traditional descriptions of the phonetic nature of these contrasts.

In general, then, the occurrence or lack of vowel interactions across a consonant is informative and suggestive. Nonetheless, it must be noted that a single instance of such effects cannot be definitive. There are many factors that may affect the occurrence or strength of vowel interactions, including the possibility of language-specific prosodic constraints and rate of speech effects. The intrinsic properties of various consonants will affect the degree of vowel interactions (Huffman 1986). Furthermore, languages may differ in how coarticulatory effects are distributed within a sequence of segments. Therefore it is overall patterns of coarticulatory effects that provide the best evidence.

6 Conclusions

Phonetic data provides detailed information about sound structure that allows us to detect and distinguish different types of assimilatory effects, or the lack of them, in segments and in their contexts. Underspecification theory provides a mechanism for describing where such effects should occur: different acoustic patterns correspond to differences in surface feature specification. Acoustic data reveal that a non-contrasting segment may behave as if it has acquired a feature value through assimilatory rules of feature spreading; or as if it has acquired one without reference to context; or as if it lacks any such value on the surface.

The claim that phonetic evidence can be used to help diagnose feature specification has an important implication for the learnability of highly underspecified representations. The occurrence of coarticulatory patterns in speech, to the extent that they are audible, could be useful evidence to a language-learner as to what features are specified in the phonetic form for particular segments, and how. With such phonetic evidence, the learner would not have to accomplish a complete phonological analysis of the language, so as to replace fully specified entries with underspecified ones. Instead, immediately accessible surface information could form the basis for an initial distinction between the three representations shown in (2) above. Of course, this information only reflects surface specification *vs.* underspecification, and therefore provides evidence only about what could have been originally unspecified. Additional analysis by the learner would be required to arrive at appropriate underlying representations. The inventory of underlying segment contrasts provides limits on what must minimally be specified; phonetic evidence about surface representations provides substantially more information to the learner.

NOTES

* This work was supported by the NSF under grant BNS84-18580. An early version appears in *UCLA Working Papers in Phonetics* **69** (February 1988). Some of this material was presented at various conferences: in Lyon, France at the June 1985 colloquium on multi-linear phonology (Keating 1985); in Seattle at the December 1985 LSA meeting; at the UCLA conference on segment structure in October 1987; and in talks elsewhere. I thank those who offered suggestions or comments on those occasions, as well as my colleagues in the UCLA Phonetics Lab.

[1] Stevens *et al.* do not give this example.

[2] Like others, I have in the past mistakenly ascribed this view to Henke.

REFERENCES

Archangeli, D. & D. Pulleyblank (forthcoming). *The content and structure of phonological representations.* Cambridge, Mass.: MIT Press.

Bell-Berti, F. & K. Harris (1976). Some aspects of coarticulation. *Haskins Laboratories Status Report on Speech Research* **SR45/46**. 197–204.

Benguerel, A.-P. & H. Cowan (1974). Coarticulation of upper lip protrusion in French. *Phonetica* **30**. 41–55.

Browman, C. & L. Goldstein (1986). Towards an articulatory phonology. *PhY* **3**. 219–254.

Clements, G. N. (1987). Toward a substantive theory of feature specification. Paper presented at the UCLA Symposium on Segment Structure, October 1987.

Derkach, M., G. Fant & A. de Serpa Leitão (1970). Phoneme coarticulation in Russian hard and soft VCV-utterances with voiceless fricatives. *Speech Transmission Laboratory Quarterly Progress and Status Report* **2–3**. 1–7.

Fowler, C. (1980). Coarticulation and theories of extrinsic timing. *JPh* **8**. 113–133.

Fowler, C. (1981). Production and perception of coarticulation among stressed and unstressed vowels. *Journal of Speech and Hearing Research* **46**. 127–139.

Gay, T. (1977). Articulatory movements in VCV sequences. *JASA* **62**. 183–193.

Gay, T. (1978). Articulatory units: segments or syllables? In A. Bell & J. B. Hooper (eds.) *Syllables and segments.* Amsterdam: North-Holland. 121–131.

Henke, W. (1966). *Dynamic articulatory model of speech production using computer simulation.* PhD dissertation, MIT.

Huffman, M. (1986). Patterns of coarticulation in English. *UCLA Working Papers in Phonetics* **63**. 26–47.

Hussein, L. (1987). Coarticulation in VCV sequences in Arabic. *JASA* **81**. S68.

Keating, P. (1985). CV phonology, experimental phonetics, and coarticulation. *UCLA Working Papers in Phonetics* **62**. 1–13.

Kent, R. (1983). The segmental organization of speech. In P. MacNeilage (ed.) *The production of speech.* New York: Springer. 57–90.

Kent, R., P. Carney & L. Severeid (1974). Velar movement and timing: evaluation of a model for binary control. *Journal of Speech and Hearing Research* **17**. 470–488.

Kent, R. & F. Minifie (1977). Coarticulation in recent speech production models. *JPh* **5**. 115–133.

Kozhevnikov, V. & L. Chistovich (1965). *Speech: articulation and perception.* Translated and distributed by Joint Publications Research Service, Washington DC.

McAdams, B. (1987). The relation between vowel variation and vowel-to-vowel coarticulation across languages. *JASA* **81**. S66.

Magen, H. (1984). Vowel-to-vowel coarticulation in English and Japanese. *JASA* **75**. S11.

Manuel, S. & R. Krakow (1984). Universal and language particular aspects of vowel-to-vowel coarticulation. *Haskins Laboratories Status Report on Speech Research* **SR77/78.** 69–78.

Moll, K. & R. Daniloff (1971). Investigation of the timing of velar movements during speech. *JASA* **50**. 678–694.

Öhman, S. (1966). Coarticulation in VCV utterances: spectrographic measurements. *JASA* **39**. 151–168.

Öhman, S. (1967). Numerical models of co-articulation. *JASA* **41**. 310–320.

Pierrehumbert, J. B. (1980). *The phonology and phonetics of English intonation.* PhD dissertation, MIT.

Pierrehumbert, J. B. & M. Beckman (forthcoming). *Japanese tone structure.* Cambridge, Mass.: MIT Press.

Purcell, E. (1979). Formant frequency patterns in Russian VCV utterances. *JASA* **66**. 1691–1702.

Sharf, D. & R. Ohde (1981). Physiological, acoustic, and perceptual aspects of coarticulation: implications for the remediation of articulatory disorders. In N. Lass (ed.) *Speech and language: advances in basic research and practice.* Vol. 5. New York: Academic Press. 153–247.

Steriade, D. (1987). Redundant values. In A. Bosch, B. Need & E. Schiller (eds.) *Papers from the parasession on metrical and autosegmental phonology.* Chicago: Chicago Linguistic Society. 339–362.

Stevens, K., S. J. Keyser & H. Kawasaki (1986). Toward a phonetic and phonological theory of redundant features. In J. Perkell & D. Klatt (eds.) *Symposium on invariance and variability of speech processes.* Hillsdale: Lawrence Erlbaum. 432–469.