

The Aims of Autosegmental Phonology

John Goldsmith

In this paper, I shall discuss several aspects of an approach to generative phonology known as "autosegmental phonology." Autosegmental phonology shares with more traditional generative phonology a commitment to a formal account of phonological processes, and also a commitment to developing a formalism in which the common, expected, or "natural" developments in a language are represented in a formally simple way. What distinguishes autosegmental phonology from the *Sound Pattern of English* type of generative phonology is, first, the development of a multi-linear phonological analysis in which different features may be placed on separate tiers, and in which the various tiers are organized by "association lines" and a Well-Formedness Condition; and, second, analysis of phonological phenomena less in terms of feature-changing rules as such, and more in terms of rules that delete and *re-organize* the various autosegments, through the readjustment of the association lines.

The general program for research has gone in a number of directions, and I shall emphasize in this paper a few of these areas. In the first part (§1), an informal account of the original motivations for the autosegmental treatment of suprasegmentals is presented. It deals with issues discussed in more detail in Goldsmith (1976b). In the second part (§2), I shall illustrate the power of an autosegmental approach in dealing with an intricate corner of the tonal system of Igbo. In the third part (§3), some of the interesting and exciting suggestions made by Clements, McCarthy, Haraguchi, and others regarding the treatment of other suprasegmental phenomena in an autosegmental way are discussed, and in the fourth section, (§4), I sketch some of the ways the general approach can deal with typically segmental assimilation phenomena. I try,

furthermore, to motivate the existence of autosegmental levels in languages classically considered to contain suprasegmentals, and even those which appear more traditional. Finally, in §5, I remark upon the view of phonology I see lying behind the autosegmental approach, and how this view may be seen to differ from that lying behind various others of the theories presented at this conference.

1. American phonology has shown—not always, but at times—a peculiar rigidity in its conception of the "shape" of phonological representations. One of the aims of autosegmental phonology is to investigate the consequences of having structures in phonology more complex, or more articulated, than a simply linear string of segments. I shall discuss some of the tools offered, and some of the problems they can deal with.

The aim, then, of autosegmental phonology is to deal with the consequences for generative phonology of multi-linear phonological analysis and representation. That is, we let go of the assumption that phonological and phonetic representations consist of a single string, or concatenation, of segments. Instead, we set up underlying and surface forms consisting of parallel strings of segments arranged in two or more tiers. Features are distributed over the various tiers in the sense that no feature may appear on more than one tier. This somewhat richer phonological representation then serves as the basis for a more enlightening, and ultimately simpler, formal phonological analysis.

Part of this idea, certainly, is not new to linguistics. The Firthians are quite serious in their attention to prosodies, those units in the phonological analysis above and beyond the sequential "phonematic units," or segments, as we might call them.

It seems to me that in their attention to such processes as vowel harmony and the complexity of tonal systems, and in their rejection of a strictly linear view of phonological representations, the Firthians are right on the mark. So too were Bernard Bloch (1948) and Zellig Harris (1944) in their discussions of the more complex possibilities of segmentation in their articles in *Language*. Apart from this major point of agreement with Firthian or prosodic analysis, the autosegmental approach differs on at least two other important points. The first is the centrality of the distinction maintained in the autosegmental approach among the notions of *segment*, *feature*, *rule* and *association*. This four-way distinction does not correspond in any natural way to the concepts and constructs found in Firthian phonology. The notion of *feature* in the generative framework corresponds in the Firthian framework to

the notion *prosody*—an element which furthermore plays the role corresponding to phonological rule and to “true suprasegmentals” (in our terms, autosegmentalized features).

But the only types of prosodies that correspond to autosegmental analyses are those in which the prosodies are phonetically homogeneous—which, as a review of the literature will show, is a small percentage of the extant prosodies. (For further discussion, see Goldsmith [in preparation a].)

The second major difference between Firthian and autosegmental analysis lies in the emphasis, in the latter, placed upon rules and their interaction with the Well-Formedness Condition. We shall see an example of this in section 2 below.

If there were several precedents for breaking away from a strictly linear representation, how this was to be done within the framework of generative grammar has been problematic.¹ The clearest generative position on this is found in Chomsky (1955:29):

(T)hough the representations that we construct on any linguistic level are unidimensional, we have not required that left-to-right order of representation correspond directly to temporal order in the represented utterance. . . . By accepting a linear system of representation, we rule out the possibility of certain kinds of discontinuity. If more general kinds of discontinuity than we can handle occur in language, a more general theory of representation will be necessary.

FN: In this study, suprasegmental features (pitch, stress, juncture) have not been seriously considered. Ultimately, of course, these phenomena must be incorporated into any full syntactic theory, and it may be that this extension still requires a more elaborate system of representation.²

A significant step was taken by William Leben (1973) and Edwin Williams (1976). Leben's and Williams' suggestion did not change the picture of phonetic representation as a linear sequence of segments—thus *not* reverting to the Bloch-Harris position—but they did suggest that the representation underlying some particular sequence of “phonetic” segments could consist of two separate and quite independent sequences of phoneme-type segments, one sequence providing the tonal information, the other sequence containing everything else. Naturally, to have a surface or phonetic representation that was entirely one-dimensional—that is, a simple sequence of segments—they each provided for “feature mapping,” which joined two separate sequences of segments into one single linear sequence.

The autosegmental approach arose out of certain inadequacies that



were brought to light explicitly and implicitly by Williams' and Leben's work. The most glaring problem was the nature of “contour-toned” vowels—that is, vowels whose surface tone is rising or falling, a situation that can often be shown to be the result of a concatenation of Low and High level tones. How can a single segment bear or carry two tonal specifications in sequence? This was the first problem to resolve.

Intimately related to this was the necessity of explaining the relationship of the left-to-right ordering *inside* the segment (it seemed) and the left-to-right order of segments themselves. A tonal assimilation occasioned by a Low tone that affects a High tone to its right, for example, will turn that High tone into a Rising tone—never a Falling tone. (1) represents this as a strictly linear system would; but, strictly speaking, this way of representing a contour tone makes no sense. We may look at the solution to this problem as a “boundary condition” on the solution to the first question, the representation of contour tones.

$$(1) \begin{bmatrix} V \\ - \text{high} \end{bmatrix} \text{Co} \begin{bmatrix} V \\ + \text{high} \end{bmatrix} \longrightarrow \begin{bmatrix} V \\ - \text{high} \end{bmatrix} \text{Co} \begin{bmatrix} V \\ - \text{high}/+ \text{high} \end{bmatrix}$$

Yet another problem in the theory of tonal representations (cf. Goldsmith 1975b) is what I have called “stability,” and, rather interestingly, its import is apparent only within a generative framework. “Stability” refers to the resistance of the tonal features of a vowel to deletion, even when the vowel that bore the tonal features is deleted or desyllabified.

The resolution of these problems was the introduction of several parallel strings of segments in the phonological *and* phonetic representations, enriched in a significant way by “association lines”—lines which, at the phonetic level, indicate the “co-registration” of the different tiers of segments. A language in which tonal features are autosegmentalized will represent a bisyllabic word with a tone pattern High - Low as in (2); a rule that would cause the second syllable to assimilate in tone would modify (2), change the relative domains of the tonal and non-tonal segments, and produce, in turn, (3).



(3) represents, it should be clear, a bisyllabic word whose tone pattern is High-Falling.

This autosegmental representation automatically solves our first three problems. The contour-toned vowels are, as (3) suggests, those in which the association is not purely one-to-one; tone-spreading rules that give rise to certain kinds of contour tones are the effect of a rule that adds an association line. We see such a change occurring between (2) and (3). We might note, furthermore, that the only rules which have the simple geometric property of adding an association line do, in fact, produce "natural"—that is, occurring—types of assimilation. Thus simplicity in the formalism does seem to match naturalness in the phonological system. Thirdly, stability of tonal features is the natural consequence of a rule deleting a segment on the tier consisting of non-tonal features (the upper tier, that is, in (2) or (3)). Now, in a language in which such an autosegmental representation is appropriate at the underlying level, two possibilities are open. Either the lexical entry for a word consists of a pair of strings of segments—much as in (2), but without the association lines—or else each tier might constitute a separate lexical entry. Thus two "morphemes" may in fact constitute the single underlying phonological level—one on the tonological tier, one on the non-tonological tier.

This kind of independent-minded behavior is familiar to linguists under the guise of "floating tones"—morphologically identifiable elements (a High tone, for example) whose position with respect to the syllables in the sentence is determined by the syntactic structure. That is to say, it can frequently be shown in tone languages that the tone pattern of a sentence is composed of the tonal patterns of the individual words in the sentence plus an additional tone, a morpheme which contributes only to the tonal pattern and not to the number or make-up of the syllables.

Conversely, it is often found in such tone languages that certain syllables are underlyingly devoid of tone, and receive a tone only through the influence of their neighboring tone-lending syllables. In general, then, a language with N level tones has the potential of $N + 1$ types of morphemes, tonally speaking. In a language with N level tones, a syllable may underlyingly be any of the N tones (or some concatenation of these) or it may be not marked for any tone. An interesting example of just such a situation is described in Wilkinson (1976).

Both of the cases described in the last two paragraphs, it should be clear, are examples of an underlying "mismatch" between the number of tonal segments and the number of syllabic segments—a difficult concept to deal with within a more familiar generative approach.

A crucial problem that arises now for any generative approach of the sort we have been considering is this: what is the technical device that relates the two tiers of segments? So far we have made allusion to association lines and their function; but at this point, we should make clear their origin. We may assume that all association lines (on segments other than boundaries) are introduced at some point in the derivation after the underlying representation. This amounts to, in effect, saying that there are no association lines in the stored or underlying representation. And this, in turn, implies that all underlying representations are linear.³


If underlying representations, then, contain no association lines, these association lines are first introduced by rule. The range of possible association rules is still undetermined, but in African tone languages, for instance, it is clear that a major rule in this respect is mapping one tone per syllable, starting from the left (see Edmundson and Bendor-Samuel (1976), Williams (1976), Leben (1973, 1977a), Goldsmith (1976b); for a far-reaching study of the interaction of accent and association rules, see Haraguchi (1976); on Sanskrit, May and Goldsmith (1975); on the interaction with the cycle, Elimelach (1976)).

A rather powerful convention—perhaps too powerful, actually—on derived forms was suggested in Goldsmith (1976a, b) and elsewhere, a convention called "the Well-Formedness Condition" (WFC). This convention has the effect of adding or deleting association lines at any point throughout the derivation. Stated informally (see Goldsmith [1976b] for a more unified statement), it reads:

Well-Formedness Condition

1. Each vowel must be associated with (at least) one toneme.
2. Each toneme must be associated with (at least) one vowel.
3. No association lines may cross.

One consequence of the WFC can be seen in (4). When a toneless suffix appears on a stem, the tone associated with the stem automatically has its domain extended to the suffixes as well (cf. (5)).




Such a Condition, it can be seen, would have the effect of placing floating tones on some vowel or other—of "docking" the floating tone, we might say. Certain additional principles—some language-specific, some


universal—may be necessary to effect the results required by the WFC.

In general, however, the autosegmental approach is not meant to deal only with tone; other languages (for example, Guarani) approach nasalization in an autosegmental way, and other features (other than the major class features)⁴ could be dealt with autosegmentally.⁵

2. Let us look at a typical autosegmental analysis, based on data from Igbo. A crucial fact about tone in Igbo is that there are three contrasting tones: High, Low, and Drop (or Mid). High, Low, and Drop contrast after a High or Drop, but the three-way contrast is reduced to a two-way contrast after a Low tone; in particular, High and Drop are indistinguishable phonetically when they follow a Low tone. Furthermore, a Drop tone is distinguished from a High tone in that the Drop tone is produced a pitch-step lower than a preceding Drop or High tone. A High tone (N.B.) will be produced at the *same* pitch as a preceding Drop or High. Schematically we can summarize this as follows:




Williams (1976) argued a certain tonal mutation in nouns (which we shall call "Object Mutation") in Igbo can be used as a diagnostic for the presence of a (sometimes floating) H tone preceding these mutated nouns. In Goldsmith (1976b) the analysis is deepened, and the floating tone is shown to account for an interesting set of alternations. Two aspects are crucially "autosegmental": first, positing a tonal deletion rule feeding the Well-Formedness Condition, and second, derived structures as in (7). Crucially, in one case ((14) below), the domain of a Drop tone is spread over two syllables, yielding a configuration as in (7).




In a context after a High or Drop tone, it has the same phonetic rendering as a sequence of Drop+High syllables but it does *not* have the same phonetic realization as a sequence of Drop + Drop syllables (cf. (6)). Thus in (14), the spread of the domain of the Drop tone *must* be viewed

autosegmentally, rather than as a segmental assimilation of tonal features.


The relevant facts are these. (i) Each verb stem is lexically marked as being underlyingly Low-toned or non-Low. In the II Root pattern, the pattern we shall consider, there appears a prefix /a/ (or its harmonic alternant). The tone of the prefix is underlyingly the opposite of the tone of the stem; thus, with Low-toned verbs, the prefix is High-toned, the stem is Low-toned, and the suffixes to the verb are all High-toned—this last being the effect of the "floating" (affixal) H that Williams argues for. Consider thus (8).



If no suffixes on the verb occur, the Stem L-tone and the suffixal H-tone converge on a single syllable, as in (9). A rule of Low-High simplification is an easily motivated rule in Igbo, given transparent alternations elsewhere in the language.






(10) Low-High Simplification




(The tone on the object *ezi* in (9) actually changes to D due to the role of Object Mutation mentioned above, but that is not directly relevant here.) Thus the surface form of (9) is (11).

(11) úñù áfù èzí

(ii) When the subject NP ends in a High tone, the tone on the verb and prefix is altered. In particular, the tone of the prefix is deleted; thus the vowel of the prefix associates with the tone of the stem, through the Well-Formedness Condition. Thus we find—totally exceptionlessly—contrasts like those in (12) and (13), where the relevant independent variable is the final tone of the subject NP.




(iii) The evidence that this schematic way—i.e., an autosegmental deletion—is the correct way to view the alternation is quite strong. First, we may observe that the shift of the tone of the prefix from High to Low (as in (12)–(13)), triggered by the High tone on the subject NP, is matched by a shift of the prefix tone from High to Drop when there are no suffixes (and concurrently the stem *appears* to shift from Drop to High). This is illustrated in (15); compare it with (9) and (11), where the prefix tone does not delete. This shows clearly that although the final H of the subject NP *triggers* the tonal change in the verb prefix, it does not determine what the new tone will be. In fact, the tone of the prefix will be simply the tone of the stem, as determined by the Well-Formedness Condition.



(The tonal accent marks over the syllables are the transcriptions given by Green and Igwe (1963), i.e., the best segmental interpretation of the tones.)

(iv) Actually, the case for viewing the domain of the tone of the verb stem as being the prefix as well as the stem in certain cases is even

stronger. There are two further arguments. The first is based on the fact that the Low-High combination on a single vowel (as in (9) and (15)) simplifies to Low if the following tone is High (otherwise it simplifies to Drop, of course). Again, this rule—let us call it “Floating H Deletion” (cf. Goldsmith [1976b])—is clearly motivated as well by the tonal properties of the nominal compound system; I will not review that material here. The process is illustrated in (16)–(17). However, the effect of this rule is to create a rather striking contrast between forms like those in (15) where the object NP begins with a Low tone, and those in (18), where the object NP begins with a High tone; observe the effect on the prefix tone in (18), precisely as expected.




(v) We can find a second, rather different, reason to support this analysis of the tonal alternations of the verb-prefix as expressed in (14). All the examples we have viewed so far have involved inherently Low-toned verbs. In this same construction, High-toned verb stems may appear, and, as (14) Prefix Tone Deletion predicts, the tone of the prefix again depends on the final tone of the Subject NP. The relevant facts are summarized in (19).


- (19) a. Subject ends in Low tone:
 $\begin{array}{c} \acute{e}z\grave{e} \quad \acute{a}c\acute{í}\acute{a} \quad \acute{a}n\grave{u} \\ \text{subject} \quad \text{verb} \quad \text{object} \end{array}$
 b. $\begin{array}{c} \grave{a}d\acute{h}\acute{a} \quad \acute{a}c\acute{í}\acute{a} \quad \acute{a}k\acute{w}\acute{a} \\ \text{subject} \quad \text{verb} \quad \text{object} \end{array}$

(We have the same verb in both (19 a and b).)


Our analysis so far has led us to the conclusion that the tone pattern on the prefix is closer to its underlying form when the subject NP ends with a Low tone, and that the underlying tone on the prefix is deleted when the subject ends in a High tone; in this latter case, the prefix vowel reassociates with the vowel of a verb-stem by the Well-Formedness Condition. All that is left to observe is that this is precisely what we do in fact observe in (19) as well, on the assumption that the underlying tone pattern for High-toned verbs in this construction is as in (19a)—that is, Low tone on the prefix, Drop tone on the stem. (20) illustrates how (19b) derives from a structure like (19a) through the operation of (14) Prefix Tone Deletion.



3. The study of other systems traditionally called "suprasegmental" has proceeded in the autosegmental framework. Haraguchi (1976) (and similarly Goldsmith [1975a] though not based on a comparably broad base of detail), show how the Japanese system of pitch and accent can be assimilated to a system of analysis originally suited to fit the intricacies of African tone languages. Ostler (in press) has recently proposed some simplifications of Haraguchi's survey in an elegant fashion. G. N. Clements has developed in several papers an autosegmental approach to traditional vowel harmony systems (Clements [1976a, b, c]). Such a treatment analyzes the harmonic feature—tongue-body position, tongue-root position, and so forth—as a feature on a separate autosegmental tier. The autosegment associates across the entire domain of the word by the Well-Formedness Condition, as in (21) (an example from Akan, from Clements [(1976a)]).



The vowel harmony process in Akan, as (21) suggests, is harmony for the feature "advanced tongue root." This approach permits, and requires, a unified statement of harmony within the stem as well as across morpheme boundaries. We thus avoid an arbitrary separation of vowel harmony into intra- and extra-stem processes. By the same token, we do not have what Kiparsky (1968a) called "a phonological use of morphological features"; the harmony autosegment is indeed a purely phonological item. More appealing yet, we can provide a natural account for words with more than one harmony domain, a capability not even available, as Clements notes, to advocates of "phonological use of morphological features," that is, to those advocating assigning a stem *as a whole* to one harmony class or the other. Clements (1976b) discusses an example from Akan, where /a/ is quite generally an "opaque" element, that is to say, it blocks vowel harmony from spreading across it. A word like *bisa* 'to ask' produces tense vowel harmony among its prefixes, and lax harmony in its suffixes. This is illustrated in (22).



The /a/ is assigned to the [−Advanced Tongue Root] segment either underlyingly (cf. note 3) or as the first association rule, and the Well-Formedness Condition adds the other association lines in (22) (cf. Goldsmith [1976b], Clements [1976b], e.g., for discussion of the principle giving precedence in association to unassociated segments).

The approach sketched here has been developed in a series of papers by Clements. The results described there, as well as the discussion of nasalization in Guarani in Goldsmith (1976b), suggest strongly that the formal properties of tonal suprasegmentals are matched in many respects by vowel harmony, though many of the details remain to be worked out.

One aspect of vowel harmony which I made a passing reference to above seemed unusual when the attempt was first made to apply autosegmental principles to vowel harmony. That is, in the vowel harmony system most discussed in the literature of generative phonology (including Akan, Hungarian, Turkish, Igbo, Nez Percé), a single harmonic domain was the rule and multiple domains the exception. This did not parallel the case in tone languages, where tone *melodies* consisting of several segments in sequence often provide crucial information about tense, mood, and other grammatical factors. Could a principled reason be found for this difference, or were the two systems in fact more similar than appeared at first?

John McCarthy, in a very interesting paper (1977) on Arabic vocalism, has argued, in effect, for the latter position. He reviews the well-known facts of Classical Arabic verbal morphology, and capitalizes on the qualitative difference found between, on the one hand, the *kind* of alternation occasioned by inflection—limited to varying the vowel patterns in a rigidly fixed syllable pattern—and, on the other hand, the derivational processes that form the often idiosyncratic conjugation or “measures,” processes that drastically alter the syllable structure as well as the *number* of consonants involved. Despite the vagaries of the derivational morphological process, however, the inflectional morphology remains constant across different forms (“measures”).

Thus, in Classical Arabic the vowel pattern for the Perfect Active is /a/—but in this case, there may be two or three syllables in the root, depending on the measure (i.e., depending on the effects of the derivational morphology). The vowel in *each* syllable will be /a/, however, in the Perfect Active. In the Perfect Passive, the final vowel will be /i/, preceded by one or two occurrences of /u/, depending on the syllable structure determined, again, by the derivational morphology. In this latter case, then, we find a vocalic melody pattern of /u i/, where /u/ has precedence in spreading. See (23), for example.

(23) $\begin{array}{ccccccc} t & V & k & V & t & t & V & b \\ & \diagdown & & \diagup & & & | & \\ & u & & i & & & i & \end{array} = \text{tukuttib}$ Form determined by derivational morphology

McCarthy considers a range of similar examples, and points out the close parallel in both the form *and* the function found in the Arabic vowel system and the Tiv tonal system described in Goldsmith (1976a, b). Although the data is complex, the parallel he draws is quite convincing, as in his analysis in general, pointing to vocalic systems like that of Arabic as being precisely the logical link between harmony systems of the sort analyzed by Clements and the tonal systems described earlier.

4. The last subject I would like to touch on is double in nature. First, can we extend the autosegmental approach to processes of a mere “local” nature—not prosodic in the usual reckoning—and if so, how? Second, what kind of answer can be provided to the deeper question, why are some features, in some languages, autosegmental?

In Goldsmith (1976b) I offered some speculations on the second question. The general thrust of the approach was to suggest that the first stage in a child’s acquisition of the phonology of a language was a stage

in which the phonological features were independent, or in our terms, autosegmental. This would suggest one ought to find rampant harmony processes in early speech, which is, as is well known, quite true. The point of the hypothesis, however, was not to analyze children’s speech, although the recent work of Lise Menn (1977) provides some striking support for this approach. The point, rather, was:

(1) those features or feature-complexes which are in fact independent in child-speech should be precisely those which may be autosegmental in the adult grammar;

(2) the process of language acquisition includes a task of “de-autosegmentalization” or, to use a less awkward term, restructuring of the phonetics into linear segments. Suppose that this is one of the tasks performed by the child: decreasing the number of independent features in the phonetics, and realizing, for example, that there is in the English phonological system an entity that we call an “n” that has both nasality and stophood as characteristics. If this is one of the child’s tasks, then there is no reason that the de-autosegmentalization should bring the number of independent tiers down to the singular number one. Thus the very same language acquisition procedures may be employed by a child learning English, Igbo, or Guarani, but the child learning Igbo, for example, will never be presented with a reason to unlearn its initial assumption that tone is an independent property (the traces of natural phonology become noticeably distinct here).

The general line of thought, then, is that the autosegmental status of tone in Igbo, or nasality in Guarani, or tongue position in Hungarian, is not a complex phonological process learned late in the process of language acquisition, but is rather the *systematization* (lending of significance of imposition of symmetries, rules, and so on) to a stage through which all learners of all languages pass. Thus the learner of English passes through a vowel harmony stage; such a learner would “have it wrong” for English, and later de-autosegmentalize English vowels. If that child had been learning a different language, though, the child would have maintained that “hypothesis,” and gone ahead to elaborate the independence of the tongue position features from those represented in the other segments.

This general line of thought has a further consequence for research. There is no reason to believe that the part of language acquisition I have called “de-autosegmentalization”—that is, the child developing a more or less single, linear model or representation for a language like En-

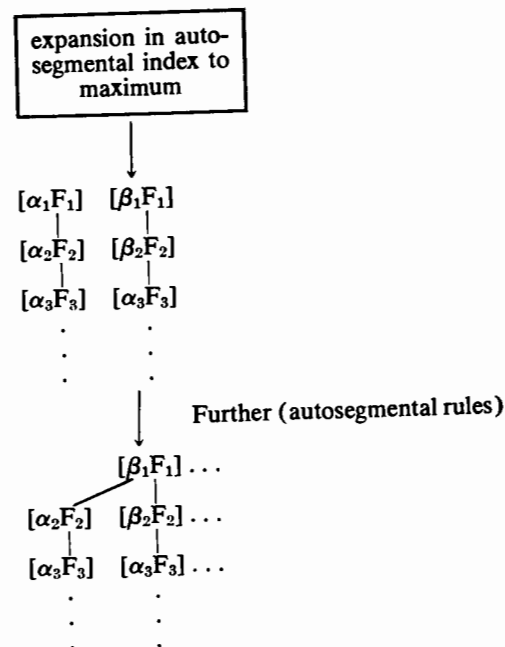

glish—should occur entirely before the child acquires more familiar sorts of phonological rules. That is, the child could begin to adopt rules of phonology for English, or whatever language, even at the point when the child's model for the underlying structure of the language is autosegmental.

If we assume, moreover, that these intermediate autosegmental levels of representation are present even in the adult grammar, then this hypothesis about language acquisition will give us a new tool for approaching the description of the adult grammar, and likewise if this new description is supported by the facts, we will be more secure in pursuing our hypothesis about language acquisition.

Now, what sort of evidence would we look for in the adult phonological grammar? Much of the evidence for autosegmental analyses of tone languages, for example, derived from "morphological" arguments—that is, a "morpheme" could be argued to be present on one autosegmental tier but not another. Such arguments will not be possible here, because we are attempting to fathom, now, an autosegmental representation which is derived from an underlyingly linear system. What evidence could we look for?

Just as much of the clear evidence for underlying autosegmental phonology comes from the existence of (underlying) morphemes which are present on only one autosegmental tier, so here the evidence we would look for would be *derived* segments present on only one *derived* autosegmental tier. One common phonological process that typifies this sort of process is the $s \rightarrow h$ change, two instances of which we consider shortly. As in the Igbo analysis, we try to show the naturalness of a large variety of phenomena when viewed as deletions.

To summarize, we are considering a model in which the number of autosegmental tiers (or "index") at the phonetic level is greater than the number of tiers underlyingly. Schematically, this presents a picture like this:



One example of this sort of process is found in many dialects of Spanish. Here I will restrict my attention to a preliminary analysis of one specific dialect, one spoken in Buenos Aires. In this dialect, all s's between a vowel and a consonant, roughly, are converted into [h].


(25) $s \rightarrow h / [+syl] - [-syl]$

This happens both inside a word ([asta] → [ahta] "until") and across word boundaries. Thus [dos] before a vowel, but [dohpapas] "two potatoes".

A restatement of (25) in terms of features fails to express the linguistic naturalness of this process. The fact that the /s/ becomes [h] rather than [r], [p], or [z] is left unexplained.

If we say, on the other hand, that the [h] derived from the /s/ does not have the same segmental status as the h in English, but is rather derived only after the autosegmentalization of the various features occurs, then a different picture emerges. The segmental process in (25) becomes a deletion of the supra-glottal features instead, and the laryngeal features of the /s/ are all that remain. Reassociation of the now-floating h-

autosegment follows. In the derivations that follow, we consider reassociation to the left.



I am assuming at the moment that the rule deleting the oral "s" segment also explicitly specifies the reassociation of the laryngeal segment. There is much variation between dialects as to whether that segment reassociates leftward or rightward, and some even with this dialect. Even under these limited conditions, however, we do find some significant advantages to this view of the process.

If the vowel preceding the underlying /s/ is an /i/, there are two possible allophones of the /i/, the normal highly tense, high front vowel found in the other occurrences of /i/, and a slightly laxer version I shall represent as /I/. Now, although i and I vary freely before s-plus-consonant, h is not the surface form of /s/ after the tense /i/. In fact, we get the surface variants as follows:



After /i/, not surprisingly, the tongue maintains its position throughout the period, just as (26c) asserts it should. Consequently, the unvoiced

flow of air causes turbulence at the tongue blade when it is in the /i/ position.

The point of all this is simply that representation (26c) says all that need be said. In particular, there is no need to add an additional rule turning [h] into [ç] after [i]—in reality, once the deletion of the oral autosegment was effected, the phonetic realization of a structure like (26c) was a *fait accompli*. Only one last point should be emphasized: the fact that the turbulence occurs at the tongue blade after [i] and not [I] is not a fact expressed formally in the phonological grammar at all; it is a consequence of the articulatory phonetics. This presents a good illustration of how some facts of phonetics may best be excluded from phonology, whereas others—like *which* articulatory organs are relatively independent and thus may be autosegmentalized—are of great significance to phonology.

A very similar analysis has been proposed by Hoskuldur Thrainsson (1978) for Icelandic. He reviews the facts concerning the well-known "preaspiration" of stops in certain contexts, and argues persuasively that the pre-aspirated stops are synchronically, as well as diachronically, derived from geminate tense stops. The first of the two stops, however, is weakened to an [h]—giving the impression of preaspiration.

In the case of Icelandic preaspiration, we could . . . say that the laryngeal articulatory gesture for the phonologically long or geminate stop is made at the "correct" point in time—i.e., when the first half of the stop is to begin—whereas the supralaryngeal articulatory gesture, namely the oral closure, is not made until the second half of the phonologically long or geminate stop begins. In other words, we have to get from a phonological sequence of two fully specified stops, as it were, to a sequence where the first member is only specified for the appropriate laryngeal features. What this means, then, is that the phonological rule of preaspiration has to wipe out or delete the set of supralaryngeal features for the first half of the stop and leave only the laryngeal feature specifications. . . . But is there any way of adequately formulating a rule of this kind within a generative framework? The answer is yes. The theory of autosegmental phonology offers a rather neat way of doing this . . . Thus it seems that in some cases a certain subset of the phonological features composing a phoneme can behave as a unit. I want to argue that preaspiration in Icelandic is just such a case [Thrains-son, p. 35].

At this point, just how these kind of analyses would be integrated into a more general view remains to be determined, but the convergence of these last two proposals seems to me to indicate a very promising direction.

5. CONCLUSIONS

In distinguishing autosegmental analysis from more traditional types of generative phonology, the emphasis here has been on the more articulated types of phonological structure proposed by autosegmental theory, disallowed by *Sound Pattern of English* postulates. On the other hand, this should not be taken as an attempt to refute the SPE tradition; *Sound Pattern* itself is clearly a working paper, a summary of results and hypotheses offered by some working phonologists at a certain point of time. Suggestions that there is a "congealed . . . [and] classically dogmatic" consensus in post-*Sound Pattern* phonology (to use Roger Lass' [1976] phrases) I find quite unrealistic. Generative phonology, quite to the contrary, has been quite supple; and I submit autosegmental theory as a proposal, which, if correct in all the ways I have suggested here, will revise in some radical ways the generative view of phonology, while remaining essentially within the general view of the goals of phonology spelled out in *Sound Pattern*.

On the other hand, I do disagree with a number of the papers presented at this conference on some basic points, even holding aside questions of immediate empirical adequacy (but cf. Goldsmith (in preparation c)).

There is a view that has been repeatedly voiced at this conference to the effect that falsifiability is, or should be taken as, a measure of a theory's adequacy and success. As such, this seems to me to be a mistaken view, possibly arising as a misinterpretation of the by-now conventional wisdom of Karl Popper. Popper was only defending the view that falsifiability is one criterion to be used in separating science from non-science, not in choosing between approaches to a subject. If occasionally it is said that even the falsifiability measure chooses between two theories "all other things being equal," it can be said only in irony since other things never are.

In any event, "falsifiability" is a sophisticated notion itself, for it surely cannot be taken in the obvious sense of producing counter-examples. Falsification at best can only be done by alternative theories, and just what the "logic" of competition between theories is is still open to considerable controversy (for some skeptical views, see Lakatos [1970] and especially Feyerabend [1975]). At all events, it is by no means obvious that the ease with which a theory can be destroyed—what we might call its "moribundity index"—is simultaneously a measure of its strength. Slightly altering Feyerabend's (1970) apt phrase, the strength of a theory might well be measured by the degree to which "the

principle of tenacity" (retaining a theory despite confounding evidence) is successfully appealed to; but this raises difficulties concerning the notion "success" which we cannot go into here.

So much for the notions of falsifiability and vulnerability. I would mention a second viewpoint, one which contends that the plausibility, attractiveness, and in the end the *value* of any theory of phonology depends on its ability to deal with a language comprehensively and in depth. That is, a theory—as opposed to a specific analysis—should lend itself to, and indeed provoke, a full-scale investigation of the details of language systems, not just considerations of a few scattered facts in various languages. When is the investigation "full-scale"? When it reveals something unexpected, surprising, and remarkable about the nature of the language being looked at, I would say. In sum, I think that the criteria for choosing a theory have little to do with falsification directly. Rather, we must adopt a theoretical viewpoint which allows for the elegant interaction of hypothesis and data that has typified science. Of course, to the extent that the hypothesis and the data must mesh to be elegant, falsification plays a role—but very much a secondary one.

Returning to phonology, we have seen that autosegmental representation provides to generative phonology one way to incorporate some more traditional and phonetically-based notions. On the other hand, as the example in section 2 has shown, it allows for revealing analyses of "suprasegmental" elements using basic techniques of generative phonology.

Whatever gains have been achieved result from the enriched descriptive power of the formalism, we should note. I have no doubt that the potentials of the formalism are not yet exhausted, and that where even a multi-tiered approach is too weak to approach the facts, a further enrichment of the formal devices can and must be developed—as, for example, that proposed by Sigurd (1972), Liberman (1975), and Liberman and Prince (1977). In short, one important way to advance phonology is to open the formal floodgates as wide as possible. I think we need not worry about drowning in any consequent torrent of formal gadgetries; as ever, the idle and useless formalisms will eventually drown themselves. The good ideas can take care of themselves.

NOTES

I am grateful to a number of people for comments on an earlier version of this paper, most especially Nick Clements, William Leben, and Carmen

Lozano, and those at the Conference in Bloomington, who, by disagreeing, helped polarize my position.

1. Curiously, no such attempt was ever made in Terence Langendoen's unsympathetic review of Firthian phonology from a generative point of view. Stephen Anderson's suggestions have been rather more serious (cf. Anderson [1974, ch. 14], [1976]). Anderson's representations are in essence identical to Bloch's, but his analysis of the advantages of a relaxation of the assumption of linearity are in close agreement with a number of the arguments in the autosegmental literature, and it seems to me we are dealing with much the same issue.

2. As I have noted in Goldsmith (1976a, b), the multi-dimensional representation is not a "discontinuous" one, in the familiar or the technical sense of the word; a careful reading shows Chomsky is not suggesting it is either, of course.

3. As stated, this assumption is possibly too strong. Leben (1977a) and Clements (1976 a, c) present quite attractive analyses in which exceptional tonal and harmonic behavior, respectively, are treated in terms of anomalous underlying association lines.

4. One criticism of autosegmental theory that has been raised (for example, in unpublished work by A. R. Walton) is that explicit criteria have not been given for determining which features are autosegmentalized within a specific language. Within the rather static picture discussed in this paper in the first three sections, and in the first three chapters of Goldsmith (1976a), in which the number of autosegmental tiers is fixed throughout the derivation, the segmental status of a particular feature is determined by its behavior with respect to such criteria as (1) whether the change in value of that feature coincides or not with the change in value of the other features; (2) whether "melodies" consisting of values for that feature must be listed in the lexicon; and several other criteria. The issue is discussed at some length in the first chapter of Goldsmith (1976b) and in Goldsmith (1976a).

Within an extended or dynamic autosegmental theory, as sketched in section 4 below, the segmental status of the various features is approached in an interesting way for the first time, I believe. There I suggest that, superficially, essentially all segments are represented on separate tiers; the more "abstract"—or, better, psychologically-oriented—levels of the derivation, however, have progressively fewer and fewer tiers, to the extent that the data of the language permit this "de-autosegmentalization". Cf. §4 below.

5. In Goldsmith (in preparation b) I try to deal with the elimination of one of the three major class features—"syllabic." It seems clear that an articulated theory of syllable structure (as discussed in the preface to Goldsmith (1976b) or Kahn (1976)), the feature *syllabic* is redundant, and many if not all of the traditional rules adjusting the value of the feature *syllabic* or inserting glides homorganic with neighboring vowels (cf. Anderson (1974), Harris (1969)) should rather be dealt with in terms of the autosegmental notions of adjustment of association lines at the syllabic level.

Siddha and Asiddha in Pāṇinian Phonology

S. D. Joshi and Paul Kiparsky

What is a grammar of Sanskrit written around 500 B.C. doing on the agenda of a conference on current phonological theories? We shall try to demonstrate that Pāṇini's grammar possesses both a theoretical content and a contemporary relevance which make it more than just a chapter in the history of linguistics. The grammar includes a generative phonology of a depth and exhaustiveness to which no modern generative phonology has even come close, which is moreover integrated with a fully worked out generative syntax and morphology, in a system of 4000 formalized rules based on very specific and elaborate principles of linguistic description. True, only a relatively small part of these underlying principles themselves are stated among the rules of the grammar. The rest must be deduced from the structure of the system and the way in which the facts of Sanskrit are analyzed there, a task which was begun by the traditional commentators and continued by modern scholars, but still remains to be carried out to the end. Thus, while the text itself is long on analysis and short on theory—just the opposite of a typical modern contribution to linguistics—it nevertheless presupposes and reflects what can be construed as a theory of grammar (though naturally differing from generative phonology in being arrived at on the basis of one language only, and in having no psychological pretensions). As for its contemporary relevance, this springs from the fact that very little in its phonological analysis of Sanskrit has to our knowledge yet been superseded, the various partial efforts at synchronic reanalysis¹ having evidently so far not led to any major revision. (Of course, Pāṇini's grammar continues to be the basic analysis of Sanskrit accepted by traditional scholars in India to this day.)