On Distinctive Features and Their Articulatory Implementation
Author(s): Morris Halle
Source: *Natural Language & Linguistic Theory*, Vol. 1, No. 1 (1983), pp. 91-105
Published by: Springer
Stable URL: http://www.jstor.org/stable/4047515
Accessed: 14-04-2018 01:04 UTC

MORRIS HALLE

ON DISTINCTIVE FEATURES AND THEIR ARTICULATORY IMPLEMENTATION

To the memory of Beatrice Hall

1. One of the first observations that students in an introductory phonetics course make is that the gestures which the vocal tract executes in producing a given sound are readily analyzable into more elementary components or sub-gestures which, in combination with other sub-gestures, are also utilized in the production of other speech sounds. Thus, we find identical lip closure in each sound of the set [pbm], whereas the sounds in the set [kgŋ] are all produced by the tongue body making contact with the velum. The two sets, moreover, each contain one consonant produced with a lowered velum [m,ŋ] and two with raised velum, [pb, kg]. Looked at from a different point of view, the consonants under discussion include four that are produced with vocal cord vibration [bgmŋ] and two without such vibration [pk]. Such observations can readily be summarized in the familiar tabular form illustrated in (1), where each sound is represented as a complex of features:

(1)

	p	b	m	k	g	ŋ
labial	+	+	+	−	−	−
velar (high)	−	−	−	+	+	+
nasal	−	−	+	−	−	+
voiced	−	+	+	−	+	+
stop (closure)	+	+	+	+	+	+

Evidence for the composite structure of speech sounds emerges also when we examine them as purely acoustic events. Thus in a speech spectrogram of the above set of sounds we can readily see the abrupt transitions in the spectral pattern that mark the beginning and end of stop sounds. Similarly, we can discern in the acoustic signal the difference between nasal and nonnasal segments as well as that between voiced and voiceless. The acoustic correlates of labiality and velarity, however, are much less simple, as they depend to a great extent on the nature of the surrounding sounds. When these consonants occur next to vowels they produce specific changes in the vowel formants – the so-called 'vowel formant transitions' – which provide crucial information about the 'point of articulation' of the consonants. Thus, labial consonants are marked by a 'negative (i.e., downward) transition' of all vowel formants, whereas unrounded velars are marked by transitions in which formants 2 and 3 come together in a single frequency. These particular cues are, of course, not available when labial or velar consonants are found

Natural Language and Linguistic Theory **1** (1983) 91–105. **0167–806X/83/0001–0091** $01.50

between other consonants rather than adjacent to a vowel, as in such words as *asps* vs. *asks* or *lisped* vs. *risked*. In such cases the acoustic cue signalling the difference between labial and velar stops (e.g., between [p] and [k]) must reside in the spectral properties of the stop burst. Since vowel formant transitions and spectral properties of a stop burst are rather different phenomena, which are unlikely to have a plausible common denominator, we conclude that in the case of labiality and velarity in stops we have two distinct acoustic cues correlated with a single articulatory property. In (2) we illustrate graphically the relationship described above as it pertains to labial sounds.

(2)

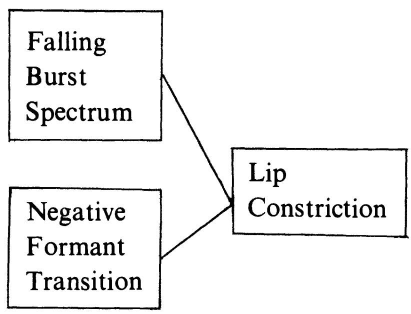

This is not the only example of a one-to-many relationship between articulatory and acoustic properties of speech. Stevens and Halle (1971) drew attention to the fact that differences in vocal cord stiffness have vastly different effects in sounds produced with a small pressure drop across the vocal cords than in sounds produced with a large pressure drop across the vocal cords. When the pressure drop is small, as it is in all obstruents, a moderate increase in vocal cord stiffness makes it impossible to sustain any vocal cord vibrations. On the other hand, when the pressure drop is relatively large – as it is in vowels – a comparable increase in vocal cord stiffness will result in an increase in the frequency of vocal cord vibrations. Thus, the articulatory gesture of increasing the stiffness of the vocal cords is correlated with the acoustic distinction of voicing vs. voicelessness in the case of obstruents, whereas in the case of vowels the same changes in vocal cord stiffness elicit changes in the fundamental pitch.

It should be noted at once that the converse relationship is also attested: i.e., there are well known instances where a given acoustic effect is produced by several distinct articulatory means. In fact, the cessation of vocal cord vibration, which is manifested by the suppression of low frequency periodicity in the signal can be produced in obstruents not only by increasing the stiffness of the cords, but also by spreading them apart. The acoustic cue of voicelessness can, therefore, be produced by two articulatory means: vocal cord stiffness and vocal cord spreading. We have thus encountered a situation that graphically may be represented as in (3)

(3)

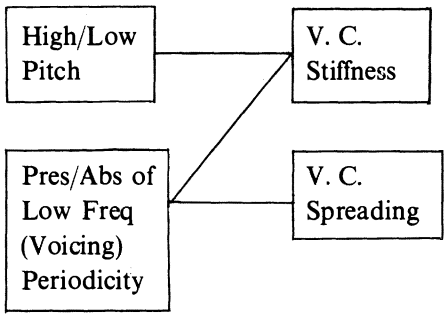

Similar situations arise with respect to other acoustic and articulatory properties of speech.[1]

In addition to articulatory and acoustic data there is a third body of facts that must be taken into account in phonetic studies; these are the facts that derive from the phonological rules of different languages. As is well known, phonological rules characteristically involve not individual speech sounds, but rather whole groups of sounds. We give two examples of rules in (4a) and (4b).

(4) a. p → f t → θ k → x in certain contexts
  b. w → u before [p, b, m]

The first of these rules is part one of Grimm's Law and describes a synchronic process that took place in Proto-Germanic. The second rule is a well known morphophonemic rule of Hebrew grammar governing the realization of the conjunction **w** 'and'.

What is noteworthy about the two rules is that they involve not just arbitrary sets of speech sounds but speech sounds that share specific phonetic properties. Thus in (4a) the affected set of sounds is the voiceless stops,

[1] K. P. Mohanan has drawn my attention to the fact that in talking of the articulatory correlate of voicing, I focus on such aspects of the phenomenon as stiffness of the vocal cords and the extent of their spreading, and have little if anything to say about the rapid opening and closing movements that are the result of the adjustments in vocal cord stiffness and spreading. The reason for the relative neglect of vocal cord vibrations here is that I am interested above all in aspects of the articulatory behavior that are under voluntary control of the speaker and only secondarily in those aspects that are the automatic consequences of the former.

whereas in rule (4b) the set involved is that of the labial consonants. And this situation is quite typical: phonological rules in the most varied languages involve groups of sounds that are readily characterized in terms of their phonetic properties, whereas rules involving such phonetically unnatural groups as [mtk] or [pbk] are unattested or extremely rare.

It might have been noticed that in the last section of this paper repeated use has been made of the phrase 'phonetic properties' as a cover term for what in earlier sections we had referred to as articulatory or acoustic properties. The reason for this substitution was that the sets of sounds found in some rules have simpler characterizations in the articulatory domain than in the acoustic domain, whereas for sets of sounds involved in other rules the converse is the case. In fact, the two rules (4a) and (4b) were chosen to provide examples of the two situations. In Grimm's Law (4a) the phonetic feature is absence of voicing which as indicated above (cf. (3)) is a single acoustic property with two articulatory correlates. On the other hand, the Hebrew rule (4b) makes crucial use of the property of labiality which has a single articulatory actualization with two distinct acoustic correlates (cf. (2)).

It might be inferred from the above that in defining the classes of sounds that figure in different phonological rules a language has a free choice between articulatory and acoustic properties. As a matter of fact this inference is not justified. When a greater number and variety of rules is considered, it becomes clear that languages never avail themselves of this freedom of choice. For instance, there are no languages that exploit acoustic distinctions between labial stops illustrated in (2), and we find no rules that differ from the Hebrew rule (4b) in that they affect stops with falling burst spectra but not stops with negative vowel formant transitions, or vice versa. Similarly there are no languages that exploit the articulatory distinction between voiceless stops illustrated in (3), and we find no rules that differ from (4a) in that they involve obstruents produced with increased vocal cord stiffness but not obstruents produced with spread vocal cords.

Considerations of this nature were much in our minds thirty years ago when Jakobson, Fant and I were working on *Preliminaries to Speech Analysis*, and it was these considerations that led us to draw a sharp distinction between distinctive features, which were abstract phonological entities, and their concrete articulatory and acoustic implementation. Thus, in *Preliminaries* we spoke not of "articulatory features" or of "acoustic features," but of "articulatory" and/or "acoustic correlates" of particular distinctive features. The model we had in mind was, therefore, of the type represented by the block diagram in (5), where the abstract distinctive features constitute the link between specific articulatory and acoustic properties of speech sounds.

(5)

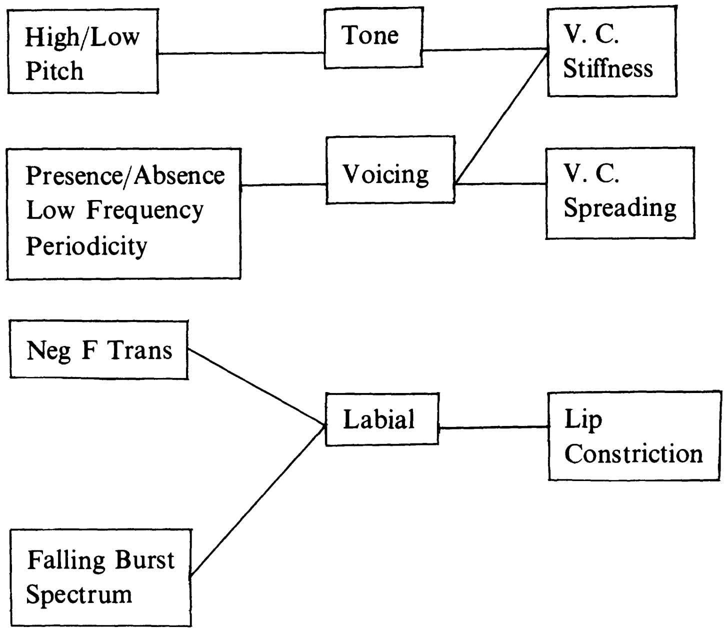

2. I want to propose at this point that the diagram (5) is more than a convenient graphic illustration of the logical structure of one theory of phonetics, namely that of *Preliminaries to Speech Analysis*, but that it should also be viewed as a proposal concerning the organization of the phonetic faculty in humans. On this view the distinctive features correspond to controls in the central nervous system which are connected in specific ways to the human motor and auditory systems. In speech perception detectors sensitive to the property detectors on the left hand side are activated, and appropriate information is provided to centers corresponding to the distinctive feature boxes in the middle of the diagram. This information is forwarded to higher centers in the nervous system where identification of the utterance takes place. In producing speech, instructions are sent from higher centers in the nervous system to the different feature boxes in the middle part of (5) about the utterance to be produced. The features then activate muscles that produce the states and configurations of different articulators listed on the right hand side of the diagram (5). Our next task is, therefore, to examine the model in some detail in order to satisfy ourselves that further exploration has a chance of yielding some worthwhile results. My discussion below focusses exclusively on speech production, i.e., on the middle and right-hand part of the diagram (5). This restriction is due not to a feeling on my part that perception is any less important than production but rather

because at this stage in our study of language, we have a somewhat better grasp of the issues in the articulatory domain than in that of speech perception and processing.

It should be noted at the outset that diagram (5) does not imply that the articulatory states in the right hand column can originate only in consequence of commands emanating from the feature centers. There are other ways in which our lips, tongue, larynx, etc., can be activated, but these are distinct and separate from what transpires in the production of speech. We recall in this connection Sapir's famous discussion of the difference between the blowing out of a candle and the production of a voiceless bilabial fricative [φ]. While the concrete physical events, i.e., the movements of the lips, tongue, larynx, may well by identical in the two cases, Sapir lists a series of essential differences between the two phenomena. Without going into the nature of these differences it is clear that the model in (5) provides us with a way to express these difference formally. When we produce a voiceless bilabial fricative [φ] the articulatory gesture is produced in response to commands from the distinctive feature centers in (5); when we blow out a candle the same (or very similar) vocal tract gymnastics are produced with commands from central nervous system centers different from the distinctive feature centers in (5).

It was remarked above with regard to the proposed model of the speaking process in (5) that the distinctive features activate muscles which move articulators into particular configurations and states. While this is a perfectly plausible way of viewing the speaking process, it is not the one generally adopted by phoneticians. In the more traditional approach exemplified, for example, by *The Principles of the International Phonetic Association* (referred to below as the IPA system), the geometric configuration of the vocal tract is characterized by means of the location of the maximal constriction (point of articulation) in the case of consonants, and by means of the location of the highest point of the tongue arch in the case of vowels. The assumption, rarely if ever stated explicitly, is that the rest of the vocal tract configuration can be deduced from this information. We note, in addition, that whereas the consonantal point of articulation is located at different landmarks on the roof of the mouth and the back wall of the pharynx, the highest point of the tongue arch in terms of which vowels are characterized, is specified with respect to a pair of rectangular coordinates: high-mid-low (also referred to as close-open) and front-back. The articulation of vowels is thus characterized in terms that are totally different from those of the consonants, and this seems rather unnatural.

The IPA system for characterizing vowels, which was originally introduced by A. M. Bell (1867), has recently been subjected to devastating criticism by

S. Wood in a number of publications collected in S. Wood (1982). Wood points out that the highest point of the tongue arch in the lax [I] is lower than in the tense vowel [e]. Although this fact has been known since the beginning of the twentieth century, textbooks of phonetics almost universally teach that [I] is a high vowel and [e] a mid vowel. This practice of the textbooks is due to the feeling on the part of most linguists that [i] and [I] belong together, regardless of the results of their measurements. The feeling on the part of linguists is, of course, not just an instance of mindless conservatism. Though never mentioned by Wood, the proposition that [i] and [I] belong together is powerfully supported by evidence from the phonological rules of the languages of the world, and in phonetics, which is the study of the sounds of language, the evidence from phonological rules can never be disregarded. The fact that the evidence from the phonological rules is not compatible with measurements of the position of the highest point of the tongue arch, raises questions about the relevance of the measurements. It suggests that we explore whether the tongue arch model cannot be replaced by one that is more appropriate in that it allows not only for the characterization of the different articulatory configurations, but is also compatible with the data from the phonological rules of the languages of the world.

I want to propose now that the model sketched at the beginning of this section is such an alternative. It was suggested there that the process of speech production consists in moving an articulator from one position to another, where by articulator is meant a recognized anatomical entity such as the lower lip, the body of the tongue, or the vocal cords, but not an entity defined purely ad hoc such as the highest point of the tongue arch which varies constantly in the course of an utterance. In the production of vowels the most important articulator is the body of tongue whose position is controlled by the extrinsic muscles of the tongue. I shall argue below that in producing speech these muscles are under the control of the three binary features [high], [low], and [back]. I shall assume that as suggested in Chomsky and Halle (1968, hereinafter SPE, pp. 304–5), the feature specification [+ high] is an instruction to raise the body of the tongue towards the roof of the mouth; the specification [+ low] is an instruction to lower the tongue body to a level below the uvula, while the specification [+ back] is an instruction to retract the tongue toward the rear wall of the pharynx. Since [+ high] and [+ low] are contradictory instructions we shall postulate that there can be no sounds that are [+ high, + low]. The three features thus define the six vowel articulations in (6), where other properties such as rounding, tensing, etc., are provisionally disregarded.

(6)	u	o	a	i	e	æ
high	+	–	–	+	–	–
low	–	–	+	–	–	+
back	+	+	+	–	–	–
	velar	uvular	pharyngal	palatal		

It was further proposed in SPE that the feature combinations in (6) were involved not only in vowels but also in consonants. Not all six combinations yield consonantal articulations, because in consonants the active articulator must make at least partial contact with some part of the stationary, passive portion of vocal tract; i.e., the rear wall of the pharynx or the roof of the mouth including the front teeth and upper lip. Such contact can be made by the body of the tongue only if it is raised ([+ high]) or retracted ([+ back]). Because of this, only four of the six vowel configurations in (6) have consonantal counterparts.

Virtually the same conclusions as those above about the articulation of vowels were arrived at by Wood (1982) without reference to SPE or the extensive literature elicited by it. Wood studied 38 sets of x-rayed vowel articulations from 16 different languages collected from the literature and from x-ray motion pictures produced by his group. He concluded that "there are four different places where the vocal tract is narrowly constricted by the tongue for vowels – along the hard palate, along the soft palate, in the upper pharynx and in the lower pharynx." (pp. 42–3). He observes that his findings confirm the important theoretical result of Stevens (1972) about the quantal nature of articulations of speech sounds; specifically, "Stevens' hypothesis that we seek to constrict the vocal tract for vowels at those places where F1 and F2 are least sensitive to variability of constriction location." (ibid).

The characterization of consonantal articulations which, as noted above, is one of the basic traits of the IPA system is not readily compatible with the model sketched above since the IPA model disregards completely the active articulator and focusses exclusively on the location of the constriction. We have already discussed our treatment of the consonants formed by the tongue body. The two other active articulators that are involved in the production of consonant types are the lower lip and the tongue blade. We follow SPE here and postulate that the tongue blade is controlled by the feature [coronal]. With regard to the lower lip, we shall deviate from SPE and postulate that it is controlled by a special feature [labial].

Consonantal occlusions are thus produced by three distinct active articulators: the lower lip, the front part of the tongue, and the tongue body. Since the position of each of these three articulators is independent of the other two it should be possible to produce consonants with more than one occlusion. Since there are three active articulators and since a given articulator can be at exactly one point at a given time there should exist three types of consonants with double occlusion and a single type of consonant with triple occlusion. As shown in (7) all double occlusion consonants are attested, but I have been unable to find an example of a consonant with triple occlusion.

(7)	labio-velar	[k͡p]	Yoruba	[ak͡pa]	"arm"
	labio-coronal	[p͡t]	Margi	[p͡təl]	"chief"
	corono-velar	[ɿ] (click)	Zulu	[ɿaɿa]	"climb"
	labio-corono-velar		unattested		

The framework proposed here implies that the configurations with multiple closures in (7) are the only ones possible. By contrast the IPA framework with its point of articulation concept makes no assertion regarding sounds with double occlusion, implying that the absence of reports in the literature concerning sounds with multiple occlusions other than those in (7) is due to a fortuitous gap in our knowledge likely to be filled in by future research. The facts in (7), therefore, constitute significant evidence in favor of the proposals that have been sketched above and against the point of articulation concept of the IPA system and other phonetic frameworks.

3. In this concluding section of my paper, I examine the muscular activity that underlies the movements of the different active articulators in speaking. My ideas on this topic have been influenced by C. R. Gallistel's (1980) *The Organization of Action*, a book which I strongly recommend to all students of articulatory phonetics.

In characterizing movements of structures connected by joints, Gallistel observes that "a stimulus that excites a muscle on one side of a joint invariably inhibits excitation of the antagonistic muscle on the other side of the joint, and vice versa" (p. 58). While the active articulators are not connected to other parts of the vocal tract by joints, their movements are in many cases controlled by paired sets of agonistic and antagonistic muscles. Thus, for example, the lowering and raising of the velum under the control of the distinctive feature [nasal] is implemented by the *palatopharyngeus* and *palatoglossus*, which together lower the velum, and the *tensor veli palatini* and *levator veli palatini*, which raise it. When the *tensor* and *levator* are excited and the *palatoglossus* and *palatopharyngeus* are inhibited, the velum is raised and no air can flow through the nasal cavities. When the former two muscles are inhibited and the latter two excited, the velum is lowered and air flows freely through the nasal cavities producing a specific acoustic effect which is referred to as nasalization. In (8) I have drawn a block diagram of a model of neuromuscular control of this feature.

(8)

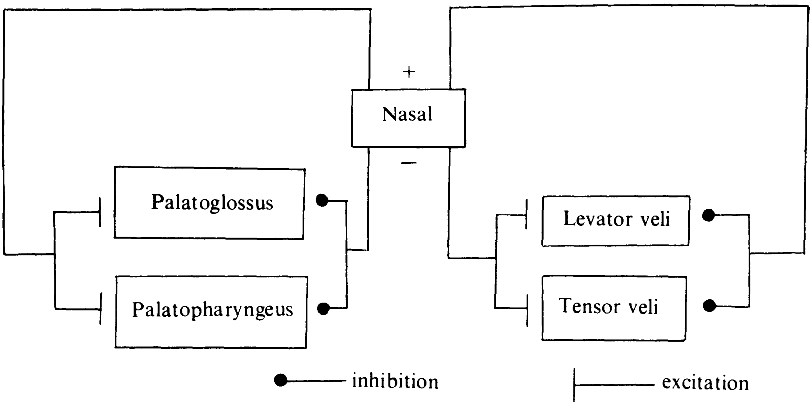

The feature [coronal] is apparently controlled in much the same fashion as nasality [i.e., by a pair of agonist–antagonist muscles]. To produce a [+ coronal] sound the tongue blade must be raised; to produce a [− coronal] the tongue blade must be lowered. Blade raising is implemented by contracting the *superior longitudinal muscles* of the tongue and relaxing the *inferior longitudinal muscles*; whereas blade lowering is produced by relaxing the *superior* pair of muscles and contracting the *inferior* pair.

The situation is more complicated in the case of features involving muscles that are not positioned so as to form agonist–antagonist pairs. Typical examples of such features are high, low, back, i.e., the main features controlling the position of the tongue body. The muscles of interest here are the extrinsic muscles of the tongue: the *anterior genioglossus* (AGG) the *posterior genio-*

(9)a.

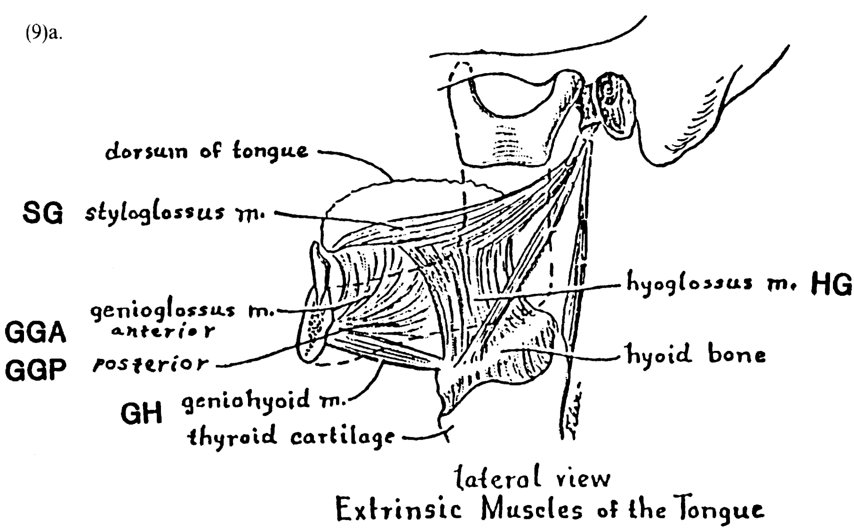

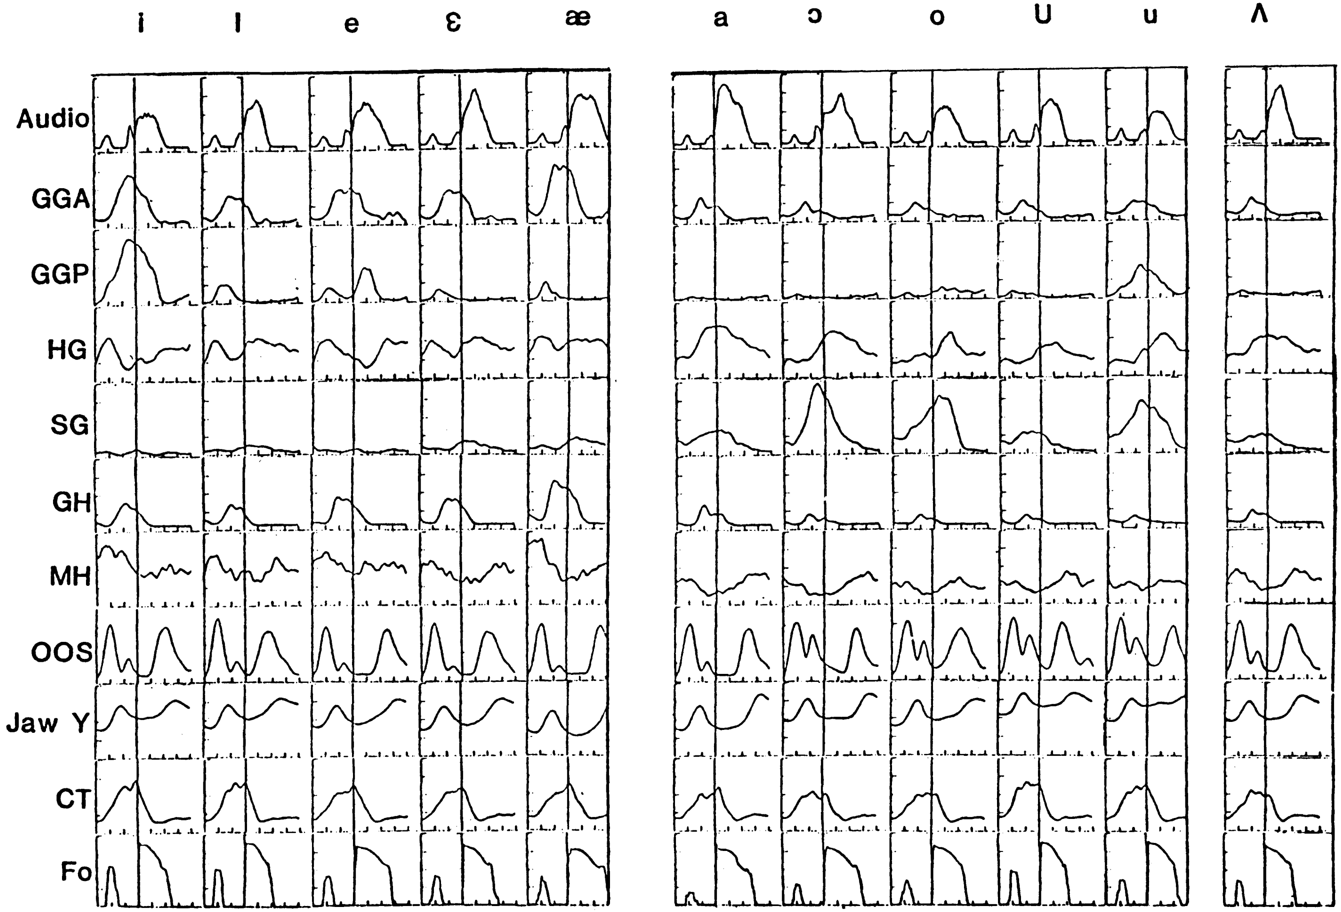

(9)b. EMG Potentials were simultaneously recorded from the muscles listed while a single speaker produced ten repetitions of the eleven vowels occurring in the [ə pVp] environment. Vertical lines represent the acoustic onset associated with the vowel. Units on the abscissa indicate 100 ms. intervals. (Reproduced with permission from Alfonso, Honda, Baer and Harris (1982)).

glossus (PGG) the *geniohyoid* (GH), *hyoglossus* (HG) and *styloglossus* (SG).

To aid in the following discussion I have reproduced, with permission, from Alfonso, Honda, Baer and Harris (1982), the anatomical drawing (9a) and the EMG recordings (9b) on pages 100 and 101 respectively.

Table (10), below, presents the postulated excitation patterns of the extrinsic tongue muscles in the articulation of the six major vowel types whose feature composition is given on the left. The table is based on the EMG data in (9b) supplemented by information from other sources. Special comment is required with regard to the following: (A) I have not included in Table (10) the behavior of the *mylohoid* since its electrical activity appears not to be correlated with the three features under discussion here. (B) In spite of substantial electrical activity in the *posterior genioglossus* in the articulation of the tense [e] and [o] shown in (9b) this muscle has been supplied with a minus mark for these vowels in table (10) indicating absence of excitation. The electrical activity shown in (9b) reflects the fact that in American English tense [e] and [o] are normally diphthongized. This fact was also noticed by Smith (1971) who remarked that "this activity corresponds to the ... high tongue position for the glide portion of the nucleus." (p. 30). (C) The data on electrical activity in the *hyoglossus* during the articulation of the front vowels is hard to interpret. T. Baer (personal communication) has kindly informed me that this is due to the effects of the bilabial plosives which precede and follow the vowel in the test utterances (9b) and that when account is taken of these contextual effects the data show a maximum of activity in the *hyoglossus* for the low vowel [æ]. (D) The EMG record in (9b) shows considerable activity in the *styloglossus* for the nonhigh back vowel [ɔ]. Since activation of the *styloglossus* exerts an upward pull on the tongue body, the recorded activity is incompatible with the low tongue position characteristic of [ɔ]. I have been informed by T. Baer (personal communication) that the subject in the EMG experiment pronounced the vowel [ɔ] in an atypical manner. I have, therefore, left this activity out of account in constructing table (10).

(10)	high	low	back	AGG	PGG	GH	HG	SG
[uɨ]	+	–	+	–	+	–	–	+
[üi]	+	–	–	+	+	+	–	–
[oγ]	–	–	+	–	–	–	–	–
[oe]	–	–	–	+	–	+	–	–
[ɔa]	–	+	+	–	–	–	+	–
[ææ]	–	+	–	+	–	+	+	–

Our next task is to outline neurologically plausible circuitry that would be capable of eliciting the excitation patterns in the five muscles given in the right part of (10) assuming that these are driven by distinctive feature centers with outputs shown on the left part of table (10). A diagram of such circuitry is given in (11).

(11)

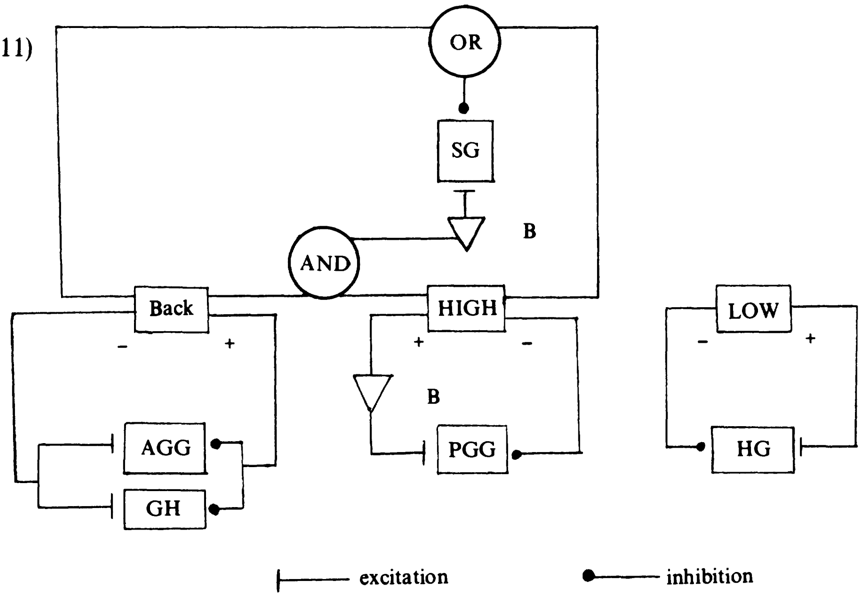

The circuits that make up the lower half of (11) are of a very simple form: the muscles represented on the bottom of (11) are excited for one output of a given feature center, and inhibited for the opposite output of the feature center. Thus, the *anterior genioglossus* and *geniohyoid* are both excited for [– back] sounds, and are both inhibited for [+ back] sounds. Similarly the *posterior genioglossus* is excited for [+ high] sounds and inhibited for [– high] sounds, whereas the *hyoglossus* is excited for [+ low] sounds and inhibited for [– low] sounds.

The situation is considerably more complicated for the *styloglosus*, which, as indicated in (10) is excited for [+ back, + high] sounds and inhibited for sounds that are either [– back] or [– high]. As shown in the upper left hand portion of (11), in order to express the fact that the *styloglossus* is excited only when the two features have positive values, we have connected these two outputs to a component labelled AND. This component represents what is known in circuit theory as an 'and-gate'; i.e., a circuit element that transmits current only if both of its inputs are excited. The output of the 'and-gate' is connected to the *styloglossus* so as the excite it. The outputs [– high] and [– back] are connected to a component labelled OR. This component, termed an 'or-gate' in circuit theory, transmits a current when either of its two inputs is excited; when this happens the *styloglossus* is inhibited. Both 'and-gates' and 'or-gates' are widely represented in the human nervous system where they are referred to by the term 'synapse'.

We have not included in table (10) and diagram (11) the activity of the *pharyngeal constrictors* which no doubt play a role in the articulation of nonhigh back vowels. The connections of the *pharyngeal constrictors* with the features [+ high] and [+ back] are of the same form as those shown in (11) for the *styloglossus.*

As was noted above, the basic articulatory difference between consonants and vowels is that in producing consonants the moving articulator – the tongue body in the cases under discussion here – makes contact with the opposite wall of the vocal tract, whereas in vowels a significant distance is maintained between the moving articulator and the stationary part of the vocal tract opposite it. The fact that the activation of the same sets of muscles produces both vowels and consonants immediately raises the question as to the mechanism that differentiates the articulation of a vowel from that of its consonantal cognate. From a purely articulatory point of view the difference seems almost trivial. Since the moving articulator makes contact with the opposite wall in consonants but not in vowels, the muscles moving the articulator in this direction must contract less in vowels than in consonants. The extent to which muscles contract can, of course, be regulated, and we have included in (11) special triangular boxes labelled B whose function it is to regulate the degree of contraction of the *styloglossus* and *posterior genioglossus.* We shall assume that in the neutral case the muscles of interest contract maximally when excited, and that in order to reduce the contraction of these muscles the elements B must be activated. These elements must, therefore, be connected to a higher level center controlling the feature [consonantal], in fact, they must be connected only to the [– consonantal] output. Since the same distinction between vowel and consonant articulations holds also for the constrictions produced by the rear of the tongue body in the pharynx, the lower lip and the tongue blade, parallel connections from the [consonantal] center must be postulated to run to the circuit exciting the *pharyngeal constrictors*, the *orbicularis oris* muscle in the lower lip and the *superior longitudinal muscles* of the tongue.

An implication of the model (11) is that when a consonant with double or triple constrictions is formed, all constrictions will be actualized as closures or occlusions. This is incorrect. Consonants with double occlusion are considerably less common than consonants with one occlusion and a second moderate constriction. Examples are the velarized and palatalized consonants of Russian and the rounded consonants of Kashmiri (Morgenstierne 1938) and Dungan Chinese (Trubetzkoy 1939). There are, of course, very obvious ways in which the proposed model could be modified to take these facts into account but since there are so few neuro-anatomical facts at my disposal to constrain the imagination, I desist from further speculation.

The rudimentary nature of the circuitry sketched in the last few paragraphs makes it likely that what has been proposed here will have to be extensively modified. Nonetheless, the fact that the rudimentary means employed above are able to account for behavior of considerable complexity should not be overlooked. It suggests that we may well be on the right road, even if still far from our goal.

Acknowledgements

In preparing this paper I have benefitted from discussions with S. Bromberger, G. Fant, N. Geschwind, J. Lettvin, M. Liberman, D. Osherson, J. Perkell, W. Poser, and K. N. Stevens, none of whom shares all of the views expressed. I am especially grateful to T. Baer of the Haskins Laboratories for granting permission to reproduce figures (9a) and (9b) which originally appeared in Alfonso, Honda, Baer and Harris (1982). I dedicate this study to the memory of Beatrice Hall, at whose invitation I presented an earlier version of this paper at the Conference on Phonological Distinctive Features at the State University of New York in Stony Brook on June 14, 1982.

References

Alfonso, P. J., K. Honda, T. Baer, and K. S. Harris: 1982, 'Multi-Channel Study of Tongue EMG during Vowel Production,' Paper presented at the 103rd Meeting of the Acoustical Society of America, April 25–30, 1982.

Bell, A. M.: 1867, *Visible Speech*, London.

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

Gallistel, C. R.: 1980, *The Organisation of Action*, Erlbaum, Hillsdale, NJ.

Jakobson, R., C. G. M. Fant and M. Halle: 1952, *Preliminaries to Speech Analysis*, MIT Acoustics Laboratory Technical Report # 13.

Morgenstierne, G.: 1938, *Indo-Iranian Frontier Languages*, Oslo.

Sapir, E.: 1925, 'Sound Patterns in Language,' *Language* **1**, 37–51.

Smith, T. S: 1971, *A Phonetic Study of the Function of the Extrinsic Tongue Muscles*, UCLA Working Papers in Phonetics # 18.

Stevens, K. N.: 1972, 'The Quantal Nature of Speech: Evidence from Articulatory-Acoustic Data,' in P. B. Denes and E. E. David, Jr. (eds.), *Human Communication, a Unified View*, 00000 pp. 51–66.

Stevens, K. and M. Halle: 1971, 'A Note on Laryngeal Features,' *MIT-RLE Quarterly Progress Report* # 101, pp. 198–213.

Trubetzkoy, N. S.: 1939, 'Aus meiner phonologischen Kartothek,' *Travaux du Cercle linguistique de Prague* **8**, 22–26.

Wood, S.: 1982, *X-Ray and Model Studies of Vowel Articulation*, Working Paper 23, Lund University, Department of Linguistics.

Received 9 October 1982

20D-219, MIT, Cambridge, MA 02139