Features, segments, and the sources of phonological primitives*

Abigail C. Cohn
Cornell University

I review the role of distinctive features in early generative theory, focusing on their multifaceted role as defined by Chomsky & Halle (1968) for capturing contrast, phonological patterns, and the phonetic realization of these patterns. Based on evidence from these multiple aspects of phonological representation, I conclude that the characterization of segments as combinations of universally-defined distinctive features is approximately, but not literally, correct. This leads to the question of how young children learn the elements of their phonology to achieve the knowledge of an adult phonological system? Crucially the evidence suggests that *how* we learn is not the same as *what* we know. Rather, an approximately categorical and compositional system is learned out of a more continuous, gradient one.

1. Introduction

When we consider the structure of adult phonological systems, we find ample evidence for the role of both segments and grouping of sounds that pattern together, characterizable in terms of their shared phonetic properties, so-called natural classes (more neutrally termed *phonologically active classes* by Mielke 2008). Distinctive features are the dimensions that characterize these groupings, or "the recurrent elementary components" (Clements 2009: 19), typically modeled as binary parameters.

Both segments and features are widely observed in the phonologies of the languages of the world, and it is common to take one or the other as "primitives"

* This is a revised version of a paper presented at the 2007 conference: Where Do Features Come From? Paris, France. Thanks to the conference audience, editors of this volume, two anonymous reviewers and students and colleagues at Cornell University for very helpful input which I hope has led to a stronger final result. I especially thank Nick Clements for many fruitful discussions about topics touched upon in this paper.

of human language. Yet it is worth considering seriously, as the contributions to this volume do, the following questions: Where do these elements come from? What is their status both as part of the adult grammar and as a mechanism of acquisition?

Implicit in the International Phonetic Association's chart of phonetic symbols (the International Phonetic Alphabet or IPA) is that the elements transcribed by the IPA are a set of potentially available sounds or phones. But the IPA chart in and of itself does not address what the status of these elements is. On the other hand, Chomsky & Halle (1968), in *The Sound Pattern of English (SPE),* explicitly take the view that distinctive features are universally available because they are part of the innate linguistic endowment of humans and therefore understood to be part of "universal grammar". An alternative line of discussion, developed most explicitly by Mielke (2005; 2008), is that evidence from natural classes for the innateness of features is not as clear cut as usually assumed, and that observed patterns can be accounted for through an "emergent" approach to sound structure.

In this paper, I review the role of distinctive features as assumed from early work by Jakobson and colleagues (e.g. Jakobson, Fant, & Halle 1952 *Preliminaries to Speech Analysis*, *PSA*) and focus on the multifaceted role of features as defined in *SPE* for capturing contrast, phonological patterns, and the phonetic realization of these patterns. The evidence from language-specific phonetics as well as phonological patterning highlights the point that there is close similarity, but not identity, between phonological categories, groupings of phonological categories, and the phonetic dimensions that characterize them across languages. This leads to the conclusion that the characterization of segments in generative phonology as combinations of universally-defined distinctive features is approximately, but not literally, correct.

This leads us to consider the relationship between adult phonological patterns and child language acquisition. How do infants or young children learn the elements of their phonological system to achieve the knowledge of an adult phonological system? I argue that crucially, *how* we learn is not the same as *what* we know. As discussed in many of the contributions to this volume, an approximately categorical and compositional system can be learned out of a more continuous, gradient one.

In short, we need to reconcile the fact that while the adult phonology looks like a categorical system built out of elements including segments and phonological features, it is likely that a significant part of what has been widely assumed to be innate is instead learnable. In an effort to provide a historical context and framing of these issues, I turn to the nature of phonological primitives in §2, and to the relationship between adult phonology and its acquisition in §3 and conclusions and future directions in §4.

2. The nature of phonological primitives

What are the primitives of human speech? Where do they come from? Typically when we characterize and analyze patterns of adult phonological systems, the vocabulary we use is cast in terms of segments and features. This vocabulary suggests that such categories and groupings are equivalent across languages. When we use these tools to practice our trade, we often assume that these elements are primitives of the system. But what does this mean? What is the relationship between the acoustic and/or articulatory dimensions that provide the substance of human speech (often characterized as either phonetic features or distinctive features) and the contrastive elements that constitute the building blocks of meaningful communication (often characterized as segments or phonemes)? There are four possible types of approaches to understanding this relationship: (i) segments as primitives, (ii) features as primitives, (iii) neither as primitives, (iv) both as primitives:[1]

i. Segments as primitives

Most American structuralist approaches to phonology took as a given the phoneme (see Anderson 1985 for a careful review of different understandings of the phoneme).[2] Since most American structural approaches were explicitly not theories of the mind, it is not clear exactly what it would mean to take something as a primitive. However a notable exception to this view was that of Sapir (e.g. Sapir 1963) who explicitly took phonemes to be mental units of native speakers. Following this sort of view, it could be argued that natural classes and therefore evidence for features can be learned from the patterning of segments. The importance of phonological grouping of phonemes is certainly implicit in Sapir's treatment.[3]

1. Clements (2009:19) draws a distinction between a "feature-mediated theory of inventory structure" and a "direct-access theory of phonological explanation". These roughly correspond to types (ii) and (iii), respectively.

2. Bloomfield (1933:79–80) considers phonemes to be made up of "lumps or bundles" of certain gross acoustic features which are distinctive, but these features have no independent existence.

3. For example, this is seen in Sapir's discussion of Nootka. Native speaker intuition leads to grouping glottalized stops and affricates together with glottalized nasals and semivowels, despite distinct phonetic realization and disparate treatment in a traditional orthography, highlighting their phonological grouping.

ii. Features as primitives

Generative phonology was built on the fundamental premises developed in Jakobson, Fant & Halle (1952: 3) in which a small set of "distinctive features" were posited and taken to be both universal and innate.

> Any minimal distinction carried by the message confronts the listener with a two-choice situation…. The choice between the two opposites may be termed distinctive feature. The distinctive features are the ultimate distinctive entities of language since no one of them can be broken down into smaller linguistic units.

This is also the view set forth by Chomsky & Halle (1968: 64):

> We take "distinctive features" to be the minimal elements of which phonetic, lexical, and phonological transcriptions are composed, by combination and concatenation. The alphabetic symbols that we use freely in the discussion below are therefore to be regarded as nothing more than convenient ad hoc abbreviations for feature bundles, introduced for ease of printing and reading but without systematic import.

This view is argued for by Archangeli (1988) and also underlies the views developed by Dresher (2008) and is widely assumed throughout the generative phonology literature (e.g. Halle & Clements 1983; Kenstowicz 1994; Clements & Hume 1995). Under this view, segments are complex units constructed as bundles of features.

Articulatory Phonology offers a compositional view of phonology, positing gestures (rather than either segments or features) as primitives. A gestural approach shares certain properties with segmentally and featurally-based approaches, but explicitly incorporates a dynamic model of implementation, offering a different resolution of a number of questions about the relationship between the contrastive elements of phonology and their implementation. (See Browman & Goldstein 1992; Goldstein 2003; Golstein & Fowler 2003; Goldstein, Byrd & Saltzman 2006 for discussion.)

iii. Neither segments nor features as primitive, but both as "emergent"

More recently a number of scholars have raised the question of the degree to which both segments and features can be learned, thus obviating the need to assume either is part of "universal grammar". Under this view, categories are learned, and phonologically active classes are derived through generalization based on phonetic similarity and analogy (see e.g. Blevins 2004; Port & Leary 2005; and Mielke 2005; 2008; Pulleyblank 2006). While some convincing evidence of statistical learning based on input has been documented for both infants and adults (see Maye, Weiss & Aslin 2008 for recent review), the question is whether these mechanisms alone

are sufficient to account for observed patterns in adult phonology. A number of contributions in this volume address this issue.

iv. Both segments and features are primitives

A fourth possible view is that both segments and features are primitives. Some interpret *SPE* as taking segments or phones (in addition to features) as primitives (Hale & Reiss 2000; Port & Leary 2005). They interpret *SPE* as defining possible phones of the languages of the world in terms of a multidimensional phonetic space predetermined by the system. Chomsky & Halle (1968: 5) posit discrete phonetic segments: "Suppose that universal phonetics establishes that utterances are sequences of discrete segments, that segments are complexes of a particular set of phonetic features…" But these are not primitives, as these too are made up of (phonetic) feature bundles.

There is another way in which both segments and features might be taken to be primitives: Cross-linguistic evidence supports the conclusion that adult speaker-hearers manipulate both segments and groupings of segments characterizable in phonological featural terms.

In the course of our discussion, I will take the view that none of these four approaches is correct. Evidence supports the conclusion that both segments and features are elements of adult phonology, but this does not address the question of whether they are "primitives". To understand the role of these structures in adult phonology, we need to start with the methodological premise that neither is a primitive of the system. That is, we need to start with an inductive baseline (Gildea & Jurafksy 1996; Hayes & Wilson 2008) against which we can add and test theoretical assumptions in order to understand how categories and groupings of categories are part of what we bring to the task of learning a phonology, and the degree to which they can be learned. Since the theme of this volume is the role of distinctive features, we focus our discussion on the role that features play in phonology and phonetics, and revisit the degree to which the widely accepted view of distinctive feature theory framed in *SPE* captures the necessary insights. We start by briefly reviewing the history of distinctive features and the fundamental assumptions about distinctive features in *SPE*.

2.1 The *SPE* view of the elements of phonology

Chomsky & Halle (1968) account for the definition of possible speech sounds, and additionally offer an explanation for natural classes, by positing a small number of elements or parameters that we are endowed with, defined as a universal set of distinctive features.

> The total set of features is identical with the set of phonetic properties that can in principle be controlled in speech; they represent the phonetic capabilities of man, and we would assume, are therefore the same for all languages. (294–295.)
> It is in this sense that the totality of phonetic features can be said to represent the speech-producing capabilities of the human vocal apparatus. (297.)

These phonetically defined properties are understood together to characterize the inventories and patterns in phonology. Much attention is given in *SPE* and since then to delineating the proper set, in terms of phonetic correlates and observed natural classes.

Chomsky & Halle's proposed feature theory grows out of earlier work by Jakobson and colleagues, which in turns builds on the idea of oppositions (contrasts) of Trubetskoy's (1939) *Grundzüge der Phonologie*. Jakobson, Fant & Halle (1952) propose roughly a dozen acoustically-based binary features. The roughly 26 binary features of *SPE* are a direct response to the *PSA* system. They are claimed to be both articulatorily and acoustically based, although the description in Chapter 7 focuses on articulation due to considerations of length.[4] However, the focus on articulation is interpreted by many to be a principled decision.

In *PSA*, the goal of distinctive features is to account for contrast. Distinctive features were explicitly understood as relational; thus a feature such as [flat] was understood to do multiple duty by capturing rounding, retroflexion, velarization, and pharyngealization – four properties claimed not to contrast within a particular language (see Anderson 1985: 123, for discussion). Distinctive features work in tandem with other kinds of features – configurational, expressive, and redundant – to account for observed patterns in language (Jakobson, Fant, & Halle 1952: 14–15).

There is a critical shift from *PSA* to *SPE*, in which the universally defined set of features accounted for more than the contrasts of phonology. Distinctive features characterize the contrastive elements of phonology as well as natural classes in their binary classificatory function. They are also the basis of phonetic implementation when translated into language-specific scalar values at the output of the phonology and then implemented by an automatic universal phonetic component. Chomsky & Halle (1968: 169) define the roles that phonological versus phonetic distinctive features play and the need for us to distinguish between them.

4. Chomsky & Halle (1968: 299) state "We shall describe the articulatory correlate of every feature and illustrate the feature by citing examples of its occurrence in different languages of the world. We shall speak of the acoustical and perceptual correlates of a feature only occasionally, not because we regard these aspects as either less interesting or less important, but rather because such discussions would make this section, which is itself a digression from the main theme of our book, much too long."

> We have used the term "phonetic distinctive features" for the universal physical scales that determine the rows of the phonetic matrices. Correspondingly, we may use the term "phonological distinctive features" to refer to the categories that label the rows of the phonological matrices. Unfortunately, the discussion and development of the theory of distinctive features has been confused by the use of the term "distinctive feature" in both senses. This is appropriate only insofar as the invariance condition is met – that is, insofar as the phonological rules simply add redundant information to lexical matrices, giving additional specification to archi-segments. As we have seen, however, this is not the case in general [….] In any event, the phonological and phonetic functions of distinctive features must be clearly distinguished.

Yet we often find a conflation of these two roles, encouraged by common usage still prevalent today, and also by explicit claims by Chomsky & Halle linking these two aspects of "distinctive features". For example, the dual roles of distinctive features were explicitly linked by Halle (1983) within a model in which "the abstract distinctive features constitute the link between specific articulatory and acoustic properties of speech sounds" (94) and where "distinctive features correspond to controls in the central nervous system which are connected in specific ways to the human motor and auditory systems" (95).

Thus, in evaluating the success of distinctive feature theory, we need to keep in mind both the posited phonological and phonetic roles of distinctive features. In terms of phonology, it is widely assumed that distinctive features should capture both contrast and phonological grouping seen in phonotactic patterns and alternations. The same set of features with scalar values is also assumed in *SPE* to provide sufficient information for phonetic implementation. An extensive body of literature has shown that distinctive features as envisioned in *SPE* and modeled by Halle (1983) are not sufficient for this latter task. We briefly review this conclusion in §2.2. We then turn to a consideration of distinctive features as an account of both contrast and phonological grouping in §2.3.

2.2 The implication of language-specific phonetics

Since *SPE* there has been much rethinking of the tight linkage between the proposed binary phonological distinctive feature set and the assumptions of an automatic phonetic implementation of the same set translated to scalar values (see Cohn 1990; 1998 for a review). Ladefoged (1980: 485) states:

> phonological features are certainly not sufficient for specifying the actual sounds of a language; nor are they in a one-to-one relationship with the minimal sets of parameters that are necessary and sufficient for this purpose.

This body of literature forces us either to decouple the phonological and phonetic roles of distinctive features or to understand distinctive features as universal in a more abstract way. We review briefly the evidence for language-specific phonetics and its implications for the *SPE* characterization of distinctive features.

One of the first cases to be carefully studied was vowel lengthening before voiced sounds (Chen 1970; Keating 1985). Patterns assumed to be universal were found to exhibit systematic differences in their implementation across languages, even when the phonological patterning was roughly the same. Similar evidence has been provided for patterns of intonation (Pierrehumbert 1980; Pierrehumbert & Beckman 1988), nasalization (Cohn 1990), and vowel-to-vowel coarticulation (Beddor, Harnsberger & Lindemann, 2002). These now widely-observed sorts of patterns led to a generally accepted view of *phonetic knowledge* (e.g. Kingston & Diehl 1994). Speaker-hearers know not only the phonological structure of their language, but also the fine phonetic details. Thus these details are neither universal nor automatic.

To highlight this point, consider the example of the vowels of English as compared to the vowels of Spanish as discussed by Bradlow (1993; 1995). Impressionistically we understand the five vowels of Spanish to be a subset of the inventory of English vowels, at least in the case of the high and mid front and back vowels, typically transcribed with the same IPA symbols in both languages [i, e, o, u] and characterized by the same distinctive features [high], [low], [back]. This is shown by Bradlow (1993: 2) in her comparison of the vowel charts for English and Spanish, reproduced here in Figure 1.

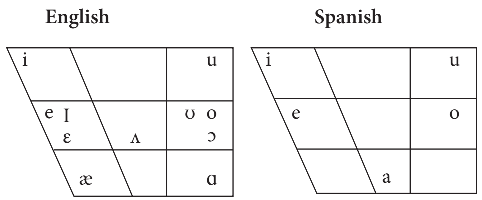

Figure 1. Charts of English vs. Spanish monophthongs, Bradlow (1993: 2, Figure 1.1)

Yet when we study these vowels more carefully by looking experimentally at their acoustic properties, we find that they are not actually the *same*, as shown in Figure 2, where Bradlow (1993: 34) compares the first and second formants of the vowels [i, e, o, u] in English, Spanish and Greek (leaving aside other differences in duration, nature of offglides, as well as individual speaker variations).

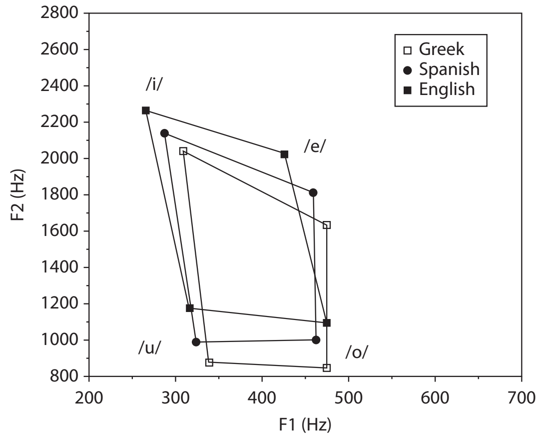

Figure 2. The /i/ - /e/ - /o/ - /u/ areas in English, Spanish, and Greek, Bradlow (1993: 3, Figure 2.7, four speakers, five repetitions averaged for each vowel for each language)

Thus, what we transcribe as [i] and represent as [+high, –back, –round] in one language is only approximately the same as in another language. (Lindau & Ladefoged 1986 also make this point.) The categories as well as the dimensions or parameters are similar but not identical across languages. In other words, the instantiation of the same distinctive feature specifications in different languages is not the same; nor can the difference be attributed to differences in scalar values. These categories are only definable in a roughly equivalent way. As stated by Pierrehumbert, Beckman & Ladd (2000: 285), the categories are language-specific in the sense that "there are no languages in which the implementation of analogous phonemes is exactly the same." The categories and grouping of categories show more variation across languages than predicted by a strict interpretation of an *SPE*-style implementation of universal distinctive features.

Such evidence from language-specific phonetics leads us to rethink both the role of distinctive features and the level of abstraction at which they can usefully be characterized. The *SPE* view of phonetic distinctive features does not provide a sufficient interface with phonetics, as there is not a universal set of phonetic elements (phones) equivalent across languages defined by the set of phonological distinctive features.[5] Thus the *SPE* view is not literally correct in the simplest sense.

5. Port & Leary (2005) take this point as one of their arguments for a blanket rejection of generative phonology and for a characterization of segments as epiphenomenal. This does not

But this limitation of distinctive feature theory is somewhat independent from the question of whether distinctive feature theory offers a useful account of phonological patterns of contrast and phonological alternations (though see Pulleyblank 2006 for discussion of this point). Is there a more abstract level of representation, or granularity (presumably either underlying phonological representations or surface phonological representations), at which featural characterizations of phonological elements and patterns are useful and at which categories and their grouping are equivalent across languages?

2.3 Distinctive features in characterizing contrast and alternations

We turn briefly now to the question of how well distinctive features characterize patterns of contrast and phonological grouping, by reviewing Clements' (2009) work on the role of distinctive features in capturing sound inventories and Mielke's (2008) consideration of distinctive features in capturing phonological alternations. These results suggest that there is a role for distinctive features in characterizing adult phonological patterns, but with a looser connection to their source than would be predicted by a universal innate set of features.

Critiques of the *SPE* set of features have led to a widely accepted, slightly modified set of features (e.g. Halle & Clements 1983); as well as hierarchical and grouping proposals (feature geometry, see Clements 1985; Clements & Hume 1995), and rethinking of processes, first in terms of autosegmental representations (e.g. Goldsmith 1976) and more recently constraint interaction (e.g. Prince & Smolensky 2004), while largely maintaining the fundamental assumption of universality.

Setting aside the finer details, we can ask how well distinctive features capture observed patterns both within and across languages. Clements (2003; 2009) has investigated this question by looking at the role of distinctive features in characterizing phonological inventories. He shows that distinctive features help capture a number of characteristics of both individual language inventories and inventories across languages, including evidence of Feature Bounding, Feature Economy, Marked Feature Avoidance, Robustness, and Phonological Enhancement.[6] He highlights the ways in which these properties are distinct from phonetic properties of dispersion, phonetic similarity, etc. For example, he shows that the predictions of Economy are quite different from either dispersion theory or gestural economy.

follow, however, as they do not fully consider other levels of abstraction at which features and segments may indeed play a role.

6. In a similar vein, Hayes (1999) and Gordon (2006) highlight the role of symmetry in the structure and organization of phonological systems, in contrast to phonetic patterns which exhibit gradient patterns.

These results support the conclusion that grouping sounds characterized in terms of phonological features provides insight into the inventory structure of phonology and suggest that this is a level of organizational structure available to the speaker-hearer.

Mielke (2008) sets out to evaluate how well distinctive features account for phonologically active classes. He compares the success of three different distinctive feature theories (Jakobson, Fant, & Halle 1952; *PSA*; Chomsky & Halle 1968 *SPE*; and Clements & Hume's (1995) Unified Feature Theory, *UFT*) against a database of phonological alternations. The database consists of 6,077 classes of sounds analyzed as targets or triggers of alternations described in 628 language varieties, culled from descriptive grammars. The main result is that "unnatural classes are widespread" (3). Of the three theories tested "no single theory is able to characterize more than 71 percent of the classes, and over 24 percent are not characterizable in *any* of the theories" (3, also 118, Table 6.2). Of the three, *SPE* fares the best. This is not surprising as it was devised in part to address empirical and structural problems with *PSA*, and as discussed by Mielke, the goals of *UFT* go beyond simple characterization of natural classes.[7]

Based on these results, Mielke argues that an innate distinctive feature theory is not adequate and duplicates other independently motivated explanations, including phonetic similarity and analogical change. He argues that emergent feature theory, where "features are abstract categories based on generalizations that emerge from phonological patterns" (9), offers a better account of natural and unnatural classes. Mielke also notes that we find similarity, not identity, across languages. The question is whether the continuum from "natural" to "unnatural" is indeed just that, or whether there is something interesting or special to say about those cases characterized as "natural" by *SPE* or another distinctive feature theory versus those that cannot be natural. Indeed, accounting for 71% of the data is a non-trivial result.

My reading of Mielke's results suggests that the distinction between "natural" and "unnatural" and the degree of similarity observed across languages requires more of an account than available from phonetic similarity and a general mechanism of generalization alone. As seen in Figure 7.6 (157) presented here as Figure 3, based on the *SPE* analysis, there is little interleaving of "natural" and "unnatural" classes, based on their relative frequency. Almost all the attested natural classes are of higher frequency than the attested unnatural classes and almost

7. For consideration of the success of the Halle & Clements (1983) feature set in this regard, see Mielke, Magloughlin & Hume (to appear).

all of the unnatural classes occur only once, suggesting that a useful or meaningful distinction might be drawn.

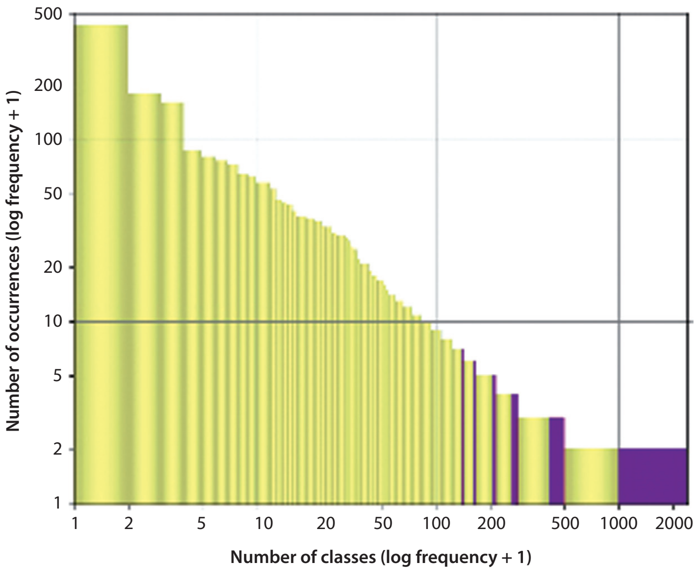

Figure 3. The distribution of frequent and infrequent natural and unnatural classes (SPE), Mielke (2008: 157, Figure 7.6). Reprinted by permission of Oxford University Press.

Second, Mielke (150, Table 7.1) characterizes the way that each feature theory accounts for the attested patterns, as "natural" (defined conjunctively by the particular feature set), involving a disjunction of 2 or more classes, or as being unnatural (even with a disjunction). Focusing on the results for *SPE*, we find the following percentages: Natural (feature conjunction): 71%; Disjunction (2 classes): 20.5%; Subtraction (2 classes): 1.2%; Disjunction (3–6 classes): 5.9%; Unnatural (even with disjunction): 2.7%. If we combine the results for Natural with Disjunction (2 classes) we find that *SPE* actually accounts for 91.5% of the phonologically active classes. Mielke (151–2) argues that treating "'unnatural' phonologically active classes with feature disjunction is not a point in their favor." Indeed, the conjunction of features rather than the disjunction of features is the very definition of a "natural" class: Halle & Clements (1983: 9) state: "'natural' classes can be specified by a single conjunction of features […]; 'unnatural' classes require a disjunction for their specification." However this raises the broader issue of the degree to

which "natural" vs. "unnatural" terminology is defined in a theory-internal way, and whether some of these disjunctions would be treated differently in a post-*SPE* system such as Halle & Clements (1983).

2.4 Distinctive feature theory as approximately correct

In reviewing the evidence of the role of distinctive feature theory both in phonetic implementation and phonological patterning, we are led to conclude that the *SPE* feature system is not literally correct as an innate universal endowment, but where does this leave us? Do we embrace what Joos (1957: 96) characterizes as "the American (Boas) tradition that languages could differ from each other without limit and in unpredictable ways"? No. Ample evidence supports the key role of categories and grouping of categories (phonologically active classes) in adult phonology. The issue is: where do these categories come from, how are they formed, and how equivalent are they cross-linguistically? What these results lead us to conclude is that in a technical sense, distinctive feature theory is *wrong*. But at a less fine-grained level, a well-defined set of distinctive features captures patterns within languages as well as strong similarities observed across languages. In other words, distinctive features are approximately correct.[8]

SPE established an important research agenda in both phonology and phonetics, and generative phonology has proved to be a successful enterprise in accounting for phonological structure, alternations, and commonalities across languages. But this alone does not address the question of the source(s) of distinctive features. The point is that unless we move away from a *literal interpretation* of feature theory as universal (and therefore innate), we will not make progress on understanding the nature of the linguistic endowment, since the literal interpretation predetermines the answers. On the other hand, we should not lose sight of the fact that generative models have proven to be excellent *approximations* of adult grammars. This is a non-trivial result when we consider the complexity of linguistic systems across the languages of the world. If we frame our discussion in terms of *right* and *wrong*, then we miss the opportunity to understand what is *almost right* about these models. We need to critique these models in a more sophisticated way. We need to understand their limitations, while also understanding the insight they offer. It may be that feature theory is right at a particular level of granularity; or it might be, as suggested by Pierrehumbert, Beckman & Ladd (2000), that feature categories capture the end state, but not how the system is formed.

8. This still leaves open a wide range of possible interpretations between assuming that features are little "more than a convenient set of labels" (Pulleyblank 2006) or a highly proscribed set.

This suggests we need to step back from the literal interpretation that distinctive features are universal because they are innate (what has been termed the "radical universal grammar (UG)" view) and take seriously the possibility of a "weak" rather than a "strong" UG. Hayes, Siptar, Zuraw & Londe (2009) show that both natural and unnatural constraints together account for the complex patterns of vowel harmony in Hungarian, and based on wug testing that speaker-hearers are sensitive to both. In modeling the role of this complex set of constraints, they show that natural constraints are disproportionately weighted over unnatural ones.[9] Recent experimental evidence supports this conclusion as well: both adults and infants can learn "unnatural classes" as well as "natural" ones (see Seidl & Buckley 2005 and Peperkamp & Dupoux 2007).

It is widely assumed that phonetic factors play a critical role in shaping category formation, that is, constraints imposed by the psycho-acoustics of the human ear and the nature of the human vocal tract. Much insight into how these constraints work is offered by Quantal Theory (Stevens 1972; 1989), as well as Dispersion Theory (e.g. Lindblom 1990), and Articulatory Phonology (e.g. Browman & Goldstein 1992). However, as stated by Fromkin (1977: 370), these are *necessary* but not *sufficient* constraints.

We need to decouple our understanding of distinctive feature theory from the assumption that it is directly linked to the phonetic primes out of which it is built, and we need to understand how and why the patterns of phonology are clearly more systematic than would be predicted from principles of phonetic similarity alone.[10] Mielke highlights the key role of "generalization" as a mechanism for learning phonologically active classes. This indeed must be part of the answer and may account for the pervasive role of principles of economy, symmetry and so forth.

The question we have to address is how the various biases in the system (including *channel bias*, defined by Moreton 2008: 83–84 as "phonetically systematic errors in transmission between speaker and hearer caused largely by subtle phonetic interactions which serve as precursors for phonologisation" and *analytic bias* defined as "cognitive biases that facilitate the learning of some phonological patterns and inhibit that of others") together with mechanisms of learning provide the building blocks of the phonological system. Evidence suggests that no simple solution privileging either analytic bias or channel bias to the exclusion of

9. To fully make sense of these results, we need a real theory of naturalness. Some evidence for what this might be is given by Hayes & Steriade (2004). See also Clements (2009) for some discussion of how we might understand markedness.

10. We need to account for the sources of phonological patterns. This is not the same thing as the sources of the explanations. See Cohn (2008) for discussion of this point.

the other, nor privileging either distributional information or "innate properties" to the exclusion of the other will offer an adequate account (see Yang 2004 and Moreton 2008 for discussion).

We pursue these issues in the next section when we turn to the relationship between adult phonology and acquisition.

3. Relationship between adult phonology and acquisition

How does the infant or young child acquire the elements of her phonological system? Chomsky (1965) and Chomsky & Halle (1968) offer one particular answer to this question by defining *language acquisition* as the source of *language universals.* Chomsky (1965: 27) defines the goal of universal grammar and the task of the child acquiring language as follows:

> A theory of linguistic structure that aims for explanatory adequacy incorporates an account of linguistic universals, and *it attributes tacit knowledge of these universals to the child* [….]
>
> Language learning would be impossible unless this were the case [….]
>
> What are the initial assumptions concerning the nature of language that the child brings to language learning, and how detailed and specific is *the innate schema (the general definition of "grammar")* that gradually becomes more explicit and differentiated as the child learns the language? [Emphasis added.]

Many would agree that the nature of phonological universals and the acquisition of phonology are two of the central questions that face the field of phonology. But does Chomsky & Halle's (1968: 4) conclusion that "The significant linguistic universals are those that must be assumed to be available to the child learning a language as an a priori, innate endowment", actually hold?[11]

With over 40 years of additional accumulated knowledge and perspective, this characterization of the child's acquisition task, its linkage to language

11. The rest of the passage defines the poverty of the stimulus argument for phonology. "That there must be a rich system of a priori properties – of essential linguistic universals – is fairly obvious from the following empirical observations. Every normal child acquires an extremely intricate and abstract grammar, the properties of which are much underdetermined by the available data. This takes place with great speed, under conditions that are far from ideal, and there is little significant variation among children who may differ greatly in intelligence and experience. The search for essential linguistic universals is, in effect, the study of the a priori *faculté de langage* that makes language acquisition possible under the given conditions of time and access to data." (Chomsky & Halle 1968: 4). See Sóskuthy (2008) for a recent critique of such arguments in phonology.

universals, and the assumptions underlying this linkage warrant close scrutiny.[12] This inextricable linkage is neither logically nor empirically warranted. The linkage between typological patterns and the acquisition of language by assuming an innate endowment offers an explanation of neither domain (see Gerken 2005). The degree to which the process of acquisition mirrors observed typological patterns is an empirical question. The ways in which these are similar must be better understood, and only then can we characterize appropriate explanations of both typology and acquisition.

3.1 Prerequisites for an understanding of language acquisition

Native speaker-hearers know the inventory, phonotactics, and patterns of alternation in their language, which in turn provide evidence for phonologically active classes. As discussed above, this involves the manipulation of approximately discrete units including segments, syllables, morphemes, words, etc. often insightfully characterized in terms of distinctive features. Speaker-hearers also control much finer phonetic details that play a role in the realization of speech, as well as marking a wide range of indexical properties. Crucially, the adult phonological system is highly practiced and polished (Pouplier 2003). It is generative in the sense that it is productive. Patterns are readily extended to nonce forms, as seen for example in wug testing (Berko 1958, and more recent work, e.g. Albright & Hayes 2003), nativization of recent borrowings (e.g. Kenstowicz & Uffmann 2006), foreign "accents" and so forth. At the same time commonalities in inventories, phonotactics, and patterns of alternation call out for an explanation. We must account for those commonalities often characterized as "universals," but in a way that captures the similarities – rather than strict identity – across languages.

How are these rich complex systems acquired? An increasing body of cross-linguistic research highlights the fact that multiple steps of acquisition need to be accounted for (see Vihman & Velleman 2000; Beckman 2003; Peperkamp 2003; and Munson, Edwards & Beckman to appear, for recent reviews). Critical steps include evidence for early language-independent discrimination (6–8 months), followed by language-specific discrimination (10–12 months, Werker & Tees 1984). These in turn set the stage for more abstract phonological and lexical learning. Small infants are already sensitive to the prosodic structure, the sounds, and the distribution of those sounds in their native language well before the development

12. See Cohn (2010) for a fuller discussion of assumptions in early generative theory and how they frame current approaches.

of a lexicon. In fact early word learning (14–16 months) impedes discrimination of similar sounds (Werker & Stager 2000). A rapid vocabulary explosion (15–18 months) suggests the beginnings of the construction of an adult-like lexicon. Yet the formation of the lexicon continues throughout childhood, and adult-like phonological categories are not fully formed for a number of years (Hazan & Barrett 2000; Menn & Vihman this volume; Munson, Edwards & Beckman to appear). As highlighted by Menn & Vihman, the broad strokes of these developments are seen across individuals and across languages; while they show much variation, they end with very similar results. (This is another way in which the characterization of phonological patterns is approximately, but not literally correct.)

Earlier views that little cognitive development happened in the infant and young toddler reinforced assumptions of a rich innate system as the foundation for language acquisition. However, more recent abundant evidence from language acquisition, as well as cognitive development more generally, demands a fundamental reevaluation of these assumptions. (See references above reviewing phonological development and Gopnik, Meltzoff & Kuhl's 1999 overview of early cognitive development.) The literal characterization of distinctive feature theory, seen in *SPE* as a fixed, universal and innate set, is not accurate. As our understanding of early cognitive development has continued to deepen, the question of what tools we bring to the task (innate linguistic abilities or structure? linguistic or cognitive biases? propensity to categorize? general or specific learning mechanisms?), and how we use these in order to construct a rich functioning phonological system over the first several years of life, are all important issues starting to get the attention they deserve. (See Pierrehumbert, Beckman & Ladd 2000; Vihman & Velleman 2000; and Gerken 2005 for discussion.)

In other words, the *SPE* view that *how* we learn is directly and inextricably linked to *what* we know is not tenable. Vihman & Velleman (2000: 307–309) point out that neither a model of "phonology all the way down" which models early acquisition in terms of adult categories and rules or constraints, nor a "phonetics all the way up" approach assuming that "phonology" emerges gradually out of the phonetics, offers an adequate account of the acquisition of a phonological system. Crucially, both language acquisition and linguistic universals need to be investigated and understood in their own right. Only then can we understand how they are related.

Once the acquisition system has gelled as an adult system, we no longer have direct insight into the steps of development. This highlights the need for a much fuller understanding of the development of both production and perception in preschool and school age children. I now briefly review one such study investigating the acquisition of initial consonants and consonant clusters in English.

3.2 Acquisition of initial consonants in English: A case of covert contrast

Recently, patterns of language acquisition have been modeled as conflicting constraints in Optimality Theory, balancing the demands of markedness and faithfulness. This approach has been applied to the oft-observed patterns of cluster reduction in child language (e.g. Gnanadesikan 1995; Pater & Barlow 2003), whereby impressionistically-observed patterns of deletion, substitutions, coalescence, etc. are modeled as particular markedness constraints outranking faithfulness to the input. Elegant as these accounts are, we need to consider whether the assumed object of study is indeed the appropriate one. Crucially, these analyses assume that the observed patterns involve complete neutralization and that these are phonological patterns akin to those seen in adult phonology. (See also Velleman & Vihman 2002 and Edwards & Beckman 2008 for a critique of the broader claim.)

I consider the phonological development of initial consonants and consonant clusters based on a phonological and phonetic study of a pair of fraternal female twins acquiring American English (Cohn & Kishel 2003). The study illustrates the limitations of the constraint reranking approach and provides a case of acquisition of categories and featural organization in young children. Of particular interest is the fact that based on acoustic data, some impressionistically-observed patterns of neutralization were indeed found to be complete, while others were incomplete. The differences observed between the acoustic patterns and the impressionistic observations (suggesting complete neutralization) offer an example of *covert contrast*, that is, cases where the child is using cues to indicate a contrast, but these cues are either not robust enough or similar enough to the cues used by adults to be identified as such by adults. The presence of such acoustic cues in the case of VOT is well documented (Macken & Barton 1980; Scobbie, Gibbon, Hardcastle & Fletcher 2000), although the evidence for covert contrast in reduced clusters is less clear. (See also recent work by e.g. Edwards, Beckman & Munson 2004; and Munson, Kurtz & Windsor 2005.)

At age 4 years, 1 month, the twins showed marked differences in their phonologies, although both perceived a full range of contrasts. Twin A had achieved a nearly adult phonology, while twin B's s phonological system differed more noticeably from an adult phonology. Twin B was just starting to produce clusters and evidenced a reduced inventory of surface contrasts, apparently due to the multiple effects of substitution, deletion, and coalescence. Most striking in twin B's speech were the range of targets intended by surface [f], [s], [fw], and [sw], as summarized in Table 1. For example, surface [f] could be the outcome for the intended target of /f/, /θ/ (substitution), or /sp/ or /sm/ (coalescence); and surface [s] could be the outcome for the intended target of /s/, /ʃ/ (substitution) or /st/, /sn/, /sk/ (deletion). These patterns of reduction and substitution resulted in apparently homophonous pairs such as [fɪl] for *fill* and *spill*, [fɔl] for *fall* and *small*, [fɪn] for *fin*

and *thin*. At the same time, the relative intelligibility of twin B's speech led to the question of whether complete neutralization was observed, with context serving to disambiguate, or whether a covert contrast was being produced.

Table 1. Surface realization of intended targets, twin B

Surface	Intended target			
	direct mapping	**substitution**	**deletion**	**coalescence**
[f]	f	θ		sp, sm
[s]	s	ʃ	st, sn, sk	
[fw]		fl, fr, θr		spl, spr
[sw]	sw	sl, ʃr	skw, str, skr	

In order to investigate whether low-level phonetic details provided cues to the intended target forms, an acoustic study of a controlled set of data was carried out and compared to careful impressionistic transcription done by the two authors. The study was designed to address the following: (1) Are these cases of true neutralization or are there systematic low-level phonetic differences? (2) If differences occur, in what phonetic dimensions are they found? and (3) Are there phonetic differences between the realization of substitution, coalescence, and deletion? Multiple repetitions of the full range of initial consonants and clusters of English (mostly real words known to the twins at the time) were recorded for both twins. We consider here just a subset of the data expected to exhibit neutralization for twin B, as presented in Table 2. Measurements were made for duration, intensity, and spectral balance. (See Cohn & Kishel 2003 for a fuller discussion of both methodology and results.)

Table 2. Subset of word list used for acoustic study in Cohn and Kishel (2003)

[f]	[aj]	[ɪ]
f	fight	fit
θ	thigh	thick
sp	spy	spit
sm	smile	Smith
[s]	**[aj]**	**[ɪ]**
s	sigh	sit
ʃ	shy	ship
st	sty	stick
sn	snipe	snip
sk	sky	skip

Overall the results showed differences for both duration and intensity between surface [s] and [f] cases, depending on their intended targets. We consider here the results for duration. Duration was measured for the initial (target) consonant(s), following vowel, and postvocalic consonant(s): C|(C)|V|(C). If there were complete neutralization, the prediction is that all surface [f]'s and [s]'s would be of comparable duration to target /f/ and /s/. If there were incomplete neutralization, the prediction is that surface [f] or [s] from target clusters should be longer than from target single consonants, e.g. [s] from target /st/ or /sk/ should be longer than [s] from target /s/, /ʃ/. Average durations for surface [f] and [s] for twin B are shown in Figure 4. (These averages are based on 3–8 tokens per target consonant type. With this small number of tokens, statistical analysis is not possible, and results are suggestive but not conclusive.)

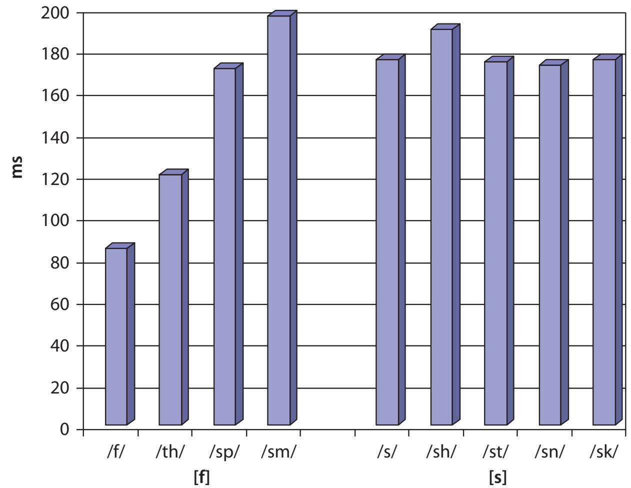

Figure 4. Average durations (in ms.) of surface [f] and surface [s] according to target sound for Twin B (/th/ = /θ/, /sh/ = /ʃ/)

On the average, surface [s] is longer than surface [f] with the mean duration for surface [s] of 177ms and mean duration for surface [f] of 140ms. This is consistent with results reported in the literature for adult fricative production (Jongman, Wayland & Wong 2000), although the differences found here are greater than those usually observed in adults, characteristic of patterns observed in children's productions. For surface [f], both target /sm/ and /sp/ are noticeably longer than target /f/ or /θ/. Thus there is not complete neutralization. These target clusters are comparable to surface [s] durations, but are not as long as surface clusters

(not shown here, with mean duration for surface [sw] of 279ms; mean duration for surface [fw] of 205ms). Target /st, sn, sk/ have durations comparable to target /s/ and /ʃ/, consistent with complete deletion. Thus a difference is seen between single target vs. cluster cases for the surface [f] cases, but not for the surface [s] cases. Interestingly, realization of /sp/ was quite variable, with some targets realized as [f] and others as [fw]. Twin B may have been on the verge of realizing this target sequence as a cluster at the time of the study.

In conclusion, for some of the apparently neutralized cases in twin B's speech, low-level phonetic differences were observable, suggesting that there was covert contrast, as has been observed for developing VOT contrasts. The fact that duration differences were found is interesting, as it is consistent with earlier less-systematic studies that noted such differences (Kornfeld 1971a, b; Menyuk 1972; Menyuk & Klatt 1968). This study suggests that children at this stage realize categories with many adult-like properties, but these cannot be equated with adult categories. We see evidence of grouping of sounds, suggesting that phonological features are playing a role at this stage, since the way different targets pattern together is characterized in terms of place of articulation (e.g. coalescence in the labial cases), and sonority (in terms of the clusters being produced). We turn to the implication of these results in the next section.

These results also highlight a methodological point about the risks of using impressionistic transcriptions to study the speech of infants and young children (a point also made by Menn & Vihman this volume; and Munson, Edwards & Beckman to appear). Since young children's speech shows greater variability than adult speech, if low-level phonetic differences exist, they may be difficult to quantify. Patterns of covert contrast may not be as systematic as overt contrast and may not involve the most obvious phonetic cues. Thus while playing a real role in differentiating target forms, they may be hard to document. Nevertheless, the existence of documented covert contrasts in at least some cases raises serious questions about analysis of child language production based on impressionistic listening alone. Only through systematic phonetic study, with a full range of possible cues investigated, can it be determined how and to what degree adult categories have been realized. A fuller understanding of covert contrast will come from more careful methods for analysis of production data as well as perceptual studies of child production. (See Munson, Edwards & Beckman to appear, for promising work in these directions.)

3.3 Learning of categories and features

This brings us back to the question: What is the relationship between the adult system and acquisition? An increasingly large body of work on the acquisition of phonology points to the conclusion that an adult-like phonological system grows out of extensive exposure to and practice with that phonological system. This is generally true of the development of the perceptual system (as shown recently in work by, e.g. Maye, Weiss & Aslin 2008) and in terms of production (Menn & Vihman this volume). In both of these areas greater variability is seen in child language than adult language. As discussed by Munson, Edwards & Beckman (to appear: 13), "Children's early words are coarse approximations of the adult forms" and the decrease in variation over times goes hand in hand with greater abstraction. Over the course of development, there is evidence first that children learn sound categories which, over time, more closely approximate adult categories; and they also group sounds into phonological classes that mirror those active in the particular adult language. The evidence for gradual development of phonological organization is clear. In Cohn & Kishel's study, we can see that twin A and B were at different stages along this trajectory.

In discussing the acquisition of categories and features, we need to be careful about terminology. We need to distinguish between phones and phonemes and between phonetically defined and phonologically defined classes. If we draw a distinction between distinctive features, that is the organizational elements of phonological structure, and the finer-grained phonetic primes from which they must be learned, then we can understand the increasing body of evidence showing how such learning takes place. We critically need to decouple the linkage assumed by Halle (1983) between abstract phonological features and their neurological basis. Once we do, we can reframe questions about the acquisition of distinctive features in more nuanced and empirically grounded ways. Contributions in this volume, in terms of both acquisition studies and models of learning, are the sorts of work that will lead to serious advances in this regard.

4. Conclusions

Ample evidence demonstrates that there are systematic groupings of sounds in the phonological inventories and alternations of the languages of the world. At the same time, it is also clear that the literal characterization of distinctive feature theory in *SPE*, as a fixed, universal and innate set, is not accurate. Rather, we conclude that this characterization of patterning is approximately, but not literally, correct. The now extensive work on early acquisition, and our increasing understanding of how a quasi-categorical phonological system is built through the process of articulatory and perceptional learning, offers new avenues to answering these questions. It also reminds us that how we learn is not the same as what we know. As amply attested by the contributions to this volume, we are on the brink of a

fundamentally new understanding of the roles of distinctive features in phonetics and phonology. If we take seriously the mechanisms involved in this process, the results become not less, but rather *more* interesting.

A part of the explanation for adult categories as approximately the same across languages is the way that the adult grammar is learned: it is mediated through experience and the child's efforts to become part of a speech community. This suggests that a model approximating adult grammar might be on the right track. Ultimately we need to remember that the sounds of language are central elements in a system of communication. For communication to be successful, all we need are individual grammars that approximate the individual grammars of the people we are communicating with.

References

Albright, Adam & Bruce Hayes. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90. 119–161.

Anderson, Stephen R. 1985. *Phonology in the Twentieth Century: Theories of Rules and Theories of Representations*. Chicago: University of Chicago Press.

Archangeli, Diana B. 1988. Aspects of underspecification theory. *Phonology* 5. 183–207.

Beckman, Mary E. 2003. Input representations (inside the mind and out). In Gina Garding & Mimu Tsujimura (eds), *WCCFL 22 Proceedings*, 101–125. Somerville, MA: Cascadilla Press.

Beddor, Patrice S., James D. Harnsberger & Stephanie Lindemann. 2002. Language-specific patterns of vowel-to-vowel coarticulation: Acoustic structures and their perceptual correlates. *Journal of Phonetics* 30. 591–627.

Berko, Jean. 1958. The child's learning of English morphology. *Word* 14. 150–177.

Blevins, Juliette. 2004. *Evolutionary Phonology: The Emergence of Sound Patterns*. Cambridge: Cambridge University Press.

Bloomfield, Leonard. 1933. *Language*. New York, NY: Henry Holt & co.

Bradlow, Ann R. 1993. *Language-Specific and Universal Aspects of Vowel Production and Perception: A Cross-linguistic Study of Vowel Inventories*. Ph.D. Dissertation, Cornell University.

Bradlow, Ann R. 1995. A comparative acoustic study of English and Spanish vowels. *Journal of the Acoustical Society of America* 97 (3). 1916–1924.

Browman, Catherine P. & Louis Goldstein. 1992. Articulatory Phonology: An overview. *Phonetica* 49. 155–180.

Chen, Matthew. 1970. Vowel length variation as a function of the voicing of the consonant environment. *Phonetica* 22. 129–159.

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge: MIT Press.

Chomsky, Noam & Morris Halle. 1968. *The Sound Pattern of English*. New York: Harper & Row.

Clements, George N. 1985. The geometry of phonological features. *Phonology* 2. 225–252.

Clements, George N. 2003. Feature economy in sound systems. *Phonology* 20.3. 287–333.

Clements, George N. 2009. The role of features in phonological inventories. In Eric Raimy & Charles E. Cairns (eds), *Contemporary Views on Architecture and Representations in Phonology*, 19–68. Cambridge: MIT Press.

Clements, George N. & Elizabeth V. Hume. 1995. The internal organization of speech sounds. In John A. Goldsmith (ed). *The Handbook of Phonological Theory*, 245–306. Oxford: Blackwells.
Cohn, Abigail C. 1990. *Phonetic and Phonological Rules of Nasalization*. Ph.D. dissertation, UCLA. Distributed as *UCLA Working Papers in Phonetics* 76. 1–31.
Cohn, Abigail C. 1998. The phonetics-phonology interface revisited: Where's phonetics? *Texas Linguistic Forum* 41. 25–40.
Cohn, Abigail C. 2008. Phonetics in phonology and phonology in phonetics. *Working Papers of the Cornell Phonetics Laboratory* 16.
Cohn, Abigail C. 2010. Laboratory Phonology: Past Successes and Current Questions, Challenges, and Goals. In Cécile Fougeron, Barbara Kühnert, Mariapaola d'Imperio & Nathalie Vallée (eds), *Laboratory Phonology 10: Variation, Phonetic Detail and Phonological Representation*, 3–29. Berlin: Mouton de Gruyter.
Cohn, Abigail C. & Emily K. Kishel. 2003. Development of Initial Clusters in American English by Fraternal Twins: An Acoustic Study. *Proceedings of the 15th International Congress of Phonetic Sciences*, 1991–1994.
Dresher, Elan B. 2008. The Contrastive Hierarchy in Phonology. In Peter Avery, B. Elan Dresher & Keren Rice (eds), *Contrast in Phonology: Theory, Perception, Acquisition*, 11–33. Berlin: Mouton de Gruyter.
Edwards, Jan & Mary E. Beckman. 2008. Some cross-linguistic evidence for modulation of implicational universals by language-specific frequency effects in phonological development. *Language Learning and Development* 4. 122–156.
Edwards, Jan, Mary E. Beckman & Benjamin Munson. 2004. The interaction between vocabulary size and phonotactic probability effects on children's production accuracy and fluency in nonword repetition. *Journal of Speech, Language, and Hearing Research* 47. 421–436.
Fromkin, Victoria A. 1977. Some questions regarding universal phonetics and phonetic representations. In Alphonse Juilland (ed.). *Linguistic Studies offered to Joseph Greenberg on the Occasion of his Sixtieth Birthday*, 365–380. Saratoga, CA: Anma Libri.
Gerken, LouAnn. 2005. What develops in language development? In Robert Kail (ed) *Advances in Child Development and Behavior*, Vol. 33, 153–192. San Diego: Elsevier.
Gildea, Daniel & Daniel Jurafsky. 1996. Learning bias and phonological rule induction. *Computational Linguistics* 22. 497–530.
Gnanadesikan, Amelia E. 1995. Markedness and Faithfulness Constraints in Child Phonology. *Rutgers Optimality Archive,* 67. http://roa.rutgers.edu.
Goldsmith, John A. 1976. An overview of autosegmental phonology. *Linguistic Analysis* 2. 23–68.
Goldstein, Louis. 2003. Emergence of discrete gestures. *Proceedings of the 15th International Congress of Phonetic Sciences*, 85–88.
Goldstein, Louis & Carol A. Fowler. 2003. Articulatory phonology: A phonology for public language use. In Schiller, Niels O. & Meyer, Antji S. (eds), *Phonetics and Phonology in Language Comprehension and Production,* 159–207. Berlin: Mouton de Gruyter.
Goldstein, Louis, Dani Byrd & Elliot Saltzman. 2006. The role of vocal tract gestural action units in understanding the evolution of phonology. In Michael A. Arbib (ed.). *Action to Language: The Mirror Neuron System*, 215–249 Cambridge: Cambridge University Press.
Gopnik, Alison, Andrew N. Meltzoff & Patricia K. Kuhl. 1999. *The Scientist in the Crib: What Early Learning Tells Us about the Mind*. New York: William Morrow & Co.
Gordon, Matthew Kelly. 2006. *Syllable Weight: Phonetics, Phonology, Typology*. New York: Routledge.

Hale, Mark & Charles Reiss. 2000. Phonology as cognition. In Noel Burton-Roberts, Philip Carr & Gerard Docherty (eds), *Phonological Knowledge: Conceptual and Empirical Issues*, 161–184. New York: Oxford University Press.
Halle, Morris. 1983. On distinctive features and their articulatory implementations. *Natural Language and Linguistic Theory* 1. 91–105.
Halle, Morris & George N. Clements. 1983. *Problem Book in Phonology*. Cambridge, MA: MIT Press.
Hayes, Bruce. 1999. Phonetically-driven phonology: The role of optimality theory and inductive grounding. In Michael Darnell, Edith Moravscik, Fredrick Newmeyer, Michael Noonan & Kathleen Wheatly (eds), *Functionalism and Formalism in Linguistics, Volume I: General Papers*, 243–285. Amsterdam: John Benjamins.
Hayes, Bruce & Donca Steriade. 2004. Introduction: the phonetic bases of phonological markedness. In Bruce Hayes, Robert Kirchner & Donca Steriade (eds), *Phonetically Based Phonology*, 1–33. Cambridge: CUP.
Hayes, Bruce & Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39. 379–440.
Hayes, Bruce, Kie Zuraw, Péter Siptar & Zsuzsa Londe. 2009. Natural and unnatural constraints in Hungarian vowel harmony. *Language* 85.4. 822–863.
Hazan, Valerie & Sarah Barrett. 2000. The development of phonemic categorization in children aged 6–12. *Journal of Phonetics* 28. 377–396.
Jakobson, Roman, Gunar M. Fant & Morris Halle. 1952. *Preliminaries to Speech Analysis: The Distinctive Features and their Correlates*. Cambridge: MIT Press.
Jongman, Allard, Ratree Wayland & Serena Wong. 2000. Acoustic characteristics of English fricatives. *JASA* 108. 1252–1263.
Joos, Martin. 1957. *Readings in Linguistics, volume I*. Washington: American Council of Learned Societies.
Keating, Patricia A. 1985. Universal phonetics and the organization of grammars. In Victoria A. Fromkin (ed.), *Phonetic Linguistic Essays in Honor of Peter Ladefoged*, 115–132. Orlando: Academic Press.
Kenstowicz, Michael J. 1994. *Phonology in Generative Grammar*. Oxford: Blackwell.
Kenstowicz, Michael J. & Christian Uffmann (eds), 2006. *Loan Phonology: Current Issues. Lingua 116.7.*
Kingston, John & Randy C. Diehl. 1994. Phonetic knowledge. *Language* 70. 419–454.
Kornfeld, J.R. 1971a. Theoretical issues in child phonology. *CLS* vol. 7, 454–69.
Kornfeld, J.R. 1971b. What initial clusters tell us about a child's speech code. *Quarterly Progress Reports, MIT RLE* vol. 101, 218–221.
Ladefoged, Peter. 1980. What are linguistic sounds made of? *Language* 56. 485–502.
Lindau, Mona & Peter Ladefoged. 1986. Variability in feature specifications. In Joseph S. Perkell & Dennis H. Klatt (eds), *Invariance and Variability in Speech Processes*. Hillsdale, NJ. Lawrence Erlbaum.
Lindblom, Björn 1990. On the notion of possible speech sound. *Journal of Phonetics* 18. 135–152.
Macken, Marlys A. & David Barton. 1980. The acquisition of voicing contrasts in English: A study of voice onset time in word-initial stop consonants. *Journal of Child Language* 7. 41–74.
Maye, Jessica, Daniel Weiss & Richard Aslin. 2008. Statistical phonetic learning in infants: facilitation and feature generalization. *Developmental Science* 11. 122–134.

Menn, Lise & Marilyn Vihman, this volume. Features in child phonology: inherent, emergent, or artifacts of analysis? George N. Clements & Rachid Ridouane (eds), *Where Do Features Come From? The Nature and Sources of Phonological Primitives.* Amsterdam: John Benjamins.

Menyuk, Paula. 1972. Clusters as single underlying consonants: Evidence from children's pronunciations. In *Proceedings of the Seventh International Congress of Phonetic Sciences,* 1161–1165.

Menyuk, Paula & Mary Klatt. 1968. Child's production of initial consonant clusters. *Quarterly Progress Reports, MIT RLE* vol. 91, 205–213.

Mielke, Jeff. 2005. Modeling Distinctive Feature Emergence. *Proceedings of WCCLF* 24 281–289.

Mielke, Jeff. 2008. *The Emergence of Distinctive Features*. Oxford: Oxford University Press.

Mielke, Jeff, Lyra Magloughlin & Elizabeth Hume. to appear. A comparison of six distinctive feature systems. In J. Goldsmith, E. Hume & L. Wetzels (eds.) Tones and Features: A Symposium in Honor of G. Nick Clements. Berlin: Mouton de Gruyter.

Moreton, Elliott. 2008. Analytic bias and phonological typology**.** *Phonology* 25. 83–127.

Munson, Benjamin, Beth A. Kurtz & Jennifer Windsor. 2005. The influence of vocabulary size, phonotactic probability, and wordlikeness on nonword repetitions of children with and without language impairments. *Journal of Speech, Language, and Hearing Research* 48 (5). 1033–1047.

Munson, Benjamin, Jan Edwards & Mary E. Beckman. to appear. Phonological representations in language acquisition: Climbing the ladder of abstraction. To appear in Abigail C. Cohn, Cécile Fougeron & Marie Huffman (eds), *The Oxford Handbook of Laboratory Phonology*. Oxford: Oxford University Press.

Pater, Joe & Jessica A. Barlow. 2003. Constraint conflict in cluster reduction. *Journal of Child Language* 30. 487–526.

Peperkamp, Sharon. 2003. Phonological acquisition: Recent attainments and new challenges. *Language and Speech* 46. 87–113.

Peperkamp, Sharon & Emmannuel Dupoux. 2007. Learning the mapping from surface to underlying representations in an artificial language. In Jennifer Cole & José Ignacio Hualde (eds), *Laboratory Phonology 9* 315–338. Berlin: Mouton de Gruyter.

Pierrehumbert, Janet B. 1980. *The Phonology and Phonetics of English Intonation*. Ph.D. Dissertation, MIT.

Pierrehumbert, Janet B. & Mary E. Beckman. 1988. *Japanese Tone Structure*. Cambridge: MIT Press.

Pierrehumbert, Janet, Mary E. Beckman & D. Robert Ladd 2000. Conceptual foundations in phonology as a laboratory science. In Noel Burton-Roberts, Philip Carr & Gerard Docherty (eds), *Phonological Knowledge: Conceptual and Empirical Issues*, 273–304. New York: Oxford University Press.

Port, Robert F. & Adam P. Leary. 2005. Against formal phonology. *Language* 81. 927–964.

Pouplier, Marianne. 2003. *Units of phonological encoding: Empirical evidence*. Ph.D. Dissertation, Yale University.

Prince, Alan & Paul Smolensky. 2004. *Optimality Theory: Constraint Interaction in Generative Grammar*. Oxford: Blackwell.

Pulleyblank, Douglas. 2006. Minimizing UG: Constraints upon Constraints. *Proceedings of the 25th West Coast Conference on Formal Linguistics* 15–39. Somerville, MA: Cascadilla Proceedings Project.

Sapir, Edward. 1963. The psychological reality of the phoneme. In David G. Mandelbaum (ed.). *Selected Writings of Edward Sapir*, 46–60. Berkeley: UC Berkeley Press.

Scobbie, James M., Fiona Gibbon, William J. Hardcastle & Paul Fletcher. 2000. Covert contrast as a stage in the acquisition of phonetics and phonology, in Michael B. Broe & Janet B. Pierrehumbert (eds), *Papers in Laboratory Phonology V: Acquisition and the Lexicon*, 194–207. CUP: Cambridge.

Seidl, Amanda & Eugene Buckley. 2005. On the learning of arbitrary phonological rules. *Language Learning and Development* 1 (3&4). 289–316.

Sóskuthy, Marton 2008. Is there a poverty of the stimulus argument. Paper presented at 16th Manchester Phonology Meeting.

Stevens, Kenneth N. 1972. The quantal nature of speech: Evidence from articulatory-acoustic data. In Edward E. David Jr. & Peter B. Denes (eds), *Human Communication: A Unified View*, 51–66. New York: McCraw-Hill.

Stevens, Kenneth N. 1989. On the quantal nature of speech. *Journal of Phonetics* 17. 3–45.

Trubetzkoy, Nikolai S. 1939. *Grundzüge der Phonologie*. Publié avec l'appui du Cercle Linguistique de Copenhague et du Ministère de l'instruction publique de la République Tchécoslovaque. Prague.

Velleman, Shelley & Marilyn Vihman. 2002. The optimal initial state. *Rutgers Optimality Archive*, 522, http://roa.rutgers.edu.

Vihman, Marilyn & Shelley Velleman. 2000. Phonetics and the origins of phonology. In Noel Burton-Roberts, Philip Carr & Gerard Docherty (eds), *Phonological Knowledge: Conceptual and Empirical Issues*, 305–339. New York: Oxford University Press.

Werker, Janet F. & Christine L. Stager. 2000. Developmental changes in infant speech perception and early word learning. In Michael Broe & Janet B. Pierrehumbert (eds), *Papers in Laboratory Phonology V: Acquisition and the Lexicon*, 181–193. Cambridge: Cambridge University Press.

Werker, Janet F. & Richard Tees. 1984. Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development* 7. 49–63.

Yang, Charles D. 2004. Universal Grammar, statistics or both? *Trends in Cognitive Science* 8, no. 10. 451–456.