

6 The Syllable

JOHN GOLDSMITH

1 Overview and Brief History

1.1 Introduction

In 1968, Ernst Pulgram began his classic monograph on the syllable with the wise words, “conscience, courtesy, and caution require that anyone wishing to concern himself with the syllable read all, or at least most, of the enormous literature on it.” The years since his study have only magnified the challenge of this suggestion. Yet life is short, and space shorter still, and so in these few pages, I will attempt to survey the range of beliefs, models, and theories regarding the syllable that have been held by linguists, and attempt to integrate and compare them. Among the studies that I have found particularly useful are those by Fischer-Jørgensen (1952), Pulgram (1970), Fudge (1969), Goldsmith (1990), Blevins (1995), van der Hulst and Ritter (1999b), and Tifrit (2005), and my goal is only to supplement them, not to replace them. I have emphasized here the historical development of the approaches to syllable structure proposed over the past century insofar as it is relevant to today’s phonologist, and I assume the reader has at least a basic familiarity with the role and usefulness of syllables in phonological analysis. For a description of a wide range of phenomena associated with syllabification, the reader is invited to consult Blevin’s chapter on the syllable in the first edition of this Handbook, or any of the other references just cited.¹

The syllable is one of the oldest constructs in the study of language, and most studies of phonology have found a place for the syllable within them. The momentous reconstruction of the behavior of Indo-European sonants, which was the greatest accomplishment of nineteenth-century linguistics, was intimately linked

to the realization that certain segments could be realized in strikingly different ways, depending on the location in which they appeared in their syllable: elements that could be identified as glides, nasals, and liquids would be realized as consonants in some contexts, but in others, when a vowel was not present for morphophonological reasons, the segment would be realized as a syllabic peak. Working out the solution to problems of historical linguistics led directly to the development of new conceptions of phonological structure, a historical event that has not yet been completed. This chapter is an overview of the evolution of the discipline's thought on this subject.

Tradition has it that a syllable consists of a vowel, usually preceded by one or more consonants, and sometimes followed by one or more consonants. In the overwhelming majority of spoken languages (though perhaps not all),² the syllable plays an important role in analyzing phonological regularities that phonologists have placed at the center of the phonological stage. The syllable is, first of all, important for the expression of statements of phonotactics, the principles of a language that describe which strings of basic sounds are found. Why does *blick* appeal to the anglophone more than *bnick*? It is, additionally, relevant for the expression of phonological conditioning for the realization of the basic sounds: the description of the realization of a *t* in American English is far more compact if the description can use the notions of "syllable onset" and "syllable coda" than if it is forced to forego them. Finally, there are few languages in the world whose prosodic systems can be adequately and compactly characterized without making reference to the syllable.³ Prosodic or suprasegmental regularities involve a wide variety of linguistic phenomena, including timing, other rhythmic effects (such as clapping and dancing), and tonal structure.

As we shall see over the course of this chapter, the study of the syllable in recent decades has been an integral part of the development of theories of phonological representation, and to a lesser degree a part of the development of the theory of rules, constraints, and their interactions. But the study of the syllable has moved in fits and starts, with movement in several directions, all at the same time. While our knowledge, on the whole, has grown, it has done so by pursuing several different ideas.

We offer in this chapter a classification of approaches to the syllable; this is its only point of originality. There are several principal views of the syllable that have dominated linguistic discussions. Each has given rise to one or more formal models which encapsulate what is appealing about it. In the end, good reasons have been adduced for each approach, as I shall try to show, and it is hardly surprising, therefore, that much work has glossed over the differences between these models.

Of the general approaches, the first, that with the greatest longevity, is the view that sees spoken language as organizing sounds into wave-like groupings of increasing and then decreasing sonority (whatever *sonority* may turn out to be), while the second view sees the chain of segments of the language as organized into constituents in a fashion similar to the way in which words of a sentence are organized as constituents. The third view focuses on local conditions on sequence

segments, seeing the syllable as a term we use to summarize the recurring pattern of segment possibilities over the course of a word or an utterance. There are two additional views. Of these, the first is the view that the syllable, rather than the segment, is the right level of analysis for production and perhaps for perception, and linked to this is the hypothesis that while syllables are inherently ordered in time, the linear ordering of segments within a syllable may be the result of general principles of construction of syllables.⁴ The second is the view that the *dynamics* of spoken language crucially depends on the syllable. This can be approached in more than one way, to be sure. Any account of speech production must offer some account of the length and timing of the sounds produced by a speaker, and framers of hypotheses have often been tempted to establish as a principle the notion that languages tend to preserve an *isochrony* – a common temporal interval – between syllables, or between stressed syllables. More recent work has developed accounts of the syllable based on models of the temporal coordination between consonantal gestures contained in the syllable onset and the gesture producing the syllable's vocalic nucleus; see, for example, Browman and Goldstein (1988), Fujimura (1992).

In this chapter, we will survey the fortunes of these accounts, but emphasizing the first three: *sonority view*, the *constituency view*, and the *segment sequence view* of the syllable, and say just a bit about the issues of articulatory gestures and of rhythm; we give none of them the full exposure they deserve, and focus on the three that are more closely tied to questions of phonological representation. It is a remarkable fact that each of these views has flourished, grown, and developed over the last 50 years, and the extant literature has not made a great effort to give a synoptic perspective on the approaches that have been taken; we hope to fill this gap. In the final section, we return to this observation, and ask whether it is a Good Thing or a Bad Thing that largely incompatible perspectives have flourished. Should not the True View eliminate the two other views after a certain period of time? We try to offer an answer to this question in the final section.

1.2 Waves of Sonority: Whitney and Saussure

The oldest perspective on the syllable may well have been inspired by the observation that the jaw opens and closes as one speaks. This perspective, the *sonority view* of the syllable, is based on the view that each segment in an utterance has a sonority value, and that there are *crests* and *troughs*, or *peaks* and *valleys* of sonority in the speech chain, with peaks coinciding with vowels and syllable nuclei, and troughs coinciding with boundaries between syllables.

If sonority rises and falls in the course of an utterance, we might expect to find a difference in the realization of consonants depending on where in the wave of sonority they appear. On this view, a consonant that appears in a context of *rising* sonority at the beginning of a syllable – that is, before the peak of the syllable – is in a qualitatively different environment compared to those that appear in the context of *falling* sonority, at the end of a syllable. Following terminology that goes back more than fifty years, we will call the first context, that of rising sonority,

the *onset*, and the second, that of falling sonority, the *coda*, and the peak of sonority the *nucleus*.

This view was well developed by the end of the nineteenth century, and is described in considerable detail in Whitney and Saussure. For example, Whitney wrote in 1874:

The ordinary definition of a syllable . . . amounts to this: a syllable is that part of a word which is uttered by a single effort or impulse of the voice. Such an account of the matter is of only the smallest value . . . The governing principle, it seems plain to me, which determines [syllabification], is that same antithesis of opener and closer sounds upon which the distinction of vowel and consonant is founded. The vowel sounds of *any* are practically identical with those that compose our *ē* (the “long *a*” of *they* . . .); and *ē* may be so protracted so as to occupy the whole time of *any*, without giving the impression of more than a single syllable; but put between the two opener vowel elements the closer consonantal *n*, and the effect is to divide them into two parts: the ear apprehends the series of utterances as a double impulse of sound. So in *lap* there are three articulated elements, of three different degrees of closeness, but the *a* (*æ*) is so much more open than either of the others that they are felt only as its introductory and closing appendages; there is a crescendo-diminuendo effect, but no violation of unity. And *alp* and *pla*, in like manner, are a *crescendo* and *diminuendo* respectively . . . [W]hen it comes to allotting to the one or the other syllable the closer sounds which intervene between the opener, there is room for much difference of opinion . . . Thus, for example, in *any*, the intervention of the *n* between the two vowels makes the dissyllable; but the *n* itself belongs as much to the one syllable as the other . . . There is, on the other hand, more reason for assigning the *p* of *apple* (*æ-pl*) to the second syllable . . . There are . . . sounds so open that they are always vowels, never occupying the position of adjuncts in the same syllable to another sound which is apprehended as the vowel of the syllable. Such is especially *a*; and *e* and *o* are of the same character. But *i* and *u* . . . become *y* and *w* on being abbreviated and slighted in utterance . . . Vowel and consonant are the two poles of a compound series, in which are included all the articulate sounds ordinarily employed by human beings for the purposes of speech. (pp. 291ff.)

Saussure⁵ was centrally concerned with the realization of the three phonological classes of segments in Indo-European: those that were always realized as consonants (what we would today call the *obstruents*), those that were always realized as vowels (*non-high vowels*), and those which were realized in the one way or the other, depending on their context (*sonorants*); this third group consisted of liquids, nasals, and glides/high vowels. Saussure (1995: 222) established a set of four ordered rules for this, which aim primarily to account for the realization of sonorants, whose surface form is heavily dependent on the phonological context:

- Rule 1: a vowel that follows a sonorant puts that sonorant into an onset (Saussure’s “explosive”) position.
- Rule 2: an obstruent inhibits a sonorant that precedes it, as does silence. Here, something “inhibited” is syllabic, in modern terms.
- Rule 3: a glide which has become inhibitive has the same effect as a vowel.
- Rule 4: a sonorant which is in an onset has the same effect as an obstruent.

Saussure indicated explicitly that these rules must be applied sequentially from right to left, from the end of the word, and illustrated the effects on forms. The influence of his familiarity with the scholarship of Panini is evident.

As we will see, most theories of the syllable formulated with any desire for rigor have begun with a set of three or four basic syllabification rules, and the cross-theoretical appearance of such principles allows for a typology of sorts of approaches to understanding syllabification. We will see this below, in the work of Pulgram, Hooper, and Kahn.

The sonority hierarchy is discussed in some detail by phonologists early in the twentieth century, such as Jespersen (1904), van Ginneken (1907), and Jones (1918).

Fischer-Jørgensen expressed the prevailing perspective in her classic paper (1952), viewing the syllable “as a unit of speech containing one relative peak of prominence. The division of the chain of speech into syllables may be due simply to the inherent loudness of the successive sounds, but the peaks may be reinforced or altered by arbitrary changes of loudness, and this means may also be used to give a clear delimitation of the units.”

1.3 *Constituents and Structure*

1.3.1 Pike, Hockett, Fudge: The Arboreal View Immediately after World War II, two studies were published which proposed a new account of the syllable on the basis of the notion of *constituent*, a concept that was being developed at the frontier of syntax at the time: Pike and Pike (1947), and Kuryłowicz (1948).⁶ This was a moment during which Bloomfield’s notion of *constituent* was coming into general use in syntax, and to some extent replacing the earlier view of syntax, according to which syntax focused on asymmetric relations between pairs of words. On this newer constituent-based view, sentences were successively cut into smaller and smaller pieces, until reaching the word. Pike and Pike argued, why stop there? Why not continue to chop up the utterance into finer-grained pieces, since we already have a name for them: syllables! The momentous step of bringing insights from the domain of syntax into the treatment of syllable-level phenomena has continued to play a major role in the development of theoretical views, notably in approaches employing the concept of government.

Within American phonology during the post-War years, there were mixed feelings about the importance of the syllable for phonology, based to some extent on the fact that in the overwhelming majority of cases, syllabification is phonologically *predictable*: given a sequence of segments in a word in a specific language, the location of the syllable boundaries is predictable, which is to say, not distinctive. In a framework which required fully predictable information to be absent from the phonological representation such as was dominant in the United States at the time, it was reasonable to draw the conclusion that syllable structure should not be present in the phonological representation. On the other hand, syllable structure is probably the single most important conditioning environment for segmental rules (we will see examples of this in Section 2 below), so it is essential

for a theory to allow for the existence of a phonological representation in which syllable structure interacts with “choice,” so to speak, of the phonetic spelling-out of underlying segments.

The question of the predictability of syllable structure is not as simple as suggested in the preceding paragraph, however. In some languages, syllabification appears to operate without any reference to morphological or word boundaries – as in Spanish – while in German, word boundaries and at least *some* morphological boundaries are critical to syllable establishment, and much the same is true of English, as we will see below (Section 2.3).⁷

1.3.2 Syntagmatic and Paradigmatic Syntacticians and phonologists agree on another point, one which is rather more abstract. It is this: when we study linguistic units – words in syntax, and segments in phonology – it is important to both analyze the specific relationship that a given unit bears to its neighbors in an utterance (and we call those syntagmatic relationships: subject, direct object; nucleus, onset), and the categories into which the inventory of units can be usefully subdivided (nouns, verbs, adjectives vocoids, fricatives, nasals). From the very earliest work we have cited, phonologists have borne in mind the fact that a particular item, such as a vocoid – a segment in which spectral resonances are its most salient properties – can be either the nucleus of a syllable, or not. In the former case, we call it a *vowel*, and in the latter we call it a *glide*. But using the convenient terms *vowel* and *glide* should not lead us to overlook where the constancy is and where the difference is between a vowel and a consonant.

Thus “being-a-vowel” – as opposed to “being-a-vocoid” – is a fact about the *role* a segment plays in a particular spot in an utterance: it is syntagmatic. How should this view, nearly universally held, be integrated into a larger or theoretical point of view? One negative conclusion is that it should not be as a binary feature (Selkirk 1984a makes this point explicitly, and she is not the only one to do so). What *is* it about its role in the syllable that is crucial: is it the fact that a syllable must have exactly one element that is nuclear or syllabic? Is this property something that inheres in the formal relationship between the syllabic element and the syllable constituent? Some have opted for this second alternative, and called the relationship *head*.⁸ In constituency-based models, it is often assumed that among the nodes depending from a given node, there is exactly one node marked as its head, and the syllable nucleus is the unique element of the syllable that is a member only of constituents marked as head, within the syllable. In other models, the nucleus is identified as the most sonorous element. What is generally agreed is that the notion of *feature* is not best suited for this job: features, as an inherent non-relational object, are ill-suited for representing the important characteristics of differences that are inherently relational, and the difference between a vowel and a corresponding glide lies in the *relation* that exists between the segment and the context in which it appears.

1.3.3 How to Parse CVC The constituent model of syllabification naturally suggests that the syllable nucleus forms a constituent with either the onset or the




Figure 6.1 How to parse CVC.

coda (as in Figure 6.1): I say “suggests,” because the flat structure of (c) is a possible analysis in a model with constituents.

Structure (a) in Figure 6.1 is the one most widely defended and used. The structure in (b) has been proposed on occasion. In Japanese, for example, Kuzobono has argued for structure of this sort (see Kubozono 1989a), for reasons that seem to be strongly linked to the centrality of the mora to Japanese phonological structure; Yi (1999) argues for a similar analysis of Korean, as did Bach and Wheeler (1981). There is a view that what has been treated as a CVC syllable should rather be analyzed as a CV syllable, followed by some kind of defective syllable, and such views involve structures more like (b) than like (a) or (c) – a syllable with an onset but no nucleus; we discuss this below (Section 1.12). The structure in (c) has been proposed (by Saporta and Contreras 1962, cited by Harris, and defended in detail by Davis 1985); so-called “flat models” of syllable structure also come close to this; see the brief discussion below in Section 1.8.

Part of the difficulty in establishing a structure conclusively lies in the fact that there are few generally accepted principles for determining constituency in phonology, and those that do exist tend not to give decisive answers when applied to this question.

Some researchers have explored the relevance of language games and tasks that can be studied by means of psychological tests; in a series of papers (Stemberger and Treiman 1986; Treiman 1986; Treiman *et al.* 2000) Treiman and colleagues have taken this approach, and been influential in turn (see, for example, Stenneken *et al.* 2005 and references there). Some of this work is based on the suggestion that manipulations of linguistic segments require what we might think of as cutting and repasting pieces of syllables to form new syllables: for example, if /krɪnt/ and /glʊpθ/ are to be merged, will a speaker give us /krʊpθ/ or /kɹɪpθ/, or something else? If we formulate a binary-branching tree structure over the set of phones, then a location between any two adjacent segments can be associated with a height in the binary tree, and we may hypothesize that locations corresponding to “high” nodes are preferred as positions for cutting: the position between the last segment of the onset and the first of the rhyme corresponds in this sense to a higher position in the syllable than the position between the first and second consonants of a syllable onset. But hierarchical structure is not the only, and certainly not the most direct, model capable of making predictions as to preferred cut-points in psychological tasks; the fact that a piece of beef may be

easier to cut at some points than at others should not be taken as evidence that the meat is hierarchically organized.

Perhaps the most widespread principle which phonologists have attempted to apply is one that says that when adjacent segments are part of smaller constituents, there should be stronger cooccurrence restrictions bearing on the two segments than when the segments are more distantly related, that is, are parts only of larger constituents. On this account, if there is a rhyme constituent, as in [_{Syllable}Onset[_{Rhyme}Nucleus Coda]], then we should find stronger cooccurrence restrictions between the nucleus and the following coda segment than between the vowel and the (final) consonant of the onset (since the latter pair of segments are in a larger constituent, the syllable, but the former are in a smaller constituent, the rhyme). This perspective is clearly presented in Pike (1967), Fudge (1969), and Selkirk (1982), and appears to be originally rooted in the early reflections on the meaning of constituency, notably Rulon Wells's 1947 paper (Wells 1947); again, the principles established on the basis of exploring syntactic structure were applied, in retrospect with relatively little reflection, to problems of phonological representation. Davis (1985) discusses a number of weaknesses of this argument.

Unfortunately, as it stands, this is a rather vague formulation, and even if the notion of "strong cooccurrence restrictions" can be clarified (as it certainly can), it is not clear why constituent structure should cause the cooccurrence restrictions in question. This imprecise notion can be interpreted as standing in for the more precise and explicit measure of *mutual information*.⁹ Mutual information measures the degree to which the frequency of any pair of successive segments departs from independence, that is, the degree to which the probability of a given pair of segments departs from the product of their independent frequencies. We can use such measures to determine whether pairs of segments that are structurally closer have, on average, greater mutual information, all this as a quantitative measure of the validity of the claim that syntagmatic structure in the syllable has an impact on the possibility of cooccurrence restrictions.

Harris (1983) argues that Spanish has a restriction limiting the number of segments in the rhyme to a maximum of three. He proposes this in a model in which the nucleus contains two segments in the case of a nucleus (e.g. *muerto*), and in other forms, the coda can contain two segments; and thus there is no purely local way of formulating the restriction that the coda has an upper limit of three elements. Harris is at pains to emphasize that he is aware of the mismatch between his finding and the ability of the hierarchical model to easily incorporate it (the problem comes from the apparent need to say that the rhyme cannot have more than one element in both the nucleus and coda, though it can have more than one element in either one separately).

1.4 Syllable Timing

Kenneth Pike (1947) was also responsible for the introducing the terms *stress-timed* and *syllable-timed* as descriptors of a language's rhythm. Pike suggested that some languages (he cited English) were stress-timed, and some were syllable-timed (he

mentioned Spanish). By this he meant that there was a strong tendency in a stress-timed language for stresses to appear equally spaced in time, while in syllable-timed languages, the tendency was greater for syllables to be equally timed (or *isochronous*). While this difference has survived many decades of usage, and a clear formulation in Abercrombie (1967: 97), it has not found experimental support over the years. Roach (1982) gives a brief summary of the issues raised and the difficulties encountered in dealing with the claims behind this distinction.

1.5 Classical Generative Phonology

During the classical period of generative phonology, many phonologists accepted the proposals in *The Sound Pattern of English* (1968), and SPE did not include the syllable within its linguistic *Weltanschauung*. On that all are agreed, but beyond it, there is little consensus regarding the relationship of the study of the syllable to generative phonology, and it is probably fair to say that there are no historical truths – only points of view. One widely held view is found, for example, in Féry and van de Vijver (2003a: 3): “In the seventies, several phonologists, such as Vennemann (1972a), Hooper (1976), and Kahn (1976), proposed including the syllable as a prosodic unit in generative phonological theory. The relevance of the syllable for linguistic theory has increased ever since.” While this is certainly true, in the sense that it contains no false statements, it does suggest something that is not true, that the syllable has no relevant history before SPE. Peter Auer (1994) suggests that “credit for the restoration of [the syllable] is . . . due to a group of phonologists who in the 1970s, in schools such as Natural Generative Phonology, Natural Phonology, and Syllable Phonology, attacked orthodox MIT generative phonology (represented by *The Sound Pattern of English*) and whose inadequacies they showed.”

1.6 Pulgram on the Syllable

Pulgram (1970), written in 1968 and published two years later, has had a great impact on the field despite the fact that it is relatively rarely cited explicitly. It offered a number of proposals (such as onset maximization) that are still widely adopted today (see Bell 1976 for an insightful review).

Pulgram offers a modern interpretation as to how to view the relationship between language-particular and universal characteristics of syllables and syllabification:

If the syllable is an operative unit of all languages, it is also a universal of language. Its definition must be . . . the same for all languages, regardless of the varying unit inventories in the different [languages] . . . there arises the interesting question whether it might not be possible to arrive at a phonotactic definition of the syllable which . . . does have universal validity for all languages. The question is, in other words, whether the phonotactic rules on syllabation might not be formulated in such a way that they are applicable to all languages, even though their implementations

in the different languages must differ because of the underlying differences of phonotactics. I believe that such general phonotactic rules on syllabicity are not only possible but also necessary for the proper syllabation of any utterance in any language. (p. 23)

His account is the first to offer a sequence of ordered rules for syllabification (p. 70ff.), beyond what we have seen Saussure had already proposed. After determining the size of the domain in which syllabification will be established (his "Rule 1"), his Rule 2 of *maximal open syllabicity* places a syllable boundary after each vowel. Rule 3, of *minimal coda*, says that if a phonotactic condition forbids certain vowels to appear syllable-finally, then the syllable boundary is shifted to the right (which is to say, one or more consonants are shifted "to the left" of the syllable boundary), but only the smallest number of shifts necessary to achieve a sequence that is possible word-finally. A sequence $C_1V_1C_2C_3V_2$ will be syllabified as $C_1V_1-C_2C_3V_2$ by Rule 2, and then as $C_1V_1C_2-C_3V_2$ if and only if V_1 may *not* appear syllable finally in the language and C_2 is a permitted syllable coda in the language. Rule 4, *maximal onset*, shifts the syllable boundary "to the right" (in the sense just discussed) if the syllabification so far has resulted in a syllable-initial sequence which cannot appear word-initially. Thus *e-mploy* is resyllabified by Rule 4 to *em-ploy*. Pulgram's Rule 5 (*principle of the irregular coda*) says that in case an interlude cannot be parsed into a legitimate word-final and a word-initial sequence, the coda must accept the material which it would not otherwise (i.e. word-finally) be forced to accept: "If the necessary transfer from syllable-initial to syllable-final position leads to an inadmissible syllable-final group of consonants, then the burden of irregularity must be borne by the coda rather than the following onset." (p. 51). Pulgram gives the example of Spanish *transcribir*, which he syllabifies as *trans-crí-bir*, despite the impossibility of word-final *ns* in Spanish.¹⁰ Although, as we will see in Section 1.17, Fischer-Jørgensen had already documented a range of languages in which the set of interludes is significantly broader than the sequences of word-final plus word-initial sequences would suggest, Pulgram's core system does not allow for that, and he attempts (in this writer's view, unsuccessfully) to come to grips with that wrong prediction.

Pulgram's account is geometrically flat; the information that is contained in the correct syllabification of an utterance is, once the nucleus has been identified, nothing but a statement of where the boundaries are between syllables. One exception to that statement must be made, however: Pulgram emphasizes the importance of finding a notational means of expressing the idea that in English, a consonant may straddle a syllable boundary, which is to say, may be ambisyllabic.

1.7 Natural Phonologies

The syllable played a central role in natural generative phonology, described by Vennemann, Bybee, among others, and natural phonology, developed by Stampe and others (see Stampe (1972), and that played an essential role in the recrudescence

of work on the syllable. Hooper (1976) presents a range of arguments within an essentially formal and traditional generative perspective in favor of syllable analysis, utilizing a formal symbol (in the event, \$, following Vennemann) among the string of phonological segments to mark syllable division. She notes, for example, that assimilation of nasals for point of articulation is found not simply when a nasal is followed by a segment specified for a point of articulation – but rather when the nasal is at the end of a syllable, and followed, in the next syllable, by an element with a point of articulation. She offers this as an explanation for the non-assimilation in words like *muevo* [mweβo] ‘I move’ despite the assimilation found in forms such as *un huevo* [unweβo] ‘an egg’.

Hooper proposes an ordered set of rules for syllabification (p. 527): (1) Place a ‘\$’ between adjacent syllabics; (2) in a VCV sequence, insert a \$ to form V\$CV; (3) place a ‘\$’ before a sequence of obstruent followed by non-nasal sonorant plus vowel – though exceptions blocking this in the case of *tl*, *dl* are discussed, and a universal set of possible blocking conditions are proposed. Rule (2) is considered universal, while rules (1) and (3) are not – the corresponding rules in languages differ in specific ways that need to be made explicitly. Syllable theories deploying a boundary symbol to demarcate syllable boundaries can be understood either as claiming that’s all there is, as far as syllables are concerned, and in particular there is no hierarchical structure, or they can be viewed as making modest positive claims, leaving open the possibility that there is further structure that the boundary symbol notation fails to indicate. Authors do not always make it clear which perspective they adopt. Selkirk (1982) notes a number of authors who use this notation (p. 354), but at least one of them (Hockett 1955) clearly indicated the presence of additional hierarchical structure.

1.8 Flat Structure

Kahn (1976) provided a number of convincing arguments for integrating the syllable into formal, generative accounts of phonology, and a number of influential studies within this framework followed, including, notably, Kiparsky (1979), Harris (1983), Selkirk (1982), Clements and Keyser (1983), all of which reflected in various ways both the formal concerns of classical generative phonology and the emphasis placed on the analysis of the syllable by natural and natural generative phonology. Kahn’s work showed the usefulness of a formal model in which syllables were represented as symbols on a distinct tier, formally parallel to an autosegmental tier, an approach that made ambisyllabicity a natural notion. We will return to this in Section 2.3 below, when we consider the process of flapping in American English.

1.9 Metrical Phonology

Metrical phonology, which began as a theory of linguistic rhythm at the syllable level and above (Liberman 1975; Liberman and Prince 1977), quickly extended its domain to syllable internal structure (Kiparsky 1979; McCarthy 1979b), and




Figure 6.2 Timing tier.

provided the means to explore the possibility of hierarchical constituent structure within the syllable, extending the analysis that had begun with Pike, Hockett, Fudge, and others. This work was part of an intensive period of work on the theory of phonological representations, which also included work on autosegmental accounts of tone, quantity, and harmony. Most important for the theory of the syllable was the usefulness of autosegmental representation in understanding the nature of long (i.e. geminate) consonants and long vowels. This new perspective allowed an account with far fewer paradoxes: long segments would henceforth be analyzed as complex 2-to-1 representations linking a single segment (a consonant or a vowel) on one tier, with two segments on a tier whose elements represent temporal or rhythmic information, often called the *timing tier* for this reason. See Figure 6.2, where a long, or geminate, consonant is represented with a one-to-many association to elements on a timing tier, here represented as filled dots.

1.10 Sonority Redux

Hankamer and Aissen (1974) was an important reminder to the American phonological community that a gradient perspective on the distinction between vowels and consonants was often critical for understanding a phonological process, a message that the natural phonologists had also tried to communicate; see Stampe (1972), Hooper (1976), and Vennemann (1972). The literature at this time shows relatively little awareness that sonority analysis has deep roots in the linguistic literature, as we have seen, although some noted the role that sonority plays in Panini's grammar of Sanskrit. James Harris (1983: 15) refers to "the familiar sonority scale $V > G > L > N > O$ ", and (p. 21) employs this scale to establish the generalization that consecutive consonants in the Spanish onset must be of increasing sonority, and not adjacent on this sonority scale. Selkirk (1984a) argues for a replacement of the major class features (sonorant, consonantal, syllabic) by a variable, called sonority, that takes on values (perhaps limited to integers, perhaps not), and perhaps in the range $[0,10]$.

A few years later, Dell and Elmedlaoui (see their 1985 and 1988) published a highly influential analysis of Imdlawn Tashlhiyt Berber (ITB) which presented the strongest argument to date for the importance of sonority in the treatment of syllabification, based on an elaborate set of principles which came down essentially to this. Syllabification is established by a sequence of ordered phonological rules, divided up into a set of “core syllabification” rules, followed by a further set of attachment rules. The core syllabification rules are a sequence of virtually identical rules that differ only in the sonority of the segment that they apply to; the rules all take the form: *Associate a [core] syllable with any sequence (Y)Z, where Z is a segment of type T*; and the algorithm passes through eight instantiations of that rule template, as the variable T proceeds through the segment classes: {the vowel *a*, high vocoids, liquids, nasals, voiced fricatives, voiceless fricatives, voiced stops, voiceless stops}, which is to say, as the variable T passes through the segment inventory from most sonorous to least sonorous.

The analysis takes into consideration the fact that in ITB, segments at any point in the sonority hierarchy can be nuclei of syllables, including even voiceless stops, in the right phonological context, and for any segment type, there are examples of segments in particular morphemes which alternate between being syllabic and being non-syllabic, based on the larger phonological context. Examples are given in Figures 6.3 and 6.4, where all and only the syllable nuclei are indicated with upper case letters (extending a notation employed by Dell and Elmedlaoui). As Dell and Elmedlaoui point out, the facts in ITB suggest that the traditional search for the language-specific boundary between possible and impossible syllables is misguided, or at least inadequate; it appears to be necessary to provide a notion of *preferred* syllable structure, since many of the incorrect syllabifications of words in ITB are composed of syllables, each of which would be possible if it occurred in some other context.

Proposing an analysis within a generative derivational account, Dell and Elmedlaoui encode their notion of preference by ordering a rule which assigns a preferred syllabification (e.g. segment plus low vowel) *earlier* in rule-ordering than a rule which assigns a less-preferred syllabification (e.g. segment + liquid), along

<i>3m.sg.</i>	<i>3f.sg.</i>	<i>Gloss</i>
IldI	tLdI	‘pull’
IrbA	tRbA	‘carry on one’s back’
IxsI	tXsI	‘go out (fire)’

Figure 6.3 IT Berber.

<i>2sg.perfective</i>	<i>3f.sg.perfective with dat. 3m.sg.object</i>	<i>Gloss</i>
tRgLt	tRgLAs	‘lock’
tSkRt	tSkRAs	‘do’
tMsXt	tMsxAs	‘transform’

Figure 6.4 IT Berber.

with a convention on rule application that blocks a later rule from syllabifying a segment that has “already” been syllabified. Other issues that seem to involve well-formedness of representations also enter into the formalism of rule application: for example, the *observation* that the only syllables without an onset are those that are word-initial becomes transformed into a constraint on syllable well-formedness; one that forces a rule not to apply if the rule’s output would violate the constraint (e.g. underlying *ihaultn* ‘he made them (m.) plentiful’ becomes first *i* [*ħa*] *ultn*, as the sequence *ħa* forms a syllable; but then *u* does not become the nucleus of an onset-less syllable, because the result of that operation would violate the constraint on word-internal onsetless syllables, and the final syllabification is in fact [*I*] [*ħA*][*w L*] [*t N*]). Dell and Elmedlaoui also note that core syllabification appears to apply from left to right, in the sense that the correct syllabification of underlying *rksx* is [*R*][*kSx*] rather than [*Rk*][*sX*] ‘I hid’, but they note that the data is far from unambiguous on this point.

The analysis proposed by Dell and Elmedlaoui was thus a demonstration that it was not just the *inherent sonority* of the segments in a language, specified along an articulated scale, that could be the primary determinant of how syllabification was determined in a language, but the *relative sonority* of each pair of neighboring segments could play a crucial rule. Neither the constituency view of the syllable, nor the sequence-based view that we consider below, is capable of providing explicit formal means for making syllabification dependent on the relative sonority of segments in a string.

A number of generalizations that use the notion of sonority have been proposed to characterize common properties of syllables. Of these, the most important are the following:

- The Sonority Sequencing Generalization: sonority rises during the onset, and falls over the rhyme. Selkirk (1984a: 116), citing Hooper (1976) and Kiparsky (1979, 1981a), writes:

In any syllable, there is a segment constituting a sonority peak that is preceded and/or followed by a sequence of segments with progressively decreasing sonority values. . . . The SSG can be viewed as imposing universal constraints on the possible form of language-particular sets of conditions on syllable structure. It in no way constitutes on its own a theory of syllable phonotactics, however, for languages will differ precisely in their choice among the various conditions on terminal positions that are consistent with [the SSG].

- Minimum Sonority Difference (or dissimilarity). Steriade (1982: 94) proposes that once an appropriate numerical sonority hierarchy has been established, along the lines of Selkirk (1984a), a language may impose the restriction that adjacent segments must be a minimum sonority distance from each other.
- Dispersion Principle, proposed by Clements (1990): all other things being equal, a language will preferentially maximize sonority difference in the syllable onset, but minimize it in the coda.

But with all this in mind, what *is* sonority? The set of answers that phonologists have provided to this question range across a wide span of opinions as to what the ultimate object of phonology's study is. Some would suggest that its fundamental motivation lies in its association with physical energy, or with the degree of opening of the mouth, or both, while another phonologist, less concrete and more abstract, might offer a different answer: sonority is the name we give to our method of organizing the segments from a language along a one-dimensional scale, with the ultimate purpose of describing permissible syllables. This latter answer raises two questions immediately: if we could identify sounds independent of the language in which they appear, would it be the case that for any such pair s_1 and s_2 which can be found in several languages, their ordering would always be the same across languages? – that is, if l is more sonorous than n in one language, is the same true in every language that has both? A second, independent question is this: if it is indeed useful to compare the sonority of two segments s_1 and s_2 in a language by modeling with arithmetic values, would we want to say that s_1 has the same sonority in every environment, or could sonority be dependent on phonological context? We will return to this question in Section 1.13.

1.11 Worrying About Slots That Hang From Trees

The Pikean, arboreal view of syllables can be pushed to the point where the terminal nodes of the tree are viewed as playing a more important role in the theory than the segments do – if by “terminal nodes” we mean (as syntacticians often do) not the symbols referring to the phonological segments directly, but some sort of node that may be “empty” of any given segment; such a view leaves open the possibility of phonologically null elements that play a significant role in the model. In this section, we will review how this reversal has an impact on the treatment of onset clusters; in the next, how it leads to more radical statements about the nature of codas and empty nuclei.

In English and many other languages, an onset can consist of a single obstruent (*pa*) or sonorant (*la*), or it can consist of an obstruent plus a sonorant, in that order (*pla*). In order to account for two-segment clusters, the arboreal view instructs us to include a phrase-structure rule of the form *onset* $\rightarrow C_1C_2$, but what should C_1 and C_2 be? It seems easy to decide that C_2 should specify that the segment in this position should be a sonorant, but what about C_1 ? The problem is that phrase structure rules are well equipped to deal with generating sets of strings like $\{\emptyset, p, b, l, r, pr, bl, pl, bl\}$, but phonologists are not always satisfied with the result that they produce. A simple phrase-structure analysis is given by the following rule:

$$(1) \text{ Onset} \rightarrow \left(\begin{pmatrix} p \\ b \end{pmatrix} \right) \left(\begin{pmatrix} r \\ l \end{pmatrix} \right)$$

But the phonologist may recoil at the conclusion that seems to follow from this: that the position of the onset consonant is different in *pa* and in *la*, that there is

no single statement that says that every class of segments can appear in the onset, and that the ordering of the obstruent and sonority must be specified in this language-particular fashion even though it reflects a widespread property of languages (e.g. Selkirk 1982: 346). Ultimately it is not clear how weighty any of these three considerations are. For example, we might argue that the first is not a valid conclusion to draw. Rule (1) will generate $[p]_{\text{onset}}$ and $[r]_{\text{onset}}$, and not something with an non-terminal node distinguishing between the two – that is, there is nothing like this to worry about: $[\text{[}_{\text{obstr}} p]_{\text{onset}}]$. To put the same point more informally, phrase-structure rules do not create slots; they determine the set of admissible tree labelings.

1.12 Government Relations

Beginning in the 1980s, a number of phonologists have explored the consequences of encouraging syllable representations with empty nuclei, along with strict constraints on what can appear in onset or coda position – constraints which have the effect not just of *allowing* for the possibility of empty nuclei, but of strongly *requiring* a wide use of empty nuclei in the analysis of real data. For example, the constraint that there be no codas at all (Harris and Gussman 1998), and the constraint that there may be only one consonant in an onset, has the consequence that there are at least as many syllables as there are consonants in a given utterance. Proponents of this view criticize earlier perspectives as being too bound to representations whose terminal elements are the observed phones of the utterance:

... we should first raise a fundamental question about the central premise of the phoneme-centred view: is it really the case that syllable structure is projected parasitically from segment strings? Suppose we entertain the alternative idea that syllable structure should be defined independently of segment strings and word structure. What empirical consequences flow from making the conceptual switch to this syllable-centered view? One immediate consequence is a rejection of the assumption that every syllabic position is necessarily occupied by a segment; there may be syllabic positions without any associated segmental content. (Harris and Gussman 1998)

One direction in which this view has been developed has been the pursuit of the hypothesis that there are no codas at all in phonology, and that all syllables are of the form CV.

From this perspective, much of the work of the phonological analysis turns to accounting for where empty vowel positions may occur, for the theory demands that there should be many of them. There is a natural similarity between this kind of phonological analysis and government-binding syntax, in the sense that both require positing a surprising number of unfilled terminal nodes, and in both, a large part of the formal account of what is grammatical and what is not is largely an account of where null nodes may appear. This perspective has been developed by a number of European phonologists, notably Kaye, Lowenstamm, Vergnaud, and Charette, to mention just four, and see Chapter 16 in this book. Considerable

care is necessary in developing a theory of syllables in this direction, if only because it is tempting to think that something has been explained when it has been labeled by a hidden variable: if, in a given language, word-final consonant clusters are permitted that match the cluster possibilities in syllable onset, one must be circumspect in determining just how much explanation is achieved by positing an abstract, silent vowel at the end of the word: the theoretical savings are no greater than what it would take to express the same constraint *without* positing the abstract syllable.¹¹

1.13 Derived Sonority

One line of work has developed the idea that one must distinguish between the inherent sonority of a segment and its sonority in a given context (Goldsmith 1993a; Goldsmith and Larson 1993; Larson 1992; Laks 1995; Tchobanov 2002). This framework of dynamic computational models employs numerical values for sonority, and is embedded within a model that includes a learning algorithm, so that adequate values for sonority can be automatically learned from a phonological sample from any given language.

There are three central ideas in this approach: first, that prosodic prominence takes on values on a numerical scale not restricted to integers; second, that there is a difference between inherent (we might say, underlying) prominence and derived, or contextual, prominence; and third, that a language can identify those elements whose prominence is a peak, that is, a local maximum in a numerical sense. "Prominence" here refers to sonority, when considering models of syllabification, and accent, when looking at models of accentuation.

We will limit our discussion to the former case. Thus we must compute for each segment, or timing tier unit, its sonority level, which will be a combination of its inherent sonority and effects that impinge on the unit from the context it finds itself in. The effects divide into two sorts: those that are specific to end units (corresponding to appendix effects: see Section 2.1) in a word, and those that result from the influence of neighboring segments. In the following equation, we define the sonority of the i^{th} segment at time $t+1$ as the sum of four terms: the segment's inherent sonority, the activation that may be due to a possible edge effect, and a weighted sum of the sonorities of the segments immediately to the left and to the right. The variables α and β specify the degree to which a segment influences its left- and right-hand neighbors; these values are typically negative, giving rise to a competition between neighboring elements. If we note the activation of the i^{th} unit after t computations as A_t^i , then the operative formula is given in (2).

$$(2) \quad A_{t+1}^i = \text{Inherent}(i) + \text{Position}(i) + \alpha \times A_t^{i+1} + \beta \times A_t^{i-1}$$

The system reaches effective equilibrium after several iterations, and all segments that are peaks of sonority are (that is, are predicted to be) the nucleus of a syllable. Languages typically put minimum sonority conditions on what elements may be

a syllable nucleus, and violations of those minima are the only way in which illicit syllabification arises within this model.

1.14 *Optimality Theory*

Optimality Theory early on offered an account of the rough typology of syllable types by proposing two syllable constraints that are specific to syllables, *ONSET* (violated by any syllable not containing a filled onset) and *NoCODA* (violated by any segment in a coda) (Féry and van de Vijver 2003); these can be usefully compared with Pulgram's Rules of *maximum open syllabicity*, and of *minimal coda*, as discussed above. One could imagine a different set of constraints, couched within an optimality theoretic framework, that would account for syllabification, but most work to date has assumed some version of these constraints, or constraint families.

If faithfulness constraints (*DEP*, *MAX*) are ranked lower than the syllable constraints, then a language will use either epenthesis or deletion to ensure that surface forms are of the form CVCVCV. If *ONSET* is ranked higher than the faithfulness constraints, which are in turn ranked higher than *NoCODA*, then some strategy, such as consonant epenthesis, will emerge to provide a consonant to precede any vowel that is not preceded by a consonant. If the faithfulness constraints outrank the syllable constraints, then codas, coda clusters, and onset clusters may emerge, if the lexicon and the morphology provide such circumstances. Such an account uses cover terms such as *NoCODA*, which can be viewed either as promissory notes, or better, as implicit hooks into whatever theory of phonological representation one chooses to use, provided it permits access to coda-labeling as such.¹²

1.15 *Must We Choose Between Sonority and Constituency?*

I have tried to emphasize in the discussion to this point that there are at least two different pictures of what the syllable is that have evolved, one focusing on the notion of sonority, and the other concerned with constituency structure at and below the level of the syllable. In the next section, we will turn to a third view, the one focusing on patterns of first, or possibly second, order restrictions on segment sequences.

Perhaps all scholars would agree that the core phenomenon lying behind the concept of the syllable is the fact that we can divide the inventory of a language's sounds into two sets, the vowels and the consonants, each of which share a number of articulatory and acoustic properties, *and* that there is a strong tendency for utterances to produce sounds successively, first one from one of these groups and then one from the other group. But things are not really so simple; there are recurring restrictions on what sequences of consonants may occur, but fewer restrictions on what vowel-consonant, or consonant-vowel, sequences may occur. The notion of sonority emerges as soon as we note that we often find a sequence X-Y that occurs at the beginning of a syllable being matched against not the

sequence X-Y, but rather the sequence Y-X, at the end of the syllable. It is this purely structural discovery of a mirror-image of segment sequence possibilities, plus a desire to associate the directionality of a permitted linear sequence of segments with a physical dimension, that brings us to the notion (or *a* notion) of sonority, because what we describe as increasing during the first part of a syllable, and decreasing during the second, *is* sonority.

Must we *choose* between such a sonority-based view of the syllable and a constituency-based view? One possible answer to the question that begins this section is “no,” on the grounds that many analyses have been offered that establish explicit connections between the sonority of the segments present and the syllable constituents (onset, coda, nucleus) to which these segments are assigned, and those analyses are not self-contradictory in any obvious way; various OT perspectives, as we have seen, are examples of such analyses. In a recent survey, Zec (2003) sums up her view of the optimality theoretic perspective with the words, “the prosodic constituency is viewed here as a hierarchy of sonority peaks” (p. 125). But there *is* no natural relationship between establishing a height over a sequence of points as in Figure 6.5, on the one hand, and a hierarchical representation, on the other – no natural connection in either direction. To make such a connection, one must add principles of one sort or another; one could offer a positive markedness constraint, as Morelli (2003: 359) does:

SONORITY SEQUENCING PRINCIPLE: In a syllable, sonority increases toward the peak and decreases toward the margins.

But a thoughtful review of whether we need both conceptions, and how well they can co-exist, leaves us with some unanswered questions, notably about the notion of the syllable as an example of constituent structure.




Figure 6.5 A sonority curve.

In the context of syntactic analysis, the notion of immediate constituent was outlined in Bloomfield (1933), and developed by Wells (1947). As Percival (1976) has discussed in detail, analysis in immediate constituents was understood as an alternative to a word-oriented notion of syntax, a view according to which a verb took a subject *noun* and object *noun*, rather than *noun phrase*, though the noun that was the subject or object could have other words modifying it in turn. The development of the theory of immediate constituents was based on two central observations. First, the structure was *hierarchical*, in that if we take a grammatical sentence like *the turtle saw the horse*, we can expand either the subject or the object into indefinitely large constituents (by adding relative clauses and the like). Second, when there are grammatical dependencies between nearby words or constituents, more often than not the dependencies can be analyzed as holding between adjacent constituents that are themselves the immediate constituents of a larger constituent. In French, the choice of determine *le* or *la* “the” depends on the gender of the following noun, and its distribution is best described as being the first of two immediate constituents that form the noun phrase in French.

If we turn then to phonology, we find nothing corresponding perfectly to the way in which sentences are found embedded within one another in syntax; if we find hierarchy, it is by way of a sequence of essentially different kinds of constituents (syllable, foot, word, and various larger phrases), rather than the recursive structure we find in syntax. We find some distributional dependencies, but the dependencies are at least as different from those found in syntax as they are similar to them; we return to this question shortly.

What is most important about constituent structure in syntax, then, does not arise in phonology; and the best-grounded fact about the syllable – that it describes an interval of rising, then falling sonority – has no natural counterpart in syntax; there is nothing which rises and falls in a syntactic sentence, even if it is of the shape SVO (i.e. *Subject-Verb-Object*) let alone if it is VSO or SOV or anything else, as there is in a phonological syllable.

1.16 *Phonotactics: Patterns in Sequences*

One of the goals of the phonological analysis of a language is to determine, once we have established the set Σ of underlying segments of the language, what sequences of such segments are permitted. What sequences – that is, what subsets of Σ^* – are found in the language, and what subsets of Σ^* are *not* found? The phonologist’s goal is to offer meaningful generalizations about the answers to both of those questions.

Phonologists have long hoped to simplify the overall phonological description of possible words by viewing words as being built up out of smaller phonological units, notably the syllable. Pike (1947) emphasized the importance of the syllable in dealing with this observation, noting that the syllable was:

the basic structural unit which serves best as a point of reference for describing the distribution of the phonemes in the language in question. (p. 144)

Not all phonotactics involve reference to the notion of syllable, however. Scholars have argued that at least some phonotactic statements are best viewed outside the purview of syllables (Lamontagne 1993; Blevins 2003; Kabak and Idsardi 2007). Blevins notes cases, for example, where a neutralization of laryngeal features on a consonant occurs before obstruents regardless of whether the two segments are in the same syllable or not. She suggests, in a similar vein, that in systems which require homorganicity of coda nasals to following obstruents, the restriction should be analyzed as a negative filter against a nasal that is specified for point of articulation and followed by any consonant (i.e. a negative filter against **[nasal, PLACE] Obstruent*) (p. 379), though we should recall Bybee's argument for the syllable, mentioned above, that the rule of nasal assimilation occurs in a fashion that is better treated by saying that the environment is across syllable-boundary.¹³

Accepting (as we have throughout this chapter) the simplification that consists of analyzing phonological representations as sequences of segments, we define the possible syllables of a language as a finite set σ of sub-strings of Σ^* , and we define a phonological word as a sequence of elements of σ . This touch of formalism is intended to make natural the following question: what is the right way to characterize the set σ ?

It is possible to focus one's theoretical account on a compact set of statements regarding what segments may follow what segments: in effect, to offer a finite state automaton as our model, where each edge of the automaton generates a segment of the language. On this account, we have a formal device that generates, as we might say, "left-to-right," – that is, a directed graph, whose paths correspond to all and only the possible sequences within a syllable. Possible clusters of three or more consonants have been analyzed in several languages as consisting of exactly those that can be analyzed as overlapping pairs of permitted two-consonant sequences (that is, $C_1C_2C_3$ can exist if and only if C_1C_2 and C_2C_3 can independently exist), and it is natural to interpret this kind of analysis as employing a finite state model.¹⁴

In studies of finite state automata (FSA), it is common to consider two ways of thinking of a string as being generated by a particular path through such an automaton from beginning (#) to end (#): either *states* are associated with symbols which may be emitted when a path goes through a state, or else *edges* (from one state to another) are associated with symbols which may be emitted when a path follows a particular edge. In Figures 6.6–6.11, I have illustrated the former style of FSA for languages with simple syllable patterns, while in Figures 6.12 and 6.13, we see the latter style of FSA, to which we now turn.

A very early effort at describing in detail the sequential structure of English syllables was made by Benjamin Lee Whorf (1940), whose goal was to illustrate the complexity of the implicit knowledge of any native speaker of English. He presented it in a format which is similar to, but by no means identical to, a finite state automaton (and he was working before the notion had entered the literature). I have therefore modified it a bit to look more like a familiar finite-state diagram. The part covering the syllable onset is given in Figure 6.12; all paths lead to a




Figure 6.6 A monstrosity? Recursive constituent structure.




Figure 6.7 CVC syllables.




Figure 6.8 CV(C) syllables.




Figure 6.9 CV(L) syllables.




Figure 6.10 (C)V(L) syllables.




Figure 6.11 C(L)V(L) syllables.




Figure 6.12 Whorf's onset.

single state (though there is an added complexity involving the *yu* diphthong), and this convergence explicitly represents the lack of dependencies between the choice of onset and the following nuclear vowel. In Figure 6.13, his rhyme is represented in a similar fashion. The intention is for there to be a one-to-one association between paths in the graph and possible onsets in English (though the graph fails to generate the cluster *kl*, apparently an overlooked flaw). Each edge is associated with a set of segments, and one such symbol is generated when taking an edge from one state to the next, if there are segments associated with the edge; if there are none, then passage along that edge does not contribute any symbol. Several edges are associated with the null set of symbols (equivalently,




Figure 6.13 Whorf's rhyme.

are associated with the null symbol), though this is essentially a characteristic that I have had to add in redoing Whorf's notation. In one place, an unusual notation is necessary: one edge is labeled as generating any consonant *other* than the velar nasal η .

An analysis of this sort does not indicate any hierarchical structure, but it obviously is one which contains a lot of structure: the structure inheres in the statement of permissible sequences, rather than in the representation of any particular sequence. In modern parlance, it is *non-deterministic*, in the sense that from several states, more than one path leading from the node can generate the same symbol (for example, a syllable-initial g can be generated by taking the top-most path, or by taking the third highest path, the one that generates any single consonant other than η). Such formal devices are in some ways better suited to express permissible sequences; the range of their abilities is different from that of phrase-structure rules. In addition, a model of this sort can easily be made probabilistic, and to illustrate a first-order Markov model. This allows us to easily indicate the difference between the probability of (for example) a syllable-initial p and a p that immediately follows an s .

1.17 Onsets, Codas, and Word-appendices

Phonologists have long hoped to simplify the overall phonological description of possible *words* by viewing words as being built up out of smaller phonological units, either syllables or, in more recent work, feet (which are, in turn, composed of syllables). After all, a rough account of possible sequences of sounds in words can often be formulated by saying that a phonotactically admissible word is one

that can be analyzed as a sequence of phonotactically admissible syllables. On this view, there should be perfect agreement between what consonant sequences can occur word-initially and syllable-initially, just as there should be perfect agreement between what consonant sequences can occur word-finally and syllable-finally.

However, it has been known for a long time that this approach is much too simple to account for the facts of language. The effort to reduce well-formed words to sequences of well-formed syllables runs into serious trouble when either of the following holds:

- (i) the set of consonant strings (*C-strings*) that appear between vowels turns out not to be identical to the set of all sequences that we obtain by concatenating a word-final *C-string* plus a word-initial *C-string*; or
- (ii) if we have some clear way of determining where the break is between syllables – let us suppose that we have an inter-vocalic sequence *C* and we can determine that it is broken between syllables as C_1 and C_2 , where $C = C_1C_2$ – and either C_1 does not occur word-finally, or C_2 does not appear word-initially, or both.

Both of these situations have long been known to exist. Languages exist, for example, in which no more than two consonants appear word-initially, and no word-final consonants, but word-internally, sequences of three consonants are found (we will see examples below). Inconsistencies may arise in the other direction: words may begin or end with sequences that are not permitted (as onsets or codas, respectively) word-internally. Similarly, while Dutch permits *str*, *spl*, and *spr* clusters word-initially, such clusters are split *s-tr*, *s-pl*, and *s-tr* when they occur word-internally, to judge from the laxing that occurs in closed syllables in the case of words such as *mistral*, *esplanade*, and *Castro* (Trommelen 1983; van Oostendorp 1995). Or again, word-initial clusters such as are found in Dutch *gnoom*, *slaaf* and *tijftjaf* are syllabified in separate syllables when the sequences are found word-internally (van der Hulst and Ritter 1999b).

These problems have been attacked in several ways. If word-initial combinations appear to be more numerous than other syllable-onset combinations, then an extra word-initial position could be proposed, and similarly, if word-final sequences are richer than syllable coda sequences, extra word-final consonant positions could be proposed. These positions are often called *appendices*. (In addition, other considerations may be brought to bear, notably the difference between what consonantal material may appear in the onset and coda of a stressed syllable and of an unstressed syllable.)

The strongest argument that has been made for approaches that employ language-specific *appendices* is based on the observation that word-final *appendices* are also inert or invisible with respect to measures of syllable weight, which in turn are relevant to stress assignment (see Chapter 5). It is also interesting to note that word-initial and word-final *appendices* often violate sonority sequencing generalizations that hold word-internally, though this observation may carry somewhat less weight (van der Hulst and Ritter 1999b: 16). In addition, it is found that in some

(perhaps most) languages with closed syllable shortening, with only a single consonant permitted in the coda of a non-final syllable and two consonants permitted in word-final position, a long vowel is permitted in a word-final syllable just in case it is followed by one, and not by two, consonants – all of which suggests that the final consonant in such a system is parsed as a word-appendix rather than the coda consonant of the final syllable. See Vaux and Wolfe (2009) for a recent in-depth look at this subject.

In some cases, there appears to be a clear link between appendix-like behavior and morphological status. Of this, English provides a simple example: except in morphologically specific cases, everywhere in English a coda nasal will be homorganic with a following obstruent in the same coda, as in *bank* and *slump*. But before the regular verbal suffix *-d* and the plural suffix *-z*, we find violations of this (*flamed*, *banged*, *flames*, *bangs*).

This general situation was noted as early as 1952, in an important survey (Fischer-Jørgensen 1952), where Fischer-Jørgensen noted that there are languages in which “some medial clusters cannot be dissolved into actually occurring final and initial clusters” (she cites Italian, Totonaco, Chontal, Yuma, and Kutenai), adding that

it is evident that the phenomenon is not rare. But generally these cases are exceptions, even within the system of the language in question . . . But there are very extreme cases of this phenomenon . . . Finnish constitutes a good example. In Finnish the only consonants admitted finally are *n*, *r*, *l*, *t*, *s* and initially genuine Finnish words have only one consonant; but medially a great diversity of clusters is found, e.g. *ks*, *rst*, *mp*, etc. (p. 306)

She notes the even starker example of Keresan, in which words obligatorily begin and end with consonants, but CVCVC and CVCVCVC are fine word-patterns, leading the analyst to the conclusion that while CV syllables are permitted word-internally, such an open syllable cannot appear word-finally.

Some work has tried to tackle the analysis of word appendices as supernumerary onsets (as word-final appendices) or codas (as word-initial appendices), a view that is very close to the proposal that these edge effects are the result of syllables that are degenerate, in the sense that they do not contain a vowel, something that is typically taken to be an *sine qua non* for a syllable, after all. Perhaps the earliest example of this is cited by van der Hulst: he notes that Kuryłowicz (1952) treats this situation as involving a stranded onset. On word-peripheral clusters, see, for example, on Dutch: Trommelen (1983), van der Hulst (1984); on Polish: Cyran and Gussman (1999), Rubach (1990), Rubach and Booij (1990a), Davis (1990); a good survey appears in Törkenczy and Siptar (1999). Kiparsky (2003b) offers an interesting account of striking differences in modern Arabic dialects based on differences in the ranking of a constraint requiring moras to be licensed by syllables; where the constraint is violated, structures are found in which phonological material appears despite it going well beyond the range of possibilities permitted by Arabic core syllables.

It has long been noted that fewer consonants and consonant sequences are permitted in codas than in onsets, though this observation has eluded precise formulation. One of the challenges to dealing adequately with this phenomenon arises from the fact that in many cases, it is not so much a segment *type* that is excluded from the coda as a segment *contrast*. In what seems to be positive terms, we can say that the coda is a position of neutralization; in negative terms, we can say that a certain segment type cannot appear in the coda unless it is the result of a generalization that expresses a neutralization! The central example of such phenomena is the appearance in many languages of geminate consonants intervocalically: in a form like Italian *detto* (Figure 6.2), the coda of the first syllable has a [t] in it, but this is possible only because the following consonant, in the syllable onset, is a [t] – which is to say, a geminate can give rise to a structure in which an obstruent appears in a syllable coda, even in a language which does not permit, in non-geminate cases, a coda obstruent. This kind of situation has been described in terms of the logic of licensing; see Itô (1986), and, for a slightly different perspective, Goldsmith (1989).

2 Syllable-based Alternations

It is a commonplace to find in the phonological literature descriptions of phonological alternations of consonants in which the crucial context is the syllable position of the consonant: it is realized in one way when in the syllable onset, and another way when in the coda. Things are often not *quite* that simple, but that serves as a central focus of a wide range of phenomena, as illustrated below.

2.1 Spanish *s*

The behavior of /s/ in New World Spanish dialects illustrates a common pattern in which a consonant is realized differently in an onset and in a coda. /s/ in onset position is realized as [s], but in many dialects, /s/ in coda position is realized as [h], which is to say, as an [s] whose oral frication is removed. This phenomenon, often referred to as *aspiration*, is widespread in New World Spanish, and attested in Peninsular Spanish. In dialects with aspiration, *esta* “this” is [ehta]; there is considerable dialectal variation with regard to the behavior of word-final /s/.

2.2 French *loi de position*

Most dialects of French have six oral mid vowels: ϵ , e, o, ɔ , ø , œ . Of these, three are open (ϵ , ɔ , œ) and three are close (e, o, ø). There are conditions on their distribution, however. A close mid vowel cannot appear in a closed syllable. For the pair (e, ϵ), we see this effect in morphophonemic alternations as well as in many loanwords from English, where a tense and diphthongized vowel in English [e^v] is borrowed as [ɛ] in closed syllables, e.g. [tɛk] English ‘teak’, [mɛl] ‘e-mail’, [stɛk] English ‘steak’, [kɛk] English ‘cake’, and in truncated forms in contemporary

speech, as in [tɪdɛʒ] from [pətɪdɛʒəne] ‘breakfast’, or [agrɛʒ] from [agregasjɔ̃] ‘teaching certification’. In Québécois French, this relationship is extended to high vowels as well: (i, y, u) appear in open syllables, and (ɪ, ʏ, ʊ) appear in closed syllables, as well as in syllables to the left of a closed syllable to which regressive vowel harmony has applied (e.g. [dzɪfsɪl] ‘difficult’). This is attested in pairs of related words, such as [ptɪsɪ] [ptsɪt] ‘small’ masc., fem.

2.3 Flapping in English

Trager and Bloch’s classic, but controversial, analysis of English phonology (1941) argues that in the case of word-internal stressed-unstressed sequences, such as *bidding*, *bedding*, *padding*, *nodding*, *budding*, and *pudding*, each of “the six short vowels is followed by an ambisyllabic voiced stop” (p. 233), and in a curious appendix to a very brief note in *Language* (see Eliason 1942), they support their view of the proper syllabification of English words such as *hitting*: “the division here is not before the medial consonant and not after it – in short, that the consonant is ambisyllabic, and that the division occurs, if at all, within the consonant itself” (p. 146). Kahn (1976) developed an analysis of flapping in American English which strongly supports their analysis.¹⁵ The following analysis departs from Kahn’s analysis in some of the specifics, but follows it in overall construction.¹⁶

One of the most striking characteristics of American English is the widespread appearance of the coronal flap [ɾ] as a realization of /t/ as well as /d/, and the principles that lie behind the distributional generalizations of the flap have led linguists to view the conditioning to be based on syllable affiliation: a coronal stop /t,d/ is realized as a flap if it is simultaneously a member of the coda of one syllable and the onset of the next syllable.

It is not hard to find phonological descriptions of American flapping which state a generalization along these lines: a coronal stop is realized as a flap when it is immediately preceded by a stressed vowel, and immediately followed by an unstressed vowel, in *Italy*. While this is true, it is only a small part of the story. In the real description of American flapping, it is first of all necessary to distinguish the conditions under which word-internal /t/s are flaps from the conditions under which word-initial /t/s and word-final /t/s are. Consider first the case of strictly intervocalic, word-internal /t/s, where the facts roughly follow the description just given. The context *ɪ̃–ɪ̃* mentioned there is, in fact, a position in which flap obligatorily appears: for example, *Italy* [ɪ̃ɾəlɪ̃]. There are three other strictly intervocalic contexts to consider: *ɪ̃–ɪ̃*, *ɪ̃–ɪ̃*, and *ɪ̃–ɪ̃*. In the first two, we do not get a flap at all; it is not possible in words such as *bótɔ̃x*, *détàil*, *rétàil*, *lâtèx*, *Útàh*; or *Ítálian*, *áttáin*, etc. In the third case, where the /t/ is surrounded by unstressed vowels, as in *sanity* or *opacity*, both flapped and unflapped variants are possible (they are equally acceptable to this writer). The generalization does not change (here as elsewhere) when we extend the context to include a preceding *r*; *parting* and *potting* have flaps just the same. The same is not true of other sonorants: a /t/ will not flap after /l/; we have *faulty* with no flap possible, for example. (There is a complication when a syllabic *n* follows the /t/, as in *Latin*, which we will ignore here.)

Word-final /t/ may always be pronounced [t̚], a glottalized and unreleased [t], associated with at least a weak phrasal boundary immediately following, but in connected speech, when the following word begins with a vowel, regardless of whether the vowel is stressed or unstressed, a flap is found – and this is true whether the vowel preceding the /t/ is stressed or unstressed. Examples of the four cases, where the [r] realizes a /t/: *Gé[r] ůs out of here! Gè[r] óut of here! A lockě[r] őf hair. The rabbi[r] áte the carrot.* In these cases, an empty onset attracts a preceding consonant, even if the consonant is “already” syllabified. The result of this is that the /t/ is ambisyllabic, and ambisyllabic /t/ś are flapped. This is not a case of maximizing onsets; a /t/ does not resyllabify before an *rV* sequence, for example (*the ha[t̚]* *ripped*, with no possibility of the /t/ being part of an onset-cluster).

The third case, that of the word-initial /t/, depends, curiously enough, on the particular word in which it is found. If the word is *to*, *today*, *tonight*, *tomorrow*, or *together*, then we find one behavior, which I will temporarily refer to as *to*-behavior; if the word is any other (for example, *tomato*, *tuba*, *Topeka*, *topology*, *Thomas*, *taste*), we find a different behavior. Furthermore, the realization is largely independent of whether the preceding vowel is stressed or not, and largely independent of whether the following syllable is stressed or not. There is no flap in *the tomato*, *a tenacious opponent*, or *a topology*, where the /t/ is in an unstressed syllable, nor in *the total* or *the toast*, where /t/ is in a stressed syllable. Thus this case is entirely different from either the word-internal or the word-final case.

However, in the case of the handful of words based historically on the preposition *to* (*to*, *today*, *tonight*, *tomorrow*, and *together*), the facts are different. In each case, flapping is possible (indeed, preferred) when the preceding word is vowel-final: *Go [r]o sleep! How'd it go [r]oday?*, etc.

A natural way to interpret this data involves two passes of syllabification. The first applies word-internally, syllabifying a segment to an immediately following vowel, regardless of stress, and a rule that adds a syllable link between any open syllable and an immediately following consonant in the word (optionally if the syllable is unstressed, and obligatorily if the syllable is stressed). This results in an ambisyllabic consonant. At the phrase level, only one rule is operative: a word-final consonant adds an affiliation to a following syllable σ if σ begins with a vowel in the same phonological phrase. That rule also results in an ambisyllabic consonant. Given these two rules, we may say that any, and only, ambisyllabic /t,d/ is realized as a flap [r]. The *to*-initial words that we noted above are all cliticized to the word that precedes, in the sense that it is treated as a single phonological word with what precedes it.

3 Conclusions

What can we conclude about the syllable, in the light of the studies that we have reviewed? There are repeating patterns of sequences of sounds in language, and these patterns define the syllables of various languages, and these patterns lie at

the base of many, or all, prosodic phenomena. But how are these repeating patterns best described, and how are they best explained? We have an embarrassment of riches in facing both of these questions. In this final section, we will reflect a bit on how this is so, and what we might do about it.

Phonology, as a field, is still struggling to deal with the consequences of the development of the phoneme, which is at the same time its greatest achievement. By “the phoneme,” we mean the abstract characterization of a set of sounds in a language which unifies all of the sounds into a relatively small inventory of elements which are then used to define contrasting morphemes and words. This insight is the beginning of all work in phonology. Yet at the same moment, two other types of analysis – analysis into syllables, and analysis into features, or in short, analysis into units both larger and smaller than the phoneme – are crucial for any descriptive account of the phonology of a language.

On the one hand, languages never offer the unconditional sequence of any phoneme followed by any phoneme: local conditions of dependence are present everywhere. Using the phrase “local dependence” suggests that a first-order Markov model might be a good model of phoneme occurrences: segments do indeed care very much what their neighbors are, so to speak. But mapping out the conditional probabilities of each phoneme, based only on what phoneme precedes, fails to capture the just slightly larger generalization that not just lurks, but looms, behind the data: to wit, that while many languages permit sequences of two consonants, very many exclude sequences of three. We could expand our vision to a second-order Markov model, allowing each phoneme’s options to be limited by the two phonemes that precede it, but we would be losing sight of the bigger generalizations. That is, if there are p different phonemes in the language, there are p^3 different parameters that need to be specified for a second-order Markov model: each phoneme’s probability after each pair of phonemes would need to be specified. But any study of a real phonology shows us that only a small portion of the universe of p^3 possibilities has a chance of being utilized by a natural language phonology, *because* there are generalizations just slightly larger in scope.

These generalizations involve what we call the syllable. But how should these generalizations be modeled and formalized? We have seen three major traditions over the course of this chapter, the syntax-based immediate constituent approach, the sonority approach, and the finite-state approach. The first specifies constituents of structure and utilizes phrase-structure rules to describe possible sequences, the second maps each element of the phonemic inventory to the real numbers, and then reconstructs conditions on possible numerical sequences, such as limiting which phonemes can appear at local peaks of sonority, while the third focuses its theoretical capability on a statement of what sequences are permitted in a given language.

On the other hand, syllabification is not simply an effect, of which the sounds are the cause: quite to the contrary, the choice of phoneme in some cases, and the choice of allophone in a very large number of cases, is determined by the location of a sound in the larger prosodic stream. Of this, the most striking special case is the difference in the realizations of consonants in syllable onset and in syllable coda.

And yet clear evidence of constituent structure in phonology is notoriously difficult to establish, certainly compared to the ease with which we can determine that choice of the allophone (realization of a phoneme) is conditioned by the immediately following phoneme, and compared to the ease with which we can distinguish between the characteristics of a consonant in the coda and in the onset of a syllable.

My conjecture is that the syllable is ultimately best regarded as the lowest level (or one of the lowest levels) of rhythmic reoccurrence of possibilities in language. Some might want to see this as the reflection of gestures made by the articulatory apparatus, a view that we have not surveyed in this chapter. For myself, I think that such a view analyzes language at the wrong level of abstraction or granularity: the correct level of abstractness for the description of language is higher than that of jaw gestures. Sonority, and the wave-like recurrence of peaks of sonority, seems to me to be the fundamental pattern of syllabification in language.

Studies that explore the consequences of optimality theory for our understanding of the syllable, and vice versa, seem to me to largely miss the point that we have discussed in this chapter, and in a sense that should not be surprising, in view of what optimality theory is: it is a theory of constraint interaction, rather than a theory of phonological representation, and it is not fundamentally a theory of how the constraints (appropriate for natural language phonology) should be formulated, even if some phonologists have implicitly, or on occasion explicitly, made some suggestions along such lines. Optimality Theory is perfectly consistent with any of the three views described here.

We began this chapter with a quotation from Ernst Pulgram, and we will end it with another. Pulgram wrote,

[The syllable] has no function, no *raison d'être*, apart from that of the syllabic segmentation of an utterance. It serves nothing but itself, as it were; it does not serve, immediately like a sign or mediately like a figura, the communicative purpose of a language . . . A syllable is . . . a phonological unit that is, as all linguistic units must be, describable and definable only on its own level of analysis exclusively.¹⁷ (p. 21ff.)

Perhaps that is the best we can do for now. But I think that it is not the last word to be said on the subject. The most important question to answer is how to develop a model that is suited precisely to capture the rhythmic character of syllables, and the striking asymmetries of onset and coda. We have amassed a great deal of knowledge in recent decades that will help us reach that goal eventually.

NOTES

- 1 I thank Peter Auer, Juliette Blevins, Diane Brentari, Stefano Canalis, Stuart Davis, François Dell, Tracy Alan Hall, Harry van der Hulst, Jonathan Kaye, Paul Newman, Marc van Oostendorp, Jason Riggle, Donca Steriade, Bert Vaux, Ilya Yakubovich, and Alan Yu for comments on an earlier draft.

- 2 Hyman (1983). Auer (1994) criticizes the view that all languages are syllable-oriented, and explores the difference between syllable languages and word languages.
- 3 I have just distinguished between phonotactic regularities and conditioning regularities, but both types of regularities describe conditions on what segments may follow one another in a language. The distinction between the two rests on the assumption that it is possible to specify, for a given language, an inventory of underlying sounds categories, the traditional phonemes or underlying segments. In the context of this book, that assumption is not controversial. Phonotactic regularities are then statements as to what sequences of phonemes are permitted in a language, while the other *conditioning* regularities are the statements in a phonology about what phonological elements may cooccur with what; the largest part of this is composed of the rules of allophony, that is, non-neutralizing rules. Syllable position plays a major role in conditioning the realization of an phonological segment. For example,
 - The realization of a consonant is frequently different depending on whether it is in a syllable onset or syllable coda.
 - The realization of a vowel is frequently different depending on whether there is a consonant in the immediately following coda or not.
- 4 The title of Sevald *et al.* (1995) gives a good flavor of this perspective: "Syllable structure in speech production: are syllables chunks or schemas?" as does the title of Cutler *et al.* (1986): "The syllable's differing role in the segmentation of French and English."
- 5 I am indebted to Bernard Laks as well as to Ali Tifrit for bringing the importance of this material to my attention. See Laks (2003) and Tifrit (2005).
- 6 It is worth bearing in mind that there *was* in fact a frontier of research in syntax at the time, and that notions that many of us take for granted today were being developed and argued about during the 1940s, and also that it was not a *discovery* that the structure of the syllable matches that of the sentence within a constituency-based theory of the syllable, since this theory was specifically created in order to have this appearance. Canalis (2007) discusses the influence of the work by Hjelmslev, with Uldall, in the framework of glossematics, noting that Hjelmslev (1939) should be cited in the development of the constituency view of the syllable. The history of this period needs to integrate Hjelmslev's influential work.
- 7 See also Haugen (1956), and the position of Clayton (1976), Shibatani (1973), and Hooper (1976).
- 8 One of the earliest explicit discussions of the significance of identifying *one* of constituents of an immediate constituent as the nucleus or head, and the other as non-head or satellite, is Pittman (1948), though he does not discuss phonology *per se*. A number of phonologists have explored this asymmetry over the last 20 years; some of them have worked in a framework influenced by government phonology, though not all have. This question is discussed in Chapter 16.
- 9 In fact, Microsoft recently patented this idea, or perhaps just something very close to it; see US Patent 20050203739, granted in September 2005. In a context such as this, the mutual information between two adjacent segments s_1 and s_2 is $\log \frac{\text{prob}(s_1, s_2)}{\text{prob}(s_1)\text{prob}(s_2)}$. See Goldsmith (2007), Goldsmith and Riggle (2012).
- 10 Pulgram presents these rules as ordered, but this reader gets the impression that his use of ordering is essentially for the purpose of indicating that a later rule has empirical precedence over an earlier one.

- 11 See the introduction in Hartmann *et al.* (2008) for a general discussion of empty categories.
- 12 Some offer this account as a success of OT, as in Féry and van de Vijver (2003: 8), where the authors take a different view: "The ability of OT to explain typological patterns as a result of the interaction of markedness and faithfulness constraints is the core of the theory, and it is to a great extent responsible for its success" (p. 8). Others might ask for an account of why the constraints are NoCODA and ONSET rather than CODA and NoONSET, or CODA and ONSET, or NoCODA and NoONSET.
- 13 See also Steriade (1999a).
- 14 See Fischer-Jørgensen (1952), who cites Bjerrum and Hjelmslev, though she gives counter-examples from Russian and Kutenai; see also Clements and Keyser (1983), as well as Hjelmslev (1939) and Butt (1992).
- 15 Some scholars were unpersuaded that the facts should be described with terms including *ambisyllabicity*; see, for example, Picard (1984), who does not appear to be familiar with the historical depth of this view, viewing it rather as an artifact of Kahn's design.
- 16 Kahn's analysis crucially involves ambisyllabicity. Such an approach has been challenged by, among others, Kiparsky (1979) and Selkirk (1982); see also Harris and Kaye (1990), Hammond (1997); also Rubach (1996), Jensen (2000) and Picard (1984). Alternatives to the ambisyllabicity approach need to appeal to using feature specifications on a segment to give it a mark indicating its syllabic position earlier in the derivation.
- 17 Pulgram's remark actually calls to mind Mark Aronoff's recent argument (1994) that what he calls *morphemes* in language have an existence that is, in many cases, for themselves and only for themselves: they are more concrete than morphemes, and play an important role in the morphologies of many languages.